

Toward Real-World High-Precision Image Matting and Segmentation

Haipeng Zhou¹, Zhaohu Xing¹, Hongqiu Wang¹, Jun Ma^{1,2}, Ping Li³, Lei Zhu^{1,2*}

¹The Hong Kong University of Science and Technology (Guangzhou)

²The Hong Kong University of Science and Technology

³The Hong Kong Polytechnic University

Abstract

High-precision scene parsing tasks, including image matting and dichotomous segmentation, aim to accurately predict masks with extremely fine details (such as hair). Most existing methods focus on salient, single foreground objects. While interactive methods allow for target adjustment, their class-agnostic design restricts generalization across different categories. Furthermore, the scarcity of high-quality annotation has led to a reliance on inharmonious synthetic data, resulting in poor generalization to real-world scenarios. To this end, we propose a Foreground Consistent Learning model, dubbed as FCLM, to address the aforementioned issues. Specifically, we first introduce a Depth-Aware Distillation strategy where we transfer the depth-related knowledge for better foreground representation. Considering the data dilemma, we term the processing of synthetic data as domain adaptation problem where we propose a domain-invariant learning strategy to focus on foreground learning. To support interactive prediction, we contribute an Object-Oriented Decoder that can receive both visual and language prompts to predict the referring target. Experimental results show that our method quantitatively and qualitatively outperforms *state-of-the-art* methods.

Code — <https://github.com/haipengzhou856/FCLM>

Introduction

High-accuracy image matting and segmentation aim to predict fine-grained masks, which are essential for vision tasks such as AR/VR, image editing, and so on (Goferman, Zelnik-Manor, and Tal 2011; Liu, Tran, and Liu 2021). Compared to conventional dense prediction, this task requires significantly finer details, making accurate boundary localization and semantic consistency challenging.

Data matters, yet real-world fine-grained annotations are extremely costly. Due to the difficulty of labeling, a large portion of existing matting datasets (Li, Zhang, and Tao 2023; Xu et al. 2017; Hou and Liu 2019) are synthetic and inharmonious. Although synthetic data can enhance the training effect (Qian et al. 2024), this is based on the premise that the distribution of the images is normal, rather than simply stitching the foreground with different backgrounds to

Methods	SAD↓	MSE↓	Grad↓	Conn↓
P3M-Net (Li et al. 2021)	8.73	0.0027	13.83	9.14
P3M-Net† (Li et al. 2021)	12.33	0.0051	16.44	11.67
MatteFormer (Park et al. 2022)	7.65	0.0020	12.34	7.89
MatteFormer† (Park et al. 2022)	10.84	0.0064	15.67	10.44
MODNet (Ke et al. 2022)	10.33	0.0057	15.35	14.35
MODNet† (Ke et al. 2022)	23.41	0.0154	26.38	18.95

Table 1: Comparison of results obtained by training on synthetic (†) and real data on P3M-500-P testing set.

create jarring images. The reliance on inharmonious synthetic data in image matting introduces visual inconsistencies and domain gaps, undermining model generalization. To demonstrate this problem, we reproduce cutting-edge methods on the P3M dataset (Li et al. 2021), replacing its original backgrounds with BG20K images (Li et al. 2021) to create inharmonious synthetic training data. As shown in Tab. 1, models trained on synthetic data exhibit a significant performance drop when evaluated on real-world images. This discrepancy highlights the limitations of inharmonious data for training. However, considering the difficulty in annotation, using such inharmonious data in matting scenarios is a data augmentation measure adopted out of necessity. Our findings underscore the critical need for techniques to properly utilize synthetic data and unleash their power.

Furthermore, most existing methods are designed to be object-specific (*e.g.*, portrait-focused) or salient-object oriented. The former (Li et al. 2021; Ke et al. 2022; Li et al. 2022b) often suffer from poor generalization, as they are trained and optimized for specific categories and fail on unseen object types. The latter approaches (Yu et al. 2024b; Qin et al. 2022; Hu et al. 2023) are category-agnostic, and they struggle to identify which object to segment when multiple salient regions coexist. Neither paradigm supports flexible user interaction or multi-instance segmentation, making them unsuitable for open-set and real-world scenarios. Multi-stage methods (Yao et al. 2024b; Li, Jain, and Shi 2024; Ye et al. 2024; Sun, Tang, and Tai 2022) can alleviate this problem, where an initial coarse mask is first generated and a refine-net is used to polish the prediction. For example, MAM (Li, Jain, and Shi 2024) integrates GroundDINO (Liu et al. 2024b) to generate semantic bounding boxes, and on top of them, SAM (Kirillov et al. 2023) can

*Corresponding author: Lei Zhu (leizhu@hkust-gz.edu.cn).
Copyright © 2026, Association for the Advancement of Artificial Intelligence (www.aaai.org). All rights reserved.

produce the initial mask. MAM then proceeds with a refinement stage to improve mask quality. Obviously, such a triple-stage approach is inherently complex. Moreover, the reliance on multiple independent modules could lead to error propagation and reduce overall robustness.

To address the aforementioned challenges, this paper introduces a Foreground Consistent Learning Model (FCLM) for high-accuracy image matting and segmentation. During training, the network takes as input a pair of images that share the same foreground, enabling the model to learn consistent representations across variations in background and context. To enhance spatial reasoning and improve foreground-background separation, we transfer depth-aware priors through a knowledge distillation strategy. Leveraging foreground consistency, we further introduce an adversarial learning loss together with an optimal transport loss to learn domain-invariant features. Additionally, to support flexible user interactions, we propose a prompt encoder that is integrated into the decoding process, enabling object-oriented and semantics-aware predictions.

In sum, our contributions are four-fold:

- We term the inharmonious synthesized images as a domain adaptation problem for better real-world generalization, in which we propose Foreground Consistent Domain Adaptation to encourage the network to focus on the domain invariant features.
- We introduce a Depth-Aware Distillation strategy for knowledge transfer, enabling our framework to separate the features between foreground and background.
- An Object-Oriented Decoder is deployed to support various prompts including visual and text guidance, making our model interactive and semantic-aware.
- Our overall framework, FCLM, achieves *state-of-the-art* performance on various public datasets, verifying the effectiveness of our approach.

Related Works

Dense Prediction

Dense prediction requires pixel-level understanding, such as semantic segmentation and depth estimation. Since deep learning era (He et al. 2016), a series of works based on CNN (Shelhamer, Long, and Darrell 2016; Ronneberger, Fischer, and Brox 2015; Fan et al. 2023; Liu, Jiang, and Ding 2024), Transformer (Xie et al. 2021; Wang et al. 2024; Ranftl, Bochkovskiy, and Koltun 2021; Jain et al. 2023; Zhou et al. 2024), Mamba (Liu et al. 2024c; Xing et al. 2024; Wu et al. 2024), or their variants (Chen et al. 2021) have emerged. To achieve higher precision in dense prediction, numerous strategies have been developed to enhance the feature representation, like multiscale approaches (Xiao et al. 2018; Zhao et al. 2017), patch tiling (Li, Bhat, and Wonka 2024), and cascade structure (Yuan et al. 2021). Recently, vision foundation model like SAM (Kirillov et al. 2023; Ravi et al. 2024) and Depth-Anything (Yang et al. 2024a,b) have shown more powerful performance. Despite the significant advancements, the existing general approaches still exhibit a substantial performance disparity when it comes to the editable level, which demands high-precision parsing.




Figure 1: Overview of the framework of our proposed FCLM. During training, we input two images sharing the same foreground. First, we perform Depth-Aware Distillation to transfer depth-related knowledge from the teacher model to the student model. Then, we apply Foreground Consistent Domain Adaptation to enhance the generalization ability of the model. Our prompt encoder supports various types of prompts, enabling multi-instance and semantic-aware high-precision prediction.

High-Precision Image Parsing

High-precision tasks such as image matting and dichotomous segmentation are significantly more challenging which require accurate estimation of fine-grained structures (e.g., hair). Previous image matting approaches (Xu et al. 2017; Hou and Liu 2019; Li and Lu 2020; Liu et al. 2021; Park et al. 2022; Yao et al. 2024a) typically rely on trimap guidance to predict the alpha matte. However, trimap generation often involves costly manual annotation. Although fully automatic networks (Li et al. 2021; Ke et al. 2022; Li et al. 2022b; Qin et al. 2022; Yu et al. 2024b; Lin et al. 2021; Sun, Tang, and Tai 2022) enable end-to-end inference without user input, their class-agnostic nature tends to produce false positive predictions. Moreover, such architectures could bias the model towards learning salient or isolated single foreground objects, thereby limiting generalization to a broader range of objects. Recently, approaches (Li, Jain, and Shi 2024; Yao et al. 2024b; Ke et al. 2023; Liu et al. 2024a) built upon SAM leverage visual prompts to enable more intuitive user interaction, where the coarse masks generated by SAM are further refined. Nevertheless, these methods heavily rely on the quality of the initial mask, and SAM itself lacks sufficient semantic understanding. Although CLIP-Mat (Li, Zhang, and Tao 2023) introduces referring image matting with textual input, it directly finetunes CLIP for the matting, inevitably suffering from catastrophic forgetting. Moreover, due to CLIP’s inherent lack of pixel-level understanding (Zhang et al. 2024), such a straightforward way struggles to achieve satisfactory performance. In contrast, our method can support both visual and textual prompts to predict the specified target accurately, instead of a single object and class-agnostic.

Data Reliability

Existing real-world datasets focusing on high-precision matting (Li et al. 2021, 2022b; Liu et al. 2021; Yu et al. 2021) or accurate dichotomous segmentation (Qin et al. 2022) only provide annotations for single objects. Although




Figure 2: Illustration of the Depth-Aware Distillation. We deploy depth map as guidance to drive foreground and background distillation.

HIM2K (Sun, Tang, and Tai 2022) contains multiple instances, it is specifically centered on human subjects. A more common design in such datasets is to annotate the foreground object and then composite it with various background images. For example, RefMatte (Li, Zhang, and Tao 2023) introduces a multiple-object matting dataset by attaching diverse foreground instances. However, this simple synthesis-based composition inevitably introduces a domain gap due to the visual inconsistency between foregrounds and backgrounds, leading to suboptimal performance. To address this limitation, we formulate it as a domain adaptation problem, aiming to better generalize from existing datasets.

Methodology

Overview

As depicted in Fig. 1, we input the synthesized image A and vanilla (real or synthesized) image B into our pipeline. We first introduce a Depth-Aware Distillation, in which we distill the robustness prior of depth feature to the student model. Next, we apply Foreground Consistent Domain Adaptation to encourage the network to focus on the domain-invariant (*i.e.*, foreground) learning. Finally, our decoder can support various prompts, including visual and text guidance, achieving a target-oriented prediction. Note that during the inference stage, only the student model and decoder will be used.

Depth-Aware Distillation

In our task, we seek to separate the foreground from the background, which is strongly related to depth estimation. Motivated by this, we propose a Depth-Aware Distillation (DAD) strategy to transfer the depth-related knowledge into a smaller student model. In our practice, we deploy Depth-Anything V2 (Yang et al. 2024b) as the teacher model while the student model is DINOv2 (Oquab et al. 2023).

Given the synthesis image A and vanilla image B , the student and teacher model will yield encoded features F_S^A , F_S^B , F_T^A , and F_T^B . Instead of aligning the knowledge at the output level which requires auxiliary student heads to predict depth, raising additional computation, we conduct the distil-

lation at the feature level. The standard distillation loss (Hinton, Vinyals, and Dean 2015) can be illustrated as:

$$\mathcal{L}_{\text{kd}} = D(F_S^A, \tau(F_T^A)) + D(F_S^B, \tau(F_T^B)), \quad (1)$$

where D is the distance function measuring the discrepancy of the features, and τ is a project function to align the dimension. However, this simple constraint can introduce feature noises due to the capacity gap (Huang et al. 2022) between teacher and student.

To better transfer valuable information, we utilize the depth map generated by the teacher model as a guide. As shown in Fig. 2, inspired by the ViT-Register (Darcet et al. 2023) we introduce two meta-nets initialized with a context token ϵ , which are used to project the teacher features into foreground and background representations, respectively. The additional token helps capture global context and reduces artifacts in the feature maps. As a result, the projected feature is produced by $F'_T = \tau(F_T + \epsilon)$, where the dimension of F'_T is aligned with the student feature F_S . To distinguish between foreground and background, we apply an empirical threshold δ to the depth map, generating corresponding weights d_+ and d_- , respectively. Specifically, we set:

$$d_+(x, y) = \begin{cases} \frac{\text{Depth}(x, y)}{\max(\text{Depth})}, & \text{if } \text{Depth}(x, y) > \delta, \\ 0, & \text{otherwise,} \end{cases} \quad (2)$$

$$d_-(x, y) = \begin{cases} \frac{\delta - \text{Depth}(x, y)}{\delta}, & \text{if } \text{Depth}(x, y) \leq \delta, \\ 0, & \text{otherwise.} \end{cases}$$

By doing so, we achieve more precise and targeted distillation of relevant features. As a result, the depth-aware distillation for image A can be written as:

$$\mathcal{L}_{\text{kd}}^A = D(d_+^A \times F_S^A, d_+^A \times \tau_+(F_T^A + \epsilon_+)) + D(d_-^A \times F_S^A, d_-^A \times \tau_-(F_T^A + \epsilon_-)). \quad (3)$$

Similarly, we can deploy this formulation for image B to conduct distillation as well.

Foreground Consistent Domain Adaptation

In this section, we perform consistent learning on features F_S^A and F_S^B , which share the same foreground. Inspired by domain adaptation methods that encourage models to focus on domain-invariant components, we refer to our approach as foreground consistent learning, emphasizing its concentration on preserving consistency in the shared foreground across different background images.

Following the conventional adversarial learning (Ganin et al. 2016), we assign domain labels to input image pairs, *i.e.*, 0 for A and 1 for B to denote their domains. A domain discriminator $h(\cdot)$ is deployed to classify the features F_S^A and F_S^B into these two domains. Notably, a Gradient Reversal Layer (GRL) (Ganin and Lempitsky 2015) is inserted between the student encoder \mathcal{E}_S and the discriminator to enforce domain-invariant feature learning. During backpropagation, the GRL will reverse the gradient from $h(\cdot)$ via multiplying the negative factor $-\lambda$. This encourages the student

encoder to produce features that confuse the discriminator, thereby reducing the discrepancy between A and B . Formally, the adversarial loss for the domain discriminator can be expressed as:

$$\mathcal{L}_{\text{adv}} = \min_{\mathcal{E}_S} \max_h \left(\mathbb{E}_A [\log h(\mathcal{E}_S(x^A))] + \mathbb{E}_B [\log (1 - h(\mathcal{E}_S(x^B)))] \right). \quad (4)$$

However, since image B also could be synthesized in a way that fails to confuse the discriminator during adversarial training, it may lead to the collapse of the training process. To mitigate this issue, we introduce a simple yet effective token exchange strategy. Specifically, we propose swapping tokens at identical indices between F_S^A and F_S^B . This design offers several intuitive benefits: (1) Exchanged foreground tokens retain contextual coherence due to the presence of the same foreground objects. (2) Exchanging background tokens introduces implicit domain adaptation perturbation signals, helping to suppress domain-specific biases. (3) Shared masks ensure consistent decision boundaries for transitional regions. In practice, we randomly swapping the visual tokens with a ratio of 25%.

For better aligning the domain-invariant distribution, we continue to optimize the learning of foreground consistent features. Firstly, we deploy the ground truth \mathcal{M} to filter the valid foreground tokens by:

$$F_{fg} = \{F_S[i] \mid \mathcal{M}_{\text{patch}}[i] > 0\}, \quad (5)$$

where the `patch` denotes we downsample to patch-level resolution for the mask grid. As a result, we have the foreground tokens $F_{fg}^A = \{f_1^A, \dots, f_K^A\}$ and $F_{fg}^B = \{f_1^B, \dots, f_K^B\}$, and K is the number of foreground tokens. Thus, we have two empirical domain distributions:

$$u_A = \sum_{i=1}^K \frac{1}{K} \delta_{f_i^A}, \quad u_B = \sum_{j=1}^K \frac{1}{K} \delta_{f_j^B}, \quad (6)$$

where δ is a Dirac delta function centered at the domain feature f_i^A or f_j^B . To align the distribution, we devise an Optimal Transport (OT) loss to optimize it. As a Kantorovich problem (Kantorovich 2006), OT seeks to find a transport plan π that minimizes the total cost \mathcal{C} of moving "mass" from different distributions, which can be written as:

$$\begin{aligned} OT(u_A, u_B) &\triangleq \min_{\pi \in \Pi(u_A, u_B)} \langle \pi, \mathcal{C} \rangle_F, \\ \text{s.t. } \Pi(u_A, u_B) &= \left\{ \pi \in \mathbb{R}_+^{K \times K} \mid \pi \mathbf{1}_K = u_A, \right. \\ &\quad \left. \pi^\top \mathbf{1}_K = u_B \right\} \end{aligned} \quad (7)$$

where $\langle \cdot, \cdot \rangle_F$ is the Frobenius inner product, equivalent to $\sum_{i,j} \pi[i, j] \mathcal{C}[i, j]$, and $\Pi(u_A, u_B)$ is the set of all transport plans π whose row and column sums match the distributions u_A and u_B , respectively. For cost matrix \mathcal{C} , to encourage feature alignment in terms of orientation in the latent space (Cheng, Tai, and Tang 2021), we use cosine dissimilarity to measure the distance. As a result, we can deploy Sinkhorn-nopp scaling algorithm (Chizat et al. 2018;




Figure 3: Illustration of our Object-Oriented Decoder. It supports various prompt inputs to orient the object, achieving interactive high-precision matting or segmentation.

Frogner et al. 2015) to minimize the OT loss:

$$\begin{aligned} \mathcal{L}_{OT} &= \min_{\pi} \sum_{i,j} \pi[i, j] \mathcal{C}[i, j] \\ &= \frac{1}{K} \times \left(1 - \frac{F_{fg}^A[i] \cdot F_{fg}^B[j]}{\|F_{fg}^A[i]\|_2 \|F_{fg}^B[j]\|_2} \right). \end{aligned} \quad (8)$$

Regarding the background, we do not perform alignment as it has limited contribution to the final prediction.

Object-Oriented Decoder

Here, we develop an Object-Oriented Decoder for supporting various prompts. Unlike the complex pipelines (Sun, Tang, and Tai 2022; Yao et al. 2024b; Li, Jain, and Shi 2024) which are based on SAM (Kirillov et al. 2023) or Mask R-CNN (He et al. 2017) for coarse mask prediction and subsequent refinement, we adopt the elegant design of Smart-Matting (Ye et al. 2024) to directly and explicitly prompt the model through visual guidance in an end-to-end manner. Moreover, our framework additionally supports language-based referring settings.

To deal with the visual prompt, we embed the coordinates using positional encoding to obtain a sparse embedding. Notably, if no visual prompt is provided, we initialize the corresponding embedding with zeros. For the language input, we utilize CLIP’s text encoder (Radford et al. 2021) to generate the text embedding. Previous works (Li et al. 2022a; Lin et al. 2023) typically compute similarity between pixel features and the CLIP text embedding for prediction. However, this paradigm is not well-suited to our task, as CLIP lacks fine-grained understanding of image details. Instead, we treat the text embedding as prompt guidance. Furthermore, we introduce an additional learnable query that serves as an implicit context prompt (Zhou et al. 2022). As illustrated in Fig. 3, following Mask2Former and SAM (Cheng et al. 2022; Kirillov et al. 2023), we employ lightweight pixel decoder and transformer decoder for mask prediction. The concatenated queries enable interaction between the prompts and hierarchical image features. As a result, our pipeline effectively facilitates both prompt integration and accurate prediction.

Optimization Objective

Here, we define the overall objective loss used to optimize our network. During knowledge distillation, we employ Kullback-Leibler (KL) divergence to measure the fea-

CONIFGS	HIM2K-NATURAL				REFMATTE-RW			
Methods	IMQ _{MSE} ↑	IMQ _{MAD} ↑	IMQ _{Grad} ↑	IMQ _{Conn} ↑	MSE↓	SAD↓	Grad↓	Conn↓
InstMatt (Sun, Tang, and Tai 2022)	81.34	70.26	<u>74.90</u>	72.60	/	/	/	/
CLIPMat (Li, Zhang, and Tao 2023)	/	/	/	/	0.0474	85.83	/	/
MatAny (Yao et al. 2024b)	75.68	62.67	46.78	68.55	0.0270	52.91	25.17	5.13
MAM (Li, Jain, and Shi 2024)	81.67	68.78	51.79	72.62	0.0151	29.23	25.85	2.83
SmartMatting (Ye et al. 2024)	<u>82.73</u>	<u>71.44</u>	68.58	<u>74.33</u>	<u>0.0120</u>	<u>25.60</u>	<u>22.62</u>	5.31
Ours	83.48	73.45	75.69	74.83	0.010	21.31	19.38	3.33

Table 2: Quantitative comparisons on real-world multi-object matting datasets. The best ones are highlighted with **bold** and runner-ups are underlined.




Figure 4: Visual comparison with on HIM2K and RefMatte real-world multi-objects datasets. Note that ours is guided by text, while others are used by visual prompt (box).

ture distance. In the domain alignment process, we incorporate Eq. 4 and Eq. 8 to constrain foreground learning. Considering the differences between tasks, we apply distinct prediction losses $\mathcal{L}_{\text{head}}$ for matting and segmentation. For the matting task, we follow ViTMatte (Yao et al. 2024a) and adopt the l_1 loss together with the Laplacian loss. In contrast, for dichotomous segmentation, we employ the commonly used combination of BCE loss and IoU loss (Yu et al. 2024b). The total term of our optimization objective is:

$$\mathcal{L} = \mathcal{L}_{\text{kd}} + \mathcal{L}_{\text{adv}} + \mathcal{L}_{\text{OT}} + \mathcal{L}_{\text{head}}. \quad (9)$$

Experiments

Setups

Datasets and Metrics. We evaluate our method on *real-world* datasets. Only HIM2K (Sun, Tang, and Tai 2022) and RefMatte (Li, Zhang, and Tao 2023) support *multiple object instances* matting configuration, and we use them to assess the referring matting capability of our approach. For other datasets, we choose DIS-5K (Qin et al. 2022) as a representative for comparison here because it has a sufficient number of samples and its scenes are more complex and challenging. Results on other datasets can be found in our supplementary material due to space constraints. Since our pipeline requires paired images (sharing the same foreground) as inputs, we use the images from BG20K (Li et al. 2022b) to replace the background yielding synthesized training images. Notably,

we use the text prompt template "a photo of {CLS}." for the referring setting, except for RefMatte, which has provided text annotations.

Following the evaluation protocol established in HIM2K, we adopt the Instance Matting Quality (IMQ) metric to comprehensively assess matting performance. For RefMatte, we report results using a set of widely metrics, including Sum of Absolute Differences (SAD), Mean Squared Error (MSE), Gradient Distortion (Grad), and Connectivity Loss (Conn), which collectively provide insights into both global and local quality of the predicted alpha mattes. In the case of dichotomous segmentation, we follow the evaluation protocol proposed in (Qin et al. 2022), and employ several commonly used criteria such as the maximum F-measure ($\max F_{\beta}$), weighted F-measure (F_{β}^w), mean absolute error (M), and structural similarity index (S_{α}).

Implement details. All the experiments are conducted via 4×Nvidia 4090 GPUs. We utilize the AdamW optimizer with a constant learning rate of 1e-5 to train the model. We deploy Depth-Anything V2 (Yang et al. 2024b) Large as teacher model. For a fair comparison, DINOv2 (Oquab et al. 2023) Base and Small are as student models for dichotomous segmentation and image matting, respectively. CLIP-Base (Radford et al. 2021) is deployed for embedding text prompts. The threshold δ in DAD is 0.25, and the metanet τ is built with a simple MLP structure (Linear-ReLU-Linear) for efficiency. We take the average performance as our final result with three runs by different seeds. For more details, please refer to our supplementary material.

Main Results

Image Matting. As shown in Tab. 2, our method achieves the best results across almost all metrics on the image matting datasets. On HIM2K-NATURAL, we outperform the runner-up (SmartMatting (Ye et al. 2024)) by a notable margin in all four IMQ-based metrics, setting new records with 83.48 (IMQ-MSE), 73.45 (IMQ-MAD), 75.69 (IMQ-Grad), and 74.83 (IMQ-Conn). Similarly, on REFMATTE-RW, our method obtains the lowest error in MSE (0.010), SAD (21.31), and Gradient (19.38), while achieving competitive performance in Connectivity, ranking second.

We present the visual comparisons in Fig. 4. Our method produces more accurate and detailed results compared to existing approaches. These improvements are particularly evident in challenging cases with complex object boundaries




Figure 5: Visual comparisons with different methods on DIS-5K dataset. Zoom in for the best view.

and ambiguous foreground-background transitions.

Accurate Dichotomous Segmentation. Our framework is also applicable for DIS5K. As shown in Tab. 3, on the DIS-VD split, our method not only surpasses the previous *state-of-the-art* by a clear margin in key metrics such as $\max F_\beta$ and M , but also places among the top two overall, often trailing only slightly behind the second-best performer in those where it does not lead. A similar trend is observed on the more comprehensive DIS-TE, where we also obtain the top performance in $\max F_\beta$ (0.922), F_β^w (0.895), M (0.026), and S_α (0.920), with competitive results in E_ϕ^m (0.947) and the lowest HCE_γ (580).

The qualitative comparisons can be found in Fig. 5. Our method achieves superior performance in handling complex real-world images, particularly in preserving fine details and maintaining accurate boundaries. For example, the ropes on the viaduct (3rd row) and cilia on the rope (last row).

Ablation Studies

Unless specification, we conduct the ablation studies on Ref-Matte (Li, Zhang, and Tao 2023) dataset.

Effectiveness of Distillation. We conduct an ablation study on our knowledge distillation strategy, as shown in Tab. 4. As shown in the table, omitting knowledge distillation (w/o KD) leads to the poorest performance. Incorporating standard knowledge distillation without depth-aware guidance (Vanilla KD) only yields a marginal improvement.




Figure 6: t-SNE visualization from the foreground tokens.

By introducing the meta-net, distilling knowledge from the foreground features of either image A or B individually improves model results, demonstrating the effectiveness of feature-specific distillation. Combining both foreground distillations yields further gains. Our final model extends this strategy by also distilling background, leading to the optimal configuration.

Effectiveness of Domain Adaptation. We evaluate the impact of different alignment losses, as shown in Tab. 5. Removing alignment entirely (w/o alignment) leads to the worst performance, highlighting its importance. While \mathcal{L}_{adv} improves global structure (lower MSE/SAD), \mathcal{L}_{OT} better preserves edge details and foreground connectivity (lower Grad/Conn). Combining both achieves the best results, leveraging their complementary strengths. Moreover, Fig.6 illustrates the effect of our alignment strategy using t-SNE visualizations. In (a) "Before Alignment", the feature clus-

Configuration		Dataset: DIS-VD						Dataset: Overall DIS-TE (1-4)					
		Metric						Metric					
Methods	Venues	$maxF_\beta \uparrow$	$F_\beta^w \uparrow$	$M \downarrow$	$S_\alpha \uparrow$	$E_\phi^m \uparrow$	$HCE_\gamma \downarrow$	$maxF_\beta \uparrow$	$F_\beta^w \uparrow$	$M \downarrow$	$S_\alpha \uparrow$	$E_\phi^m \uparrow$	$HCE_\gamma \downarrow$
IS-Net (Qin et al. 2022)	ECCV ₂₂	0.791	0.717	0.074	0.813	0.856	1116	0.799	0.726	0.070	0.819	0.858	1016
PGNet (Xie et al. 2022)	CVPR ₂₂	0.798	0.733	0.067	0.824	0.879	/	0.809	0.746	0.063	0.830	0.885	/
FP-DIS (Zhou et al. 2023)	IJCAI ₂₃	0.823	0.763	0.062	0.843	0.891	1309	0.831	0.770	0.057	0.847	0.895	1164
HitNet (Hu et al. 2023)	AAAI ₂₃	0.805	0.757	0.061	0.828	0.890	1550	0.815	0.767	0.057	0.836	0.894	1001
UDUN (Pei et al. 2023)	MM ₂₃	0.823	0.763	0.059	0.838	0.892	1097	0.831	0.772	0.057	0.844	0.892	977
HQ-SAM (Ke et al. 2023)	NIPS ₂₃	0.739	0.706	0.120	0.740	0.818	1553	0.830	0.804	0.618	0.835	0.902	1299
Pi-SAM (Liu et al. 2024a)	MM ₂₄	0.883	0.866	0.035	0.889	<u>0.945</u>	1322	0.890	0.873	0.033	0.893	0.948	1191
BiRefNet (Zheng et al. 2024)	AIR ₂₄	0.891	0.854	0.038	0.898	0.931	989	0.896	0.858	0.035	0.901	0.934	916
FSA-Net (Jiang et al. 2024)	TNNLS ₂₄	0.825	0.776	0.057	0.845	0.894	1094	0.841	0.794	0.052	0.857	0.903	973
MV-ANet (Yu et al. 2024b)	CVPR ₂₄	0.904	0.861	0.035	0.909	0.937	<u>878</u>	0.916	0.855	0.035	0.905	0.938	<u>790</u>
MaskFactory (Qian et al. 2024)	NIPS ₂₄	0.835	0.759	0.072	0.866	0.923	/	0.839	0.773	0.069	0.859	0.907	/
DiffDIS (Yu et al. 2024a)	NIPS ₂₄	<u>0.918</u>	<u>0.888</u>	<u>0.029</u>	<u>0.904</u>	0.948	/	<u>0.921</u>	<u>0.892</u>	<u>0.028</u>	<u>0.905</u>	<u>0.947</u>	/
Gerpercept (Xu et al. 2024)	ICLR ₂₅	0.877	0.859	0.035	0.887	0.941	1262	0.875	0.856	0.036	0.885	0.939	1176
Ours	/	0.924	0.900	0.025	0.909	<u>0.945</u>	796	0.922	0.895	0.026	0.920	<u>0.947</u>	580

Table 3: Quantitative comparisons on DIS5K (Qin et al. 2022) validation and testing sets. The best performing results are highlighted with **bold** and runner-up results are underlined.

Methods	MSE↓	SAD↓	Grad↓	Conn↓
w/o KD	0.018	26.34	23.48	4.38
Vanilla KD	0.014	23.48	22.84	3.99
τ_+^A	0.016	23.40	23.02	3.90
τ_+^B	0.014	22.84	21.92	3.88
$\tau_+^A + \tau_+^B$	0.013	22.18	20.29	3.60
$\tau_+^A + \tau_+^B + \tau_-^A + \tau_-^B$ (Ours)	0.010	21.31	19.38	3.33

Table 4: Ablation study on the distillation.

Methods	MSE↓	SAD↓	Grad↓	Conn↓
w/o alignment	0.157	24.33	26.87	6.07
\mathcal{L}_{adv}	0.128	22.79	23.01	4.83
\mathcal{L}_{OT}	0.134	22.47	20.72	4.34
$\mathcal{L}_{adv} + \mathcal{L}_{OT}$ (Ours)	0.010	21.31	19.38	3.33

Table 5: Ablation study on the loss of domain alignment.

ters are distinctly separated and highly scattered, indicating poor consistency between different semantic regions. After applying our alignment loss, the feature clusters become more compact and exhibit significant overlap, demonstrating improved coherence and discriminability. This qualitative improvement aligns with the quantitative results in Tab. 5, where removing alignment leads to a sharp performance drop across all metrics.

Prompt Probing. We also evaluate the effectiveness of different prompts. As shown in Tab. 6, visual prompts yield better results compared to text-based guidance, likely due to their stronger spatial constraints, which help more accurately localize the foreground object. However, these visual prompts are class-agnostic, limiting their ability to distinguish between semantically meaningful regions. Therefore,

Methods	MSE↓	SAD↓	Grad↓	Conn↓
point	0.009	18.46	15.44	3.01
box	0.006	16.37	13.89	2.48
text (Ours)	0.010	21.31	19.38	3.33

Table 6: Ablation study on the prompt types.

we report the results based on text prompts.

Conclusions

In this work, we introduce a Foreground Consistent Learning Model (FCLM) for high-accuracy image matting and segmentation toward real-world. Motivated by the limitations of current datasets, we formulate the use of synthetic training data as a domain adaptation problem and introduce distillation and alignment strategies to more effectively learn discriminative foreground features. To accommodate a variety of interaction modes, our network supports both visual and textual prompts, facilitating semantic-aware and target-specific predictions. Our method achieves *state-of-the-art* results on publicly accessible real-world datasets, validating its effectiveness in practical scenarios.

Acknowledgments

This work is supported by OPPO Research Fund, the Guangdong Science and Technology Department (No. 2024ZDZX2004), the Guangzhou Industrial Information and Intelligent Key Laboratory Project (No. 2024A03J0628), the China Southern Power Grid Science and Technology Project (Project No.: 030117KC23120005), and the Nansha Key Area Science and Technology Project (No. 2024ZD006).

References

- Chen, J.; Lu, Y.; Yu, Q.; Luo, X.; Adeli, E.; Wang, Y.; Lu, L.; Yuille, A. L.; and Zhou, Y. 2021. Transunet: Transformers make strong encoders for medical image segmentation. *arXiv preprint arXiv:2102.04306*.
- Cheng, B.; Misra, I.; Schwing, A. G.; Kirillov, A.; and Girshick, R. 2022. Masked-attention mask transformer for universal image segmentation. In *CVPR*, 1290–1299.
- Cheng, H. K.; Tai, Y.-W.; and Tang, C.-K. 2021. Rethinking space-time networks with improved memory coverage for efficient video object segmentation. *NeurIPS*, 34: 11781–11794.
- Chizat, L.; Peyré, G.; Schmitzer, B.; and Vialard, F.-X. 2018. Scaling algorithms for unbalanced optimal transport problems. *Mathematics of computation*, 87(314): 2563–2609.
- Darcet, T.; Oquab, M.; Mairal, J.; and Bojanowski, P. 2023. Vision transformers need registers. *ICLR*.
- Fan, D.-P.; Ji, G.-P.; Xu, P.; Cheng, M.-M.; Sakaridis, C.; and Van Gool, L. 2023. Advances in deep concealed scene understanding. *Visual Intelligence*, 1(1): 16.
- Frogner, C.; Zhang, C.; Mobahi, H.; Araya, M.; and Poggio, T. A. 2015. Learning with a Wasserstein loss. *NeurIPS*, 28.
- Ganin, Y.; and Lempitsky, V. 2015. Unsupervised domain adaptation by backpropagation. In *ICML*, 1180–1189. PMLR.
- Ganin, Y.; Ustinova, E.; Ajakan, H.; Germain, P.; Larochelle, H.; Laviolette, F.; March, M.; and Lempitsky, V. 2016. Domain-adversarial training of neural networks. *JMLR*, 17(59): 1–35.
- Goferman, S.; Zelnik-Manor, L.; and Tal, A. 2011. Context-aware saliency detection. *TPAMI*, 34(10): 1915–1926.
- He, K.; Gkioxari, G.; Dollár, P.; and Girshick, R. 2017. Mask r-cnn. In *ICCV*, 2961–2969.
- He, K.; Zhang, X.; Ren, S.; and Sun, J. 2016. Deep residual learning for image recognition. In *CVPR*, 770–778.
- Hinton, G.; Vinyals, O.; and Dean, J. 2015. Distilling the knowledge in a neural network. *arXiv preprint arXiv:1503.02531*.
- Hou, Q.; and Liu, F. 2019. Context-aware image matting for simultaneous foreground and alpha estimation. In *ICCV*, 4130–4139.
- Hu, X.; Wang, S.; Qin, X.; Dai, H.; Ren, W.; Luo, D.; Tai, Y.; and Shao, L. 2023. High-resolution iterative feedback network for camouflaged object detection. In *AAAI*, volume 37, 881–889.
- Huang, T.; You, S.; Wang, F.; Qian, C.; and Xu, C. 2022. Knowledge distillation from a stronger teacher. *NeurIPS*, 35: 33716–33727.
- Jain, J.; Li, J.; Chiu, M. T.; Hassani, A.; Orlov, N.; and Shi, H. 2023. Oneformer: One transformer to rule universal image segmentation. In *CVPR*, 2989–2998.
- Jiang, Q.; Cheng, J.; Wu, Z.; Cong, R.; and Timofte, R. 2024. High-Precision Dichotomous Image Segmentation With Frequency and Scale Awareness. *TNNLS*.
- Kantorovich, L. V. 2006. On the Translocation of Masses. *Journal of mathematical sciences*, 133(4).
- Ke, L.; Ye, M.; Danelljan, M.; Tai, Y.-W.; Tang, C.-K.; Yu, F.; et al. 2023. Segment anything in high quality. *NeurIPS*, 36: 29914–29934.
- Ke, Z.; Sun, J.; Li, K.; Yan, Q.; and Lau, R. W. 2022. Modnet: Real-time trimap-free portrait matting via objective decomposition. In *AAAI*, volume 36, 1140–1147.
- Kirillov, A.; Mintun, E.; Ravi, N.; Mao, H.; Rolland, C.; Gustafson, L.; Xiao, T.; Whitehead, S.; Berg, A. C.; Lo, W.-Y.; et al. 2023. Segment anything. In *ICCV*, 4015–4026.
- Li, B.; Weinberger, K. Q.; Belongie, S.; Koltun, V.; and Ranftl, R. 2022a. Language-driven semantic segmentation. *ICLR*.
- Li, J.; Jain, J.; and Shi, H. 2024. Matting anything. In *CVPR*, 1775–1785.
- Li, J.; Ma, S.; Zhang, J.; and Tao, D. 2021. Privacy-preserving portrait matting. In *ACM MM*, 3501–3509.
- Li, J.; Zhang, J.; Maybank, S. J.; and Tao, D. 2022b. Bridging composite and real: towards end-to-end deep image matting. *IJCV*, 130(2): 246–266.
- Li, J.; Zhang, J.; and Tao, D. 2023. Referring image matting. In *CVPR*, 22448–22457.
- Li, Y.; and Lu, H. 2020. Natural image matting via guided contextual attention. In *AAAI*, volume 34, 11450–11457.
- Li, Z.; Bhat, S. F.; and Wonka, P. 2024. PatchRefiner: Leveraging Synthetic Data for Real-Domain High-Resolution Monocular Metric Depth Estimation. In *ECCV*.
- Lin, S.; Ryabtsev, A.; Sengupta, S.; Curless, B. L.; Seitz, S. M.; and Kemelmacher-Shlizerman, I. 2021. Real-time high-resolution background matting. In *CVPR*, 8762–8771.
- Lin, Y.; Chen, M.; Wang, W.; Wu, B.; Li, K.; Lin, B.; Liu, H.; and He, X. 2023. Clip is also an efficient segmenter: A text-driven approach for weakly supervised semantic segmentation. In *CVPR*, 15305–15314.
- Liu, C.; Jiang, X.; and Ding, H. 2024. Primitivenet: decomposing the global constraints for referring segmentation. *Visual Intelligence*, 2(1): 16.
- Liu, F.; Tran, L.; and Liu, X. 2021. Fully understanding generic objects: Modeling, segmentation, and reconstruction. In *CVPR*, 7423–7433.
- Liu, M.; Wang, M.; Ding, H.; Xu, Y.; Zhao, Y.; and Wei, Y. 2024a. Segment anything with precise interaction. In *ACM MM*, 3790–3799.
- Liu, S.; Zeng, Z.; Ren, T.; Li, F.; Zhang, H.; Yang, J.; Jiang, Q.; Li, C.; Yang, J.; Su, H.; et al. 2024b. Grounding dino: Marrying dino with grounded pre-training for open-set object detection. In *ECCV*, 38–55. Springer.
- Liu, Y.; Tian, Y.; Zhao, Y.; Yu, H.; Xie, L.; Wang, Y.; Ye, Q.; Jiao, J.; and Liu, Y. 2024c. Vmamba: Visual state space model. *NeurIPS*, 37: 103031–103063.
- Liu, Y.; Xie, J.; Shi, X.; Qiao, Y.; Huang, Y.; Tang, Y.; and Yang, X. 2021. Tripartite information mining and integration for image matting. In *ICCV*, 7555–7564.

- Oquab, M.; Darcet, T.; Moutakanni, T.; Vo, H.; Szafraniec, M.; Khalidov, V.; Fernandez, P.; Haziza, D.; Massa, F.; El-Nouby, A.; et al. 2023. Dinov2: Learning robust visual features without supervision. *arXiv preprint arXiv:2304.07193*.
- Park, G.; Son, S.; Yoo, J.; Kim, S.; and Kwak, N. 2022. Matteformer: Transformer-based image matting via prior-tokens. In *CVPR*, 11696–11706.
- Pei, J.; Zhou, Z.; Jin, Y.; Tang, H.; and Heng, P.-A. 2023. Unite-divide-unite: Joint boosting trunk and structure for high-accuracy dichotomous image segmentation. In *ACM MM*, 2139–2147.
- Qian, H.; Chen, Y.; Lou, S.; Khan, F.; Jin, X.; and Fan, D.-P. 2024. Maskfactory: Towards high-quality synthetic data generation for dichotomous image segmentation. In *NeurIPS*.
- Qin, X.; Dai, H.; Hu, X.; Fan, D.-P.; Shao, L.; and Van Gool, L. 2022. Highly accurate dichotomous image segmentation. In *ECCV*, 38–56. Springer.
- Radford, A.; Kim, J. W.; Hallacy, C.; Ramesh, A.; Goh, G.; Agarwal, S.; Sastry, G.; Askell, A.; Mishkin, P.; Clark, J.; et al. 2021. Learning transferable visual models from natural language supervision. In *ICML*, 8748–8763. PmLR.
- Ranftl, R.; Bochkovskiy, A.; and Koltun, V. 2021. Vision transformers for dense prediction. In *ICCV*, 12179–12188.
- Ravi, N.; Gabeur, V.; Hu, Y.-T.; Hu, R.; Ryali, C.; Ma, T.; Khedr, H.; Rädle, R.; Rolland, C.; Gustafson, L.; et al. 2024. Sam 2: Segment anything in images and videos. *arXiv preprint arXiv:2408.00714*.
- Ronneberger, O.; Fischer, P.; and Brox, T. 2015. U-net: Convolutional networks for biomedical image segmentation. In *MICCAI*, 234–241. Springer.
- Shelhamer, E.; Long, J.; and Darrell, T. 2016. Fully convolutional networks for semantic segmentation. *TPAMI*, 39(4): 640–651.
- Sun, Y.; Tang, C.-K.; and Tai, Y.-W. 2022. Human instance matting via mutual guidance and multi-instance refinement. In *CVPR*, 2647–2656.
- Wang, H.; Yang, G.; Zhang, S.; Qin, J.; Guo, Y.; Xu, B.; Jin, Y.; and Zhu, L. 2024. Video-instrument synergistic network for referring video instrument segmentation in robotic surgery. *IEEE TMI*.
- Wu, H.; Yang, Y.; Xu, H.; Wang, W.; Zhou, J.; and Zhu, L. 2024. Rainmamba: Enhanced locality learning with state space models for video deraining. In *ACM MM*, 7881–7890.
- Xiao, T.; Liu, Y.; Zhou, B.; Jiang, Y.; and Sun, J. 2018. Unified perceptual parsing for scene understanding. In *ECCV*, 418–434.
- Xie, C.; Xia, C.; Ma, M.; Zhao, Z.; Chen, X.; and Li, J. 2022. Pyramid grafting network for one-stage high resolution saliency detection. In *CVPR*, 11717–11726.
- Xie, E.; Wang, W.; Yu, Z.; Anandkumar, A.; Alvarez, J. M.; and Luo, P. 2021. SegFormer: Simple and efficient design for semantic segmentation with transformers. *NeurIPS*, 34: 12077–12090.
- Xing, Z.; Ye, T.; Yang, Y.; Liu, G.; and Zhu, L. 2024. Segmamba: Long-range sequential modeling mamba for 3d medical image segmentation. In *MICCAI*, 578–588. Springer.
- Xu, G.; Ge, Y.; Liu, M.; Fan, C.; Xie, K.; Zhao, Z.; Chen, H.; and Shen, C. 2024. What Matters When Repurposing Diffusion Models for General Dense Perception Tasks? In *ICLR*.
- Xu, N.; Price, B.; Cohen, S.; and Huang, T. 2017. Deep image matting. In *CVPR*, 2970–2979.
- Yang, L.; Kang, B.; Huang, Z.; Xu, X.; Feng, J.; and Zhao, H. 2024a. Depth anything: Unleashing the power of large-scale unlabeled data. In *CVPR*, 10371–10381.
- Yang, L.; Kang, B.; Huang, Z.; Zhao, Z.; Xu, X.; Feng, J.; and Zhao, H. 2024b. Depth anything v2. *NeurIPS*, 37: 21875–21911.
- Yao, J.; Wang, X.; Yang, S.; and Wang, B. 2024a. Vitmatte: Boosting image matting with pre-trained plain vision transformers. *Information Fusion*, 103: 102091.
- Yao, J.; Wang, X.; Ye, L.; and Liu, W. 2024b. Matte anything: Interactive natural image matting with segment anything model. *Image and Vision Computing*, 147: 105067.
- Ye, Z.; Liu, W.; Guo, H.; Liang, Y.; Hong, C.; Lu, H.; and Cao, Z. 2024. Unifying automatic and interactive matting with pretrained vits. In *CVPR*, 25585–25594.
- Yu, Q.; Jiang, P.-T.; Zhang, H.; Chen, J.; Li, B.; Zhang, L.; and Lu, H. 2024a. High-Precision Dichotomous Image Segmentation via Probing Diffusion Capacity. In *ICLR*.
- Yu, Q.; Zhang, J.; Zhang, H.; Wang, Y.; Lin, Z.; Xu, N.; Bai, Y.; and Yuille, A. 2021. Mask guided matting via progressive refinement network. In *CVPR*, 1154–1163.
- Yu, Q.; Zhao, X.; Pang, Y.; Zhang, L.; and Lu, H. 2024b. Multi-view aggregation network for dichotomous image segmentation. In *CVPR*, 3921–3930.
- Yuan, Y.; Fu, R.; Huang, L.; Lin, W.; Zhang, C.; Chen, X.; and Wang, J. 2021. HRFormer: high-resolution transformer for dense prediction. In *NeurIPS*, 7281–7293.
- Zhang, Y.; Guo, M.-H.; Wang, M.; and Hu, S.-M. 2024. Exploring regional clues in CLIP for zero-shot semantic segmentation. In *CVPR*, 3270–3280.
- Zhao, H.; Shi, J.; Qi, X.; Wang, X.; and Jia, J. 2017. Pyramid scene parsing network. In *CVPR*, 2881–2890.
- Zheng, P.; Gao, D.; Fan, D.-P.; Liu, L.; Laaksonen, J.; Ouyang, W.; and Sebe, N. 2024. Bilateral Reference for High-Resolution Dichotomous Image Segmentation. *CAAI Artificial Intelligence Research*.
- Zhou, H.; Wang, H.; Ye, T.; Xing, Z.; Ma, J.; Li, P.; Wang, Q.; and Zhu, L. 2024. Timeline and boundary guided diffusion network for video shadow detection. In *ACM MM*, 166–175.
- Zhou, K.; Yang, J.; Loy, C. C.; and Liu, Z. 2022. Learning to Prompt for Vision-Language Models. *IJCV*.
- Zhou, Y.; Dong, B.; Wu, Y.; Zhu, W.; Chen, G.; and Zhang, Y. 2023. Dichotomous Image Segmentation with Frequency Priors. In *IJCAI*, 1822–1830.