

TransMatting: Enhancing Transparent Objects Matting with Transformers

Huanqia Cai*, Fanglei Xue*, Lele Xu**, and Lili Guo

University of Chinese Academy of Sciences, Beijing, China
Key Laboratory of Space Utilization, Technology and Engineering Center
for Space Utilization, Chinese Academy of Sciences, Beijing, China
{caihuanqia19,xuefanglei19}@mails.ucas.ac.cn {xulele,guolili}@csu.ac.cn

Abstract. Image matting refers to predicting the alpha values of unknown foreground areas from natural images. Prior methods have focused on propagating alpha values from known to unknown regions. However, not all natural images have a specifically known foreground. Images of transparent objects, like glass, smoke, web, etc., have less or no known foreground. In this paper, we propose a Transformer-based network, TransMatting, to model transparent objects with a big receptive field. Specifically, we redesign the trimap as three learnable tri-tokens for introducing advanced semantic features into the self-attention mechanism. A small convolutional network is proposed to utilize the global feature and non-background mask to guide the multi-scale feature propagation from encoder to decoder for maintaining the contexture of transparent objects. In addition, we create a high-resolution matting dataset of transparent objects with small known foreground areas. Experiments on several matting benchmarks demonstrate the superiority of our proposed method over the current state-of-the-art methods.

Keywords: Image Matting; Vision Transformer; Deep Learning

1 Introduction

Image matting is a technique to separate the foreground object and the background from an image by predicting a precise alpha matte as a result. It has been widely used in many applications, such as image and video editing, background replacement, and virtual reality [5,46,18]. Image matting assumes that every pixel in the image I is a linear combination of the foreground object F and the background B by an alpha matte α :

$$I = \alpha F + (1 - \alpha)B, \alpha \in [0, 1] \quad (1)$$

* The first two authors contributed equally.

** Corresponding author. This work was supported by the National Natural Science Foundation of China under Grant 61901454.

Project page: <https://github.com/AceCHQ/TransMatting>

As only the image I is known in this equation, the image matting is an ill-posed problem. So many existing methods [5,18,38,43,46,29,22] take a trimap as an auxiliary input. The trimap segments the image into three parts: known foreground and background, and unknown area, indicated as white, black, and gray, separately.

Most traditional methods, including both the sampling-based methods [2,7,43,12,15,36] and propagation-based methods [5,17,38,18], utilize the known area samples to find candidate colors or propagate the known alpha value. They heavily rely on the information from known areas, especially the known foreground areas. Recently, learning-based methods directly predict alpha mattes by neural network learning from well-annotated datasets. Although these methods take a great improvement in image matting, they also need specific information from known areas to predict unknown areas. However, according to [25], more than 50% pixels in the unknown areas cannot be correlated to pixels in the known regions due to the limited reception field of deep learning methods. LFPNet [25] proposes a Center-Surround Pyramid Pooling module to propagate the context feature from the known regions to the near unknown regions. However, not all natural images have a salient and opaque object as the known foreground [21]. Images of the glass, bonfires, plastic bags, etc., have salient foregrounds but with transparent or meticulous interiors; images of the web, smoke, water drops, etc., have non-salient foregrounds. The corresponding trimaps of these kinds of images will have very few or even no foreground areas. Most of the areas will be divided into the unknown regions. It is very challenging for existing models to learn long-range features with little known information. Furthermore, with the development of modern cameras, picture resolution is becoming higher and higher. However, the reception fields of existing models could not increase as the resolution of input images do, which makes the problem even worse.

To address this issue, we make the first attempt to introduce Vision Transformer (ViT) [9] to extract features with a large receptive field. The Transformer model is first proposed in natural language processing (NLP) and has achieved great performance in computer vision tasks, such as classification [9,41,27], segmentation [53,30], and detection [4,47]. It mainly consists of a multi-head self-attention (MHSA) module and a multi-layer perception module. The MHSA module could mine information in a global scope. Thus, the ViT model could learn global semantic features of the foreground object with high-level position relevance. To further help the model integrate the low-level appearance features (*e.g.*, texture) with high-level semantic features (*e.g.*, shape), a Multi-scale Global-guided Fusion (MGF) module is proposed. The MGF takes three adjacent scales of features as input, uses the non-background mask to guide the low-level feature, and employs the high-level feature to guide the information integration. With this new MGF module, only foreground features could be transmitted to the decoder, reducing the influence of background noises.

Since the DIM [46] concatenates the trimap and RGB image to feed into the network, almost all subsequent trimap-based methods follow this strategy. However, compared with the RGB image, the trimap is very sparse and has

some high-level positional relevance [26]. Most areas in the trimap have the same value, making convolution neural networks with small kernels inefficient in extracting features. Inspired by the [cls] token in ViT, we propose a new form of trimap named the tri-token map. Three learnable tokens are used to indicate the foreground, background, and unknown categories. We denote them as tri-tokens. Based on these tri-tokens, we propose a Tri-token Guided Transformer Block (TGTB), which adds the query with the corresponding tri-tokens for introducing the trimap information into the self-attention mechanism. With this high-level position information, the Transformer module could identify which features are from the known areas and which are from the unknown areas.

Besides, there has not been any testbed for images with transparent or non-salient foreground objects. Previous datasets mainly focus on salient and opaque foregrounds, like animals [20] and portraits [37,26], which have significantly been investigated. To further help the community to dig into the transparent and non-salient cases, we collect 460 high-resolution natural images with large unknown areas and manually label their alpha mattes.

Our main contributions can be summarized as follows:

1. We propose a TGTB module, introducing the Vision Transformer module to extract global semantic features with a big receptive field. We also redesign the trimap as a tri-token map to directly bring location information to the self-attention mechanism.
2. A MGF module is proposed to integrate multi-scale features, and the global information is well organized to guide the integration with low-level Transformer features.
3. We build a high-resolution matting dataset with 460 images of the transparent or non-salient foreground. The dataset will be released to promote the development of matting technology.
4. Experiments on three matting datasets demonstrate that the proposed TransMatting method outperforms the current SOTA methods, indicating the effectiveness of our proposed modules.

2 Related Works

In this section, we first briefly review matting from two perspectives: traditional methods and deep-learning methods. Then, we further give an overview of Vision Transformer models, as the Tri-token Guided Transformer Block (TGTB) is one of the main contributions of this work.

2.1 Traditional Matting

Traditional matting methods can be divided into two categories: sampling-based and propagation-based methods. These methods mainly rely on low-level features, like color, location, etc. The sampling-based methods [2,7,43,12,15,36] first predict the colors of the foreground and background by evaluating the similarity of colors between the known foreground, background, and unknown area

in samples, and then predict alpha mattes. Various sampling techniques have been investigated, including color cluster sampling [36], edge sampling [15], ray casting [12], etc. The propagation-based methods [5,17,38] propagate the information from the known foreground and background to the unknown area by solving the sparse linear equation system [18], the Poisson equation system [13], etc., to obtain the best global optimal alpha.

2.2 Deep-Learning Matting

In recent decades, deep learning technologies have boomed in various fields of computer vision. The same goes for the image matting task. [40] combines the sampling and deep neural network to improve the accuracy of alpha matting prediction. The Indices matter method [29] proposes an index-guided method for up-sampling and down-sampling to make the detailed information in the prediction graph more complete. Based on providing a larger dataset Composition-1k [46], DIM utilizes an encoder-decoder model to directly predict alpha mattes, which effectively improves the accuracy. [39] introduces semantic classification information of the matting region and uses learnable weights and multi-class discriminators to revise the prediction results. [51] proposes a general matting framework, which is conducive to obtaining better results under the guidance of different qualities and forms. [26] further mines the information of the RGB map and trimap and fuses the global information from these maps for obtaining better alpha mattes. All of the above methods use trimap as guidance. Some trimap-free methods can predict alpha mattes without using trimap. However, the accuracy of these trimap-free methods still has a big gap compared to that of the trimap-guided ones [6,49,35,48], indicating that the trimap could help the model to capture information efficiently.

2.3 Vision Transformer

The Transformer is firstly proposed in [42] to model long-range dependencies for machine translation and has demonstrated impressive performance on NLP tasks. Inspired by this, numerous attempts have been made to adapt transformers for vision tasks, and promising results have been shown for vision fields such as image classification, objection detection, semantic segmentation, etc. In particular, ViT [9] divides the input image into patches with a size of 16×16 and feeds the patch sequences to the vanilla Transformer model. To help the training process and improve the performance, DeiT [41] proposes a teacher-student strategy, which includes a distillation token for the student to learn from the teacher. Later, Swin [27], PVT [44], Crossformer [45], and HVT [32] combine the Transformer and pyramidal structure to decrease the number of patches progressively for obtaining multi-scale feature maps. To reduce computing and memory complexity, Swin, HRFormer [52], and CrossFormer apply local-window self-attention in Transformer, which also shows superior or comparable performance compared to the counterpart CNNs. The powerful self-attention mechanism in Transformer shows great advantages over CNN by capturing global

Table 1. Comparison between different public matting datasets.

Image Matting Dataset	total num	TT num	resolution
DAPM[37]	2000	0	800×600
Composition-1k[46]	481	86	1297×1082
Distinction-646[34]	646	79	1727×1565
AIM-500[21]	500	76	1260×1397
Transparent-460 (Ours)	460	460	3820×3766

attention of the whole image. However, some researchers [23] argue that locality and globality are both essential for vision tasks. Therefore, various researchers have tried combining the locality of CNN with the globality of Transformer to improve performance further. LocalViT [23] brings depth-wise convolutions to vision transformer to combine self-attention mechanism with locality, and shows great improvement compared to the pure Transformer, like DeiT, PVT, and TNT [14].

3 Matting Dataset

According to the transparency of foregrounds, we could divide the images of matting into two types: 1) Transparent partially (TP): TP refers to that there are significant foreground and uncertainty areas, and the foreground areas can provide information for the prediction of uncertainty areas. For example, when the foreground is human, the opaque and unknown regions are the hair or clothes. 2) Transparent totally (TT): there are minor or non-salient foreground areas, and the entire image is semi-transparent or high transparent. These images include glass, plastic bags, fog, water drops, etc.

As illustrated in Tab. 1, we select four popular image matting datasets for comparison, including DAPM [37], Composition-1k, Distinctions-646 [34], and AIM-500 [21]. The DAPM dataset consists only of portraits with no translucent or transparent objects. The Composition-1k dataset contains multiple categories, while most images are portraits (227 out of 481, TP-type). The Distinctions-646 dataset also mainly consists of portraits (343 out of 646, TP-type) [26]. The AIM-500 dataset contains only 76 TT-type images (correspond to the Salient Transparent/Meticulous type and the Non-Salient type in the original dataset) but 424 TP-type images.

As we can see, the transparent objects in the above datasets only occupy a small portion. This may be because it is much more difficult to label transparent objects than other objects, limiting the progress of transparent objects in the matting field. In this work, we propose the first large-scale dataset targeting various high transparent objects called Transparent-460 dataset. Our Transparent-460 dataset includes 460 high-quality manually-annotated alpha mattes, where 410 images are for training and 50 for testing. Furthermore, to our best knowledge, the resolution of our Transparent-460 is the highest (the average resolution

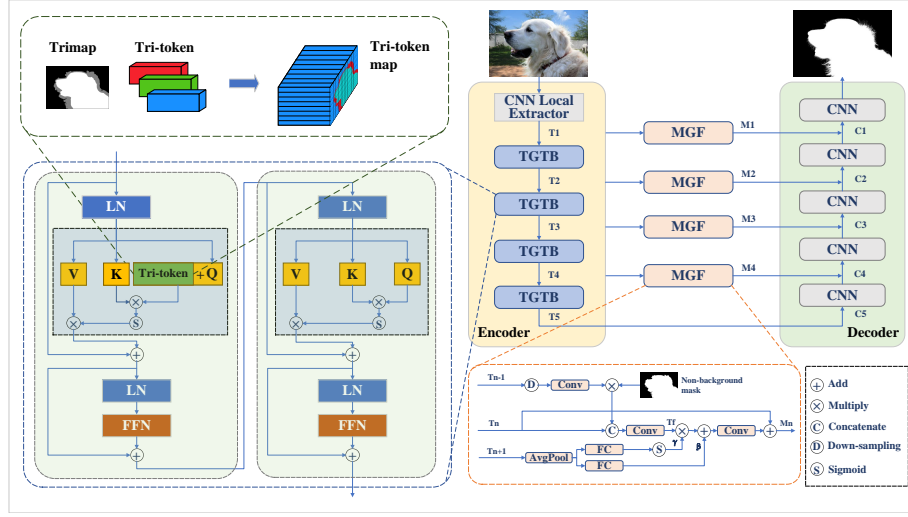


Fig. 1. The structure of our TransMatting.

is up to 3820×3766) among all datasets with high transparent objects. We believe this new matting dataset will greatly advance the matting research on objects with massive transparent areas.

4 Methodology

4.1 Motivation

By evaluating the results of some SOTA methods on TT and TP objects separately on the Composition-1k dataset (Tab. 2), we find that the results of TT, which denotes the total transparent objects, are much worse than TP, indicating that TT objects are the key to affecting the overall evaluation results. Furthermore, we find that most of the existing methods rely on the information of the foreground region for predicting the unknown region [5,17,19,22,1]. However, such methods will become useless or ineffective when facing images with no definite known regions. For example, [22] borrows features from both the known and unknown regions; when the unknown region is overwhelming in the images, the opacity propagation and the mattes prediction will face difficulties. Therefore, global information with a large or global receptive field and local features with inherent representation are needed to enhance the understanding and recognition capacity for objects with totally unknown regions. Although we can stack CNN layers to enlarge the receptive field, the information that covers the whole image is still hard to be obtained [25]. Besides, CNN also lacks global connectivity [23]. By contrast, the Transformer is good at modeling long-range connectivity with its attention mechanism.

Table 2. Performance of TT and TP objects on the Composition-1k dataset.

Methods	MSE↓			SAD↓		
	TT	TP	TT+TP	TT	TP	TT+TP
IndexNet [29]	22.87	8.9	13	110.3	18.08	45.8
GCAMatting [22]	15.89	6.2	9.1	85.72	13.68	35.3
MGMatting [51]	13.01	4.65	7.18	77.88	11.87	31.76
TransMatting(Ours)	7.49	3.4	4.58	59.37	10.35	24.96

Moreover, most existing SOTA trimap-guided methods directly concatenate the trimap with RGB image as input. However, the huge gap between the two modalities of RGB image and trimap brings great difficulties in semantic feature extraction. At the same time, the trimap cannot effectively help the model focus on the region of interest. Therefore, a more efficient way to promote the guiding role of trimap is needed.

In short, to improve the performance of TT objects, more global and local features should be captured, and an effective guidance method for the trimap should be developed.

4.2 Baseline Structure

To extract both the local and global features, we combine CNN and the Transformer model as our encoder. Specifically, the first part, like [51,22] is the same as the first two stages of ResNet34-UNet (denoted as CNN Local Extractor in Fig. 1). The second part consists of a stack of our proposed Tri-token Guided Transformer Block (TGTB) based on the Swin Transformer [27]. As the decoder, we adopt the original ResNet34-UNet, a widespread network in the matting field [51,22].

4.3 Trimap Guided Methods

Almost all SOTA methods [26,51,22,46,50,39,3] use trimap as a guide and directly concatenate the RGB image and the annotated trimap as the model’s input. However, the modalities of the RGB image and trimap are quite different. The RGB image scales from 0 to 255 and shows fine low-level features like texture, color similarity, etc. The trimap includes three values, containing high-level semantic information, like shape, location, etc., [26]. Thus, the direct concatenation between them is not the most efficient way to extract features.

Although trimap can explicitly indicate the region of interest, it is still hard to take full advantage of this information. To the best of our knowledge, we are the first to attempt to harmonize the RGB image and trimap rather than simply concatenating them. We insert a learnable trimap into the Transformer module to guide the model to concentrate on the valuable area, making the network learning more efficient and robust.

4.4 Tri-token

Inspired by the [cls] token in Vision Transformer, we design a new tri-token (shown in Fig. 1) structure, aiming to introduce the high-level semantic information directly into the self-attention mechanism to replace the inefficient concatenation methods. Given a vanilla $Trimap \in \mathbb{R}^{H \times W}$, we generate three **learnable** tri-tokens (denoted as $Token_i, i = \{0, 1, 2\}$) with different initialization to represent the known foreground, known background, and unknown areas, respectively. Every tri-token is a 1D vector, that is, $Token_i \in \mathbb{R}^C$. Then we replace every pixel in the trimap with the corresponding tri-token to generate the tri-token map, formulated as:

$$Trimap[Trimap == i] = Token_i, i = \{0, 1, 2\} \quad (2)$$

In this manner, the tri-token map can directly guide the self-attention process in the Transformer to pay more attention to the unknown areas for self-updating.

4.5 Tri-token Guided Transformer Block

Global connectivity is much more important for the prediction of total transparent objects. CNN does not have global attention, and its receptive field cannot cover the whole image [25], which leads to poor estimation for pixels outside receptive fields, while Transformer has global attention, and its receptive field can cover every pixel at the first layer.

The Transformer consists of multi-head self-attention (MHSA) and Multi-layer Perceptron (MLP) blocks. The self-attention mechanism can be thought of as a mapping between a query and a collection of key-value pairs. The output is a weighted sum of the values, and the weights are assigned by the compatibility function between the query and the relevant key. This can be implemented by Scaled Dot-Product Attention [42], in which a softmax function is used to activate the dot products of query and all keys for obtaining the weights. MHSA means that more than one self-attention is performed in parallel.

Like [27,45,52], we use non-overlapping windows whose size is $M \times M$ to divide the feature maps. The MHSA is performed within each window. The formulations of vanilla attention and our tri-token attention in a specific window are shown as follows:

$$Attention(Q, K, V) = Softmax(QK^T/\sqrt{d})V \quad (3)$$

$$Tri-token\ Attention(Q, K, V) = Softmax((Q + Tri-token)K^T/\sqrt{d})V \quad (4)$$

where $Q, K, V \in \mathbb{R}^{M^2 \times d}$ represent the query, key, and value in the attention mechanism, respectively. d is the query/key dimension. In the Tri-token Attention formulation, Q, K , and V are the same as that in the standard self-attention. The *Tri-token* is our proposed learnable *trimap* that adds to the query for forming a new tri-token query. In this way, our tri-token attention mechanism can

selectively aggregate contexts and evaluate which region should be paid more attention to with the guidance of our learnable tri-tokens.

In this way, we combine the self-attention and tri-tokens to focus on more valuable regions by considering the relationship between non-background and background areas, and finally achieve the best performance. We use our tri-token attention every five blocks in each Tri-token Guided Transformer Block (TGTB).

4.6 Multi-scale Global-guided Fusion Module

In the multi-scale feature pyramid structure, in-depth features contain more global information, while shallow features have rich local information like texture, color similarity, etc. Fusing these features is vital for accurately predicting alpha mattes for high transparent objects [35]. Although the direct sum operation can realize feature fusion, the details in the shallow features may attenuate the impact of the advanced semantics, resulting in some subtle regions missing [35]. To address this issue, we propose a Multi-scale Global-guided Fusion (MGF) module in the decoder process (see Fig. 1 for details), with both the non-background information and the advanced semantic features as guidance, to fuse the high-level semantic information and the lower ones effectively.

Specifically, we denote three adjacent features from shallow to deep as T_{n-1} , T_n , and T_{n+1} . The T_{n-1} is first down-sampled, then the Hadamard product is employed between the non-background mask and T_{n-1} to extract the low-level features of non-background, which helps to reduce the impact of complex background influence. This can guide the network to pay more attention to the foreground and unknown areas. After that, the T_{n-1} is concatenated with T_n , and a convolution layer is performed to align the channel of fused features. We mark this feature as T_f .

For the T_{n+1} , we first perform a global average pooling to generate channel-wise statistics and then use two fully connected (FC) layers to squeeze channels. As shown in Fig. 1, features output from the two FC layers are denoted as γ and β , separately. To fully capture channel-wise dependencies, we add a sigmoid function to activate γ and perform broadcast multiplication with T_f for channel re-weighting. After that, broadcast addition is performed between the channel-weighted feature and β . A convolution layer is used to fuse information from different groups. Notably, a skip connection from T_n is employed for obtaining the final fused features of MGF.

In short, considering that fusing low-level features directly may cause a negative impact on the advanced semantics [35], two techniques are proposed here. Firstly, the non-background mask is introduced into the fusion process to filter out the complex background information and further help to concentrate more attention on the foreground and unknown areas. Secondly, the global channel-wise attention from higher-level features is used for re-weighting and enhancing the important information in the fused features.

4.7 Loss Function

Following [51], we use three losses, including the alpha loss (\mathcal{L}_α), Compositional loss [46] (\mathcal{L}_{comp}), and Laplacian loss [16] (\mathcal{L}_{lap}). As formulated below, their weights are set as 0.4, 1.2, and 0.16, respectively.

$$\mathcal{L}_{final} = 0.4 * \mathcal{L}_\alpha + 1.2 * \mathcal{L}_{comp} + 0.16 * \mathcal{L}_{lap} \quad (5)$$

5 Experiments

In this section, we show our experimental settings and compare our evaluation results on the test set of Composition-1k [46], Distinction-646 [34], and our Transparent-460 datasets with other state-of-the-art methods.

5.1 Dataset

Composition-1k contains 431 and 50 unique foreground objects and manually labeled alpha mattes as training and test sets, respectively. Every foreground object is composited with 100 (for training set) and 20 (for test set) background images from COCO [24] and Pascal VOC [10]. As a result, there are 43,100 images for training and 1,000 images for testing.

Distinction-646 comprises 646 distinct foreground objects. Similar to the Composition-1k, 50 objects are divided as the test set. Following the same composition rule, there are 59,600 and 1000 images for training and testing, respectively.

Our Transparent-460 mainly consists of transparent and non-salient objects as the foreground, like water drops, jellyfish, plastic bags, glass, crystals, etc. We collect 460 high-resolution images and carefully annotate them with Photoshop. Considering the transparent objects are very meticulous, we keep the original resolution of all collected images (up to 3820×3766 pixels on average). To our best knowledge, this is the first transparent object matting dataset in such a high resolution.

5.2 Evaluation Metrics

Following [16,3,29,26], we use four metrics for evaluation, including the Sum of Absolute Differences (SAD), Mean Squared Error (MSE), Gradient error (Grad.) and Connectivity error (Conn). It is notable that the unit of MSE value is set to $1e-3$ for easy reading.

5.3 Implementation Details

We use PyTorch [33] to implement our proposed method. All the experiments are trained for 200,000 iterations. We initialize our network with ImageNet [8] pre-trained weights. The ablation experiments in Tab. 3, 4, 5 are done with 2

Table 3. The effectiveness of our proposed TGTB and MGF modules on the Composition-1k dataset.

	TGTB	MGF	SAD ↓	MSE ↓	Grad. ↓	Conn. ↓
			29.14	6.34	12.06	25.21
✓			27.45	5.66	11.77	24.30
		✓	27.21	5.57	11.23	23.25
✓	✓		26.83	5.22	10.62	22.14

NVIDIA Tesla V100 GPU with a batch size of 32. Moreover, to compare our method with the existing SOTA methods, we use a batch size of 64 with 4 NVIDIA Tesla V100 GPU to train our proposed method in Tab. 6, 7, 8. The Adam optimizer is utilized, and the initial learning rate is set to 1e-4 with the same learning rate decay strategy as [51,28]. For a fair comparison, we follow the data augmentation methods used in [22], like random crop, rotation, scaling, shearing, etc. Moreover, the trimaps for training are generated using dilation and erosion ways on alpha images by random kernel sizes from 1 to 30. Finally, we crop 512×512 patches on the center of the unknown area of alpha and composite them with the background from COCO. We use the same training conditions on the Composition-1k and Distinction-646 datasets.

5.4 Ablation Study

To evaluate the effectiveness of our new proposed modules of TGTB and MGF, and the performance with different hyper-parameters, we design the ablation study on the Composition-1k dataset.

Evaluate the effectiveness of our proposed modules. The quantitative results under the SAD, MSE, Gradient, and Connectivity errors with and without our proposed TGTB and MGF modules are illustrated in Tab. 3. As we can see, with the TGTB module, the four metrics listed above decrease to 27.45, 5.66, 11.77, and 24.30, respectively. The main reason is that our redesigned tri-token map is more suitable for propagating location information than simply concatenating to the input image. The MGF module could solely achieve similar performance, indicating that our proposed multi-scale feature fusion strategy can also help the decoder to make better use of the local and global information. When combined with the TGTB and MGF modules, the model achieves the best performance, indicating the effectiveness of the two new proposed modules.

Determine where to introduce tri-tokens. There are four TGTB stages in our encoder model. Tab. 4 reports the performance with different positions to introduce tri-tokens. As the position goes deep, the feature map size decreases, making more position information lose. On the other hand, deep stages have learned more abstract semantic features, which is suitable for mutual learning with tri-tokens. As shown in Tab. 4, both shallow and deep stages benefit from tri-tokens, indicating that the tri-tokens in TGTB modules could guide the encoder to focus on the right regions.

Table 4. Ablation results on the Composition-1k dataset with different positions to introduce the proposed tri-tokens.

Table 5. Ablation results on the Composition-1k dataset with local or (and) global features in the proposed MGF module.

Position	SAD↓	MSE↓	Grad.↓	Conn.↓
1	31.68	7.24	14.20	27.42
4	29.50	6.20	13.18	25.23
1,2,3,4	26.83	5.22	10.62	22.14

	local	global	SAD↓	MSE↓	Grad.↓	Conn.↓
			27.45	5.66	11.77	24.30
✓			27.16	5.34	11.03	22.60
	✓		27.39	5.46	11.43	23.20
✓	✓		26.83	5.22	10.62	22.14

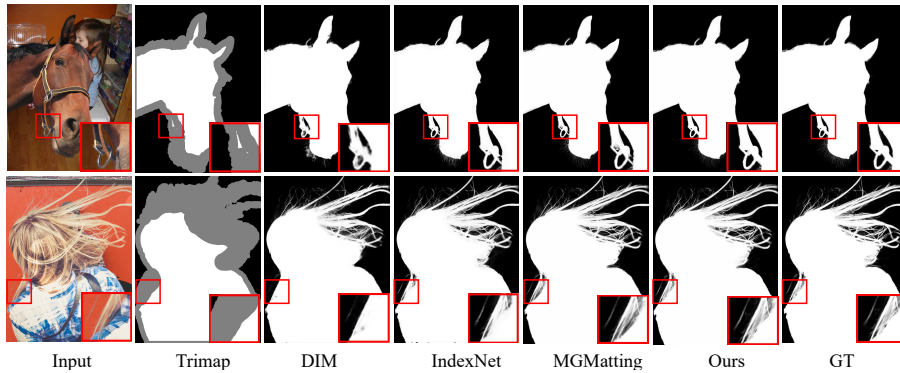


Fig. 2. Visual comparison of our TransMatting against SOTA methods on the Composition-1k test set.

The impact of local and global features in MGF. Tab. 5 reports the effectiveness of our MGF module with and without local or global branches. The local branch is proposed to integrate T_{n-1} with the non-background mask, and the global branch is responsible for introducing global features from T_{n+1} to guide the feature flow. As we can see from Tab. 5, combining local and global branches could achieve the best performance compared with using one of them solely. The main reason is the effectiveness of our MGF in fusing local (texture, border) and global (semantic, location) features for modeling unknown regions.

5.5 Comparison with Prior Work

To evaluate our method’s performance, we compare it with other state-of-the-art models on the following three datasets. Notably, we achieve the best performance on all three datasets.

Testing on Composition-1k. We show the quantitative and visual results on Tab. 6 and Fig. 2. Without any test-time augmentations, our proposed TransMatting outperforms other SOTA methods on all four evaluation metrics by only using the Composition-1k training set for training. As illustrated in Tab. 6, our

Table 6. The quantitative results on the Composition-1k test set [46]. [†] denotes results with test-time augmentation.

Methods	SAD↓	MSE↓	Grad↓	Conn↓
AlphaGAN [31]	52.4	30	38	53
DIM [29]	50.4	14	31.0	50.8
IndexNet [29]	45.8	13	25.9	43.7
AdaMatting [3]	41.7	10	16.8	-
ContextNet [16]	35.8	8.2	17.3	33.2
GCAMatting [22]	35.3	9.1	16.9	32.5
MGMatting [51]	31.5	6.8	13.5	27.3
TIMI-Net [26]	29.08	6.0	12.9	27.29
FBAMatting [11] [†]	25.8	5.2	10.6	20.8
TransMatting(Ours)	24.96	4.58	9.72	20.16

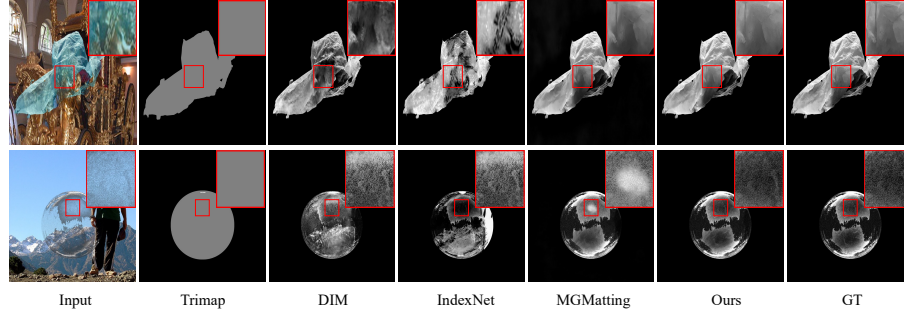


Fig. 3. Visual comparison of our TransMatting against SOTA methods on our Transparent-460 test set.

model decreases the MSE and Grad metrics heavily: from 5.2, 10.6 to 4.58 and 9.72, respectively, indicating the effectiveness of our TransMatting.

Testing on Distinction-646. Tab. 7 compares the performance of our TransMatting with other state-of-the-art methods on Distinction-646. For a fair comparison, we follow the whole inference protocol in [34,51] to calculate the metrics based on the whole image. Without any additional tuning, our method outperforms all the SOTA methods.

Testing on our Transparent-460 . Based on their release codes, we train IndexNet and MGMatting methods on our dataset and compare them with ours in Tab. 8. Our Transparent-460 dataset mainly focuses on transparent and non-salient foregrounds, which is very difficult for existing image matting methods. Surprisingly, as illustrated in Tab. 8, our TransMatting achieves promising results with only a 4.02 MSE error. Furthermore, to evaluate the generalization performance of our model. We train our TransMatting on the Composition-1k training set and directly test it on the Transparent-460 test set. The results are shown in Tab. 9. Thanks to the big receptive field and well-designed multi-scale

Table 7. The quantitative results on the Distinction-646 test set.

Methods	SAD↓	MSE↓	Grad.↓	Conn.↓
KNNMatting [5]	116.68	25	103.15	121.45
DIM [46]	47.56	9	43.29	55.90
HAttMatting [34]	48.98	9	41.57	49.93
GCAMatting [22]	27.43	4.8	18.7	21.86
MGMatting [51]	33.24	4.51	20.31	25.49
TransMatting (Ours)	25.65	3.4	16.08	21.45

Table 8. The quantitative results on our proposed Transparent-460 test set.

Methods	SAD↓	MSE↓	Grad.↓	Conn.↓
IndexNet [29]	573.09	112.53	140.76	327.97
MGMatting [51]	111.92	6.33	25.67	103.81
TransMatting (Ours)	88.34	4.02	20.99	82.56

fusion module, our model reduces nearly half of the SAD, MSE, and Conn. errors compared to the SOTA methods.

Table 9. The generalization results on our proposed Transparent-460 test set.

Methods	SAD↓	MSE↓	Grad.↓	Conn.↓
DIM [46]	356.2	49.68	146.46	296.31
IndexNet [29]	434.14	74.73	124.98	368.48
MGMatting [51]	344.65	57.25	74.54	282.79
TIMI-Net [26]	328.08	44.2	142.11	289.79
TransMatting(Ours)	192.36	20.96	41.8	158.37

6 Conclusion

In order to generalize to transparent and non-salient foregrounds, matting algorithms must have the ability to mine long-range features and utilize the semantic features in trimap. In this paper, we propose a novel Transformer-based network by redesigning a tri-token map to introduce the trimap semantic features into the long-range dependencies of the self-attention mechanism. Furthermore, a multi-scale global-guided fusion module is proposed to take the global information and local non-background mask as a guide to fuse multi-scale features for better modeling the unknown regions in transparent objects. Experiments on the Composition-1k, Distinctions-646, and our proposed Transparent-460 datasets demonstrate that our TransMatting outperforms the state-of-the-art methods.

References

1. Aksoy, Y., Aydın, T.O., Pollefeys, M.: Information-flow matting. arXiv preprint arXiv:1707.05055 (2017)
2. Berman, A., Dadourian, A., Vlahos, P.: Method for removing from an image the background surrounding a selected object (Oct 17 2000), uS Patent 6,134,346
3. Cai, S., Zhang, X., Fan, H., Huang, H., Liu, J., Liu, J., Liu, J., Wang, J., Sun, J.: Disentangled image matting. In: Proceedings of the IEEE/CVF International Conference on Computer Vision. pp. 8819–8828 (2019)
4. Carion, N., Massa, F., Synnaeve, G., Usunier, N., Kirillov, A., Zagoruyko, S.: End-to-End Object Detection with Transformers. In: European Conference on Computer Vision. Springer. https://doi.org/10.1007/978-3-030-58452-8_13, <http://arxiv.org/abs/2005.12872>
5. Chen, Q., Li, D., Tang, C.K.: Knn matting. IEEE transactions on pattern analysis and machine intelligence **35**(9), 2175–2188 (2013)
6. Chen, Q., Ge, T., Xu, Y., Zhang, Z., Yang, X., Gai, K.: Semantic human matting. In: Proceedings of the 26th ACM international conference on Multimedia. pp. 618–626 (2018)
7. Chuang, Y.Y., Curless, B., Salesin, D.H., Szeliski, R.: A bayesian approach to digital matting. In: Proceedings of the 2001 IEEE Computer Society Conference on Computer Vision and Pattern Recognition. CVPR 2001. vol. 2, pp. II–II. IEEE (2001)
8. Deng, J., Dong, W., Socher, R., Li, L.J., Li, K., Fei-Fei, L.: Imagenet: A large-scale hierarchical image database. In: 2009 IEEE conference on computer vision and pattern recognition. pp. 248–255. Ieee (2009)
9. Dosovitskiy, A., Beyer, L., Kolesnikov, A., Weissenborn, D., Zhai, X., Unterthiner, T., Dehghani, M., Minderer, M., Heigold, G., Gelly, S., et al.: An image is worth 16x16 words: Transformers for image recognition at scale. arXiv preprint arXiv:2010.11929 (2020)
10. Everingham, M., Van Gool, L., Williams, C.K., Winn, J., Zisserman, A.: The pascal visual object classes (voc) challenge. International journal of computer vision **88**(2), 303–338 (2010)
11. Forte, M., Pitié, F.: f , b , α matting. arXiv preprint arXiv:2003.07711 (2020)
12. Gastal, E.S., Oliveira, M.M.: Shared sampling for real-time alpha matting. In: Computer Graphics Forum. vol. 29, pp. 575–584. Wiley Online Library (2010)
13. Grady, L., Schiwietz, T., Aharon, S., Westermann, R.: Random walks for interactive alpha-matting. In: Proceedings of VIIP. vol. 2005, pp. 423–429 (2005)
14. Han, K., Xiao, A., Wu, E., Guo, J., Xu, C., Wang, Y.: Transformer in transformer. Advances in Neural Information Processing Systems **34** (2021)
15. He, K., Rhemann, C., Rother, C., Tang, X., Sun, J.: A global sampling method for alpha matting. In: CVPR 2011. pp. 2049–2056. IEEE (2011)
16. Hou, Q., Liu, F.: Context-aware image matting for simultaneous foreground and alpha estimation. In: Proceedings of the IEEE/CVF International Conference on Computer Vision. pp. 4130–4139 (2019)
17. Lee, P., Wu, Y.: Nonlocal matting. In: CVPR 2011. pp. 2193–2200. IEEE (2011)
18. Levin, A., Lischinski, D., Weiss, Y.: A closed-form solution to natural image matting. IEEE transactions on pattern analysis and machine intelligence **30**(2), 228–242 (2007)
19. Levin, A., Rav-Acha, A., Lischinski, D.: Spectral matting. IEEE transactions on pattern analysis and machine intelligence **30**(10), 1699–1712 (2008)

20. Li, J., Zhang, J., Maybank, S.J., Tao, D.: End-to-end animal image matting. arXiv e-prints pp. arXiv-2010 (2020)
21. Li, J., Zhang, J., Tao, D.: Deep automatic natural image matting. arXiv preprint arXiv:2107.07235 (2021)
22. Li, Y., Lu, H.: Natural image matting via guided contextual attention. In: Proceedings of the AAAI Conference on Artificial Intelligence. vol. 34, pp. 11450–11457 (2020)
23. Li, Y., Zhang, K., Cao, J., Timofte, R., Van Gool, L.: Localvit: Bringing locality to vision transformers. arXiv preprint arXiv:2104.05707 (2021)
24. Lin, T.Y., Maire, M., Belongie, S., Hays, J., Perona, P., Ramanan, D., Dollár, P., Zitnick, C.L.: Microsoft coco: Common objects in context. In: European conference on computer vision. pp. 740–755. Springer (2014)
25. Liu, Q., Xie, H., Zhang, S., Zhong, B., Ji, R.: Long-range feature propagating for natural image matting. In: Proceedings of the 29th ACM International Conference on Multimedia. pp. 526–534 (2021)
26. Liu, Y., Xie, J., Shi, X., Qiao, Y., Huang, Y., Tang, Y., Yang, X.: Tripartite information mining and integration for image matting. In: Proceedings of the IEEE/CVF International Conference on Computer Vision. pp. 7555–7564 (2021)
27. Liu, Z., Lin, Y., Cao, Y., Hu, H., Wei, Y., Zhang, Z., Lin, S., Guo, B.: Swin transformer: Hierarchical vision transformer using shifted windows. In: Proceedings of the IEEE/CVF International Conference on Computer Vision. pp. 10012–10022 (2021)
28. Loshchilov, I., Hutter, F.: Sgdr: Stochastic gradient descent with warm restarts. arXiv preprint arXiv:1608.03983 (2016)
29. Lu, H., Dai, Y., Shen, C., Xu, S.: Indices matter: Learning to index for deep image matting. In: Proceedings of the IEEE/CVF International Conference on Computer Vision. pp. 3266–3275 (2019)
30. Lu, Z., He, S., Zhu, X., Zhang, L., Song, Y.Z., Xiang, T.: Simpler Is Better: Few-Shot Semantic Segmentation With Classifier Weight Transformer. In: ICCV 2021. pp. 8741–8750. https://openaccess.thecvf.com/content/ICCV2021/html/Lu_Simpler_Is_Better_Few-Shot_Semantic_Segmentation_With_Classifier_Weight_Transformer_ICCV_2021_paper.html
31. Lutz, S., Amplianitis, K., Smolic, A.: Alphagan: Generative adversarial networks for natural image matting. arXiv preprint arXiv:1807.10088 (2018)
32. Pan, Z., Zhuang, B., Liu, J., He, H., Cai, J.: Scalable visual transformers with hierarchical pooling. arXiv e-prints pp. arXiv-2103 (2021)
33. Paszke, A., Gross, S., Massa, F., Lerer, A., Bradbury, J., Chanan, G., Killeen, T., Lin, Z., Gimelshein, N., Antiga, L., et al.: Pytorch: An imperative style, high-performance deep learning library. *Advances in neural information processing systems* **32** (2019)
34. Qiao, Y., Liu, Y., Yang, X., Zhou, D., Xu, M., Zhang, Q., Wei, X.: Attention-guided hierarchical structure aggregation for image matting. In: Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition. pp. 13676–13685 (2020)
35. Qiao, Y., Liu, Y., Zhu, Q., Yang, X., Wang, Y., Zhang, Q., Wei, X.: Multi-scale information assembly for image matting. In: *Computer Graphics Forum*. vol. 39, pp. 565–574. Wiley Online Library (2020)
36. Shahrian, E., Rajan, D., Price, B., Cohen, S.: Improving image matting using comprehensive sampling sets. In: Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition. pp. 636–643 (2013)

37. Shen, X., Tao, X., Gao, H., Zhou, C., Jia, J.: Deep automatic portrait matting. In: European conference on computer vision. pp. 92–107. Springer (2016)
38. Sun, J., Jia, J., Tang, C.K., Shum, H.Y.: Poisson matting. In: ACM SIGGRAPH 2004 Papers, pp. 315–321 (2004)
39. Sun, Y., Tang, C.K., Tai, Y.W.: Semantic image matting. In: Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition. pp. 11120–11129 (2021)
40. Tang, J., Aksoy, Y., Oztireli, C., Gross, M., Aydin, T.O.: Learning-based sampling for natural image matting. In: Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition. pp. 3055–3063 (2019)
41. Touvron, H., Cord, M., Douze, M., Massa, F., Sablayrolles, A., Jégou, H.: Training data-efficient image transformers & distillation through attention. In: International Conference on Machine Learning. pp. 10347–10357. PMLR (2021)
42. Vaswani, A., Shazeer, N., Parmar, N., Uszkoreit, J., Jones, L., Gomez, A.N., Kaiser, L., Polosukhin, I.: Attention is all you need. *Advances in neural information processing systems* **30** (2017)
43. Wang, J., Cohen, M.F.: Optimized color sampling for robust matting. In: 2007 IEEE Conference on Computer Vision and Pattern Recognition. pp. 1–8. IEEE (2007)
44. Wang, W., Xie, E., Li, X., Fan, D.P., Song, K., Liang, D., Lu, T., Luo, P., Shao, L.: Pyramid vision transformer: A versatile backbone for dense prediction without convolutions. In: Proceedings of the IEEE/CVF International Conference on Computer Vision. pp. 568–578 (2021)
45. Wang, W., Yao, L., Chen, L., Cai, D., He, X., Liu, W.: Crossformer: A versatile vision transformer based on cross-scale attention. *arXiv e-prints* pp. arXiv–2108 (2021)
46. Xu, N., Price, B., Cohen, S., Huang, T.: Deep image matting. In: Proceedings of the IEEE conference on computer vision and pattern recognition. pp. 2970–2979 (2017)
47. Yang, J., Li, C., Zhang, P., Dai, X., Xiao, B., Yuan, L., Gao, J.: Focal self-attention for local-global interactions in vision transformers. *arXiv preprint arXiv:2107.00641* (2021)
48. Yang, X., Qiao, Y., Chen, S., He, S., Yin, B., Zhang, Q., Wei, X., Lau, R.W.: Smart scribbles for image matting. *ACM Transactions on Multimedia Computing, Communications, and Applications (TOMM)* **16**(4), 1–21 (2020)
49. Yang, X., Xu, K., Chen, S., He, S., Yin, B.Y., Lau, R.: Active matting. *Advances in Neural Information Processing Systems* **31** (2018)
50. Yu, H., Xu, N., Huang, Z., Zhou, Y., Shi, H.: High-resolution deep image matting. *arXiv preprint arXiv:2009.06613* (2020)
51. Yu, Q., Zhang, J., Zhang, H., Wang, Y., Lin, Z., Xu, N., Bai, Y., Yuille, A.: Mask guided matting via progressive refinement network. In: Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition. pp. 1154–1163 (2021)
52. Yuan, Y., Fu, R., Huang, L., Lin, W., Zhang, C., Chen, X., Wang, J.: Hrformer: High-resolution transformer for dense prediction. *arXiv preprint arXiv:2110.09408* (2021)
53. Zheng, S., Lu, J., Zhao, H., Zhu, X., Luo, Z., Wang, Y., Fu, Y., Feng, J., Xiang, T., Torr, P.H., Zhang, L.: Rethinking Semantic Segmentation from a Sequence-to-Sequence Perspective with Transformers. In: 2021 IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR). pp. 6877–6886.

IEEE. <https://doi.org/10.1109/CVPR46437.2021.00681>, <https://ieeexplore.ieee.org/document/9578646/>