







# Diffusion for Natural Image Matting

Yihan Hu<sup>1,2,5</sup>, Yiheng Lin<sup>1,2</sup>, Wei Wang<sup>1,2</sup>, Yao Zhao<sup>1,2,3</sup>  
Yunchao Wei<sup>1,2,3</sup>, and Humphrey Shi<sup>4,5</sup>

<sup>1</sup> Institute of Information Science, Beijing Jiaotong University

<sup>2</sup> Visual Intelligence + X International Joint Laboratory of the Ministry of Education

<sup>3</sup> Pengcheng Laboratory, Shenzhen, China

<sup>4</sup> Georgia Institute of Technology

<sup>5</sup> Picsart AI Research (PAIR)

Yihan.hu@bjtu.edu.cn

**Abstract.** Existing natural image matting algorithms inevitably have flaws in their predictions on difficult cases, and their one-step prediction manner cannot further correct these errors. In this paper, we investigate a multi-step iterative approach for the first time to tackle the challenging natural image matting task, and achieve excellent performance by introducing a pixel-level denoising diffusion method (DiffMatte) for the alpha matte refinement. To improve iteration efficiency, we design a lightweight diffusion decoder as the only iterative component to directly denoise the alpha matte, saving the huge computational overhead of repeatedly encoding matting features. We also propose an ameliorated self-aligned strategy to consolidate the performance gains brought about by the iterative diffusion process. This allows the model to adapt to various types of errors by aligning the noisy samples used in training and inference, mitigating performance degradation caused by sampling drift. Extensive experimental results demonstrate that DiffMatte not only reaches the state-of-the-art level on the mainstream Composition-1k test set, surpassing the previous best methods by **8%** and **15%** in the SAD metric and MSE metric respectively, but also show stronger generalization ability in other benchmarks. The code will be open-sourced for the following research and applications. Code is available at <https://github.com/YihanHu-2022/DiffMatte>.

**Keywords:** Image matting · Diffusion process · Iterative refinement

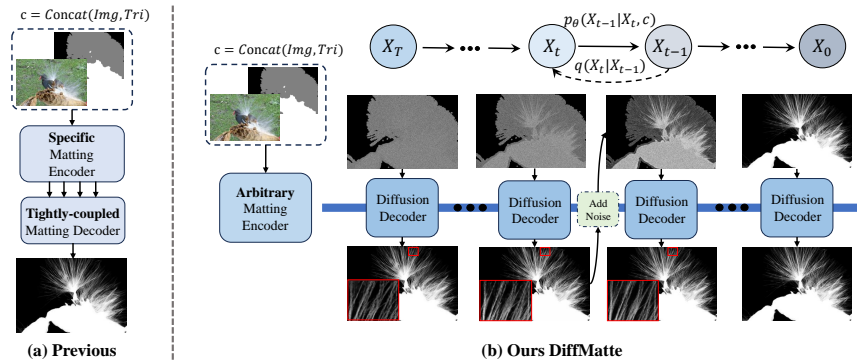
## 1 Introduction

Natural image matting is an important task in computer vision, serving the purpose of isolating foreground objects from their backgrounds. Mathematically, a natural image can be expressed as a linear combination of the foreground  $F \in \mathbb{R}^{H \times W \times C}$ , background  $B \in \mathbb{R}^{H \times W \times C}$ , and the alpha matte  $\alpha \in \mathbb{R}^{H \times W}$ , described as:

$$I_i = \alpha_i F_i + (1 - \alpha_i) B_i, \quad \alpha \in [0, 1], \quad (1)$$

---

✉ Corresponding author.



**Fig. 1:** DiffMatte introduces the diffusion process to solve the natural image matting problem. By iteratively correcting the prediction, our method obtains state-of-the-art matting accuracy. DiffMatte can be embedded into arbitrary matting encoders, which makes its application scenarios more flexible and versatile

Since the foreground color  $F_i$ , the background color  $B_i$ , and the alpha value  $\alpha$  are left unknown, solving for alpha matte is a highly ill-posed problem. To tackle this, the manually labeled trimaps are used to guide the extraction of foreground opacity with modern deep neural networks [27, 50, 66, 72].

Although the performance of previous one-step matting models keeps on increasing, they still cannot perfectly predict the alpha matte on complex cases, producing conspicuous artifacts or fine area flaws. Considering the success of iterative approaches on segmentation tasks [8, 61, 70], it is an intuitive idea to introduce the iterative mechanism for alpha matte refinement. However, coarse masks are enough to provide semantic priors to guide the iteration of segmentation methods, while matting methods rely on trimap to specifically exploit opacity information in images, which cannot be replaced by coarse alpha mattes. In addition, there is no good way to convert alpha matte into usable trimap. This inconsistency in the form of guidance and prediction prevents the matting task from borrowing the mature iteration framework of the segmentation method, hindering the exploration of iterative alpha matte refinement.

Recently, the advent of denoising diffusion models [25, 62, 63] provide an iterative process for high-fidelity and fine-grained generation [13, 52]. We notice that this unique noising and denoising process of the diffusion method naturally forms an iterative paradigm, which avoids the inconsistency between trimap and alpha matte by utilizing the noised alpha matte as additional prior information. In addition, the diffusion method can control this guidance prior by changing the input scaling of noise, and customizing the entire iterative process through the noise schedule. This allows a more flexible utilization of the alpha matte to improve matting quality.

However, applying the diffusion process to natural image matting is non-trivial for the following reasons. First, low iteration efficiency. Existing matting models, to take into account both the specific matting feature and low-level in-

formation, usually adopt a tightly coupled network design with specific feature encoders that are bulky in terms of computational overhead. This will lead to redundant calculations while receiving original-sized images (resolution above 2K to ensure clarity) during inference. This difficulty will be further exacerbated by the direct introduction of such a matting model into the iterative process. Second, the performance decline caused by the sampling drift [12]. Due to the recursive nature of the diffusion process, the flawed alpha matte will form the guidance for the next prediction during inference, which deviates from the training samples generated using the ground truth alpha matte. This problem is prominent in matting tasks because the prediction of the alpha matte requires faithful utilization of pixel-level information and demands high accuracy.

To address the challenges that arise when adapting the diffusion process to matting models, we propose the DiffMatte in this paper. Specifically, to reduce the high computational overhead, DiffMatte decouples the image encoder and decoder. As shown in Fig.1, unlike past tightly-coupled matting predictors, diffusion decoder  $\mathcal{D}$  only receives the top-level features of arbitrary matting encoder  $\mathcal{B}$  without shortcut connections. During the reverse process of inference, only the lightweight  $\mathcal{D}$  performs iteratively, and  $\mathcal{B}$  acts only once to generate high-dimensional context knowledge, which brings the benefit of a significant reduction in computational overhead. To tackle the sampling drift, DiffMatte includes a modified self-aligned strategy in the later stages of training. We add noise to the alpha matte predicted by the model as a guide to align the samples during training. This helps the model adapt to errors in previous predictions and correctly trade off the prior information for prediction in the current step. Furthermore, our strategy can handle the cumulative effect caused by multi-step errors and maintain stable performance gains during the iteration process.

We perform extensive experiments on a series of composited image matting benchmarks [55, 65, 72] and in-the-wild benchmark AIM-500 [41] to validate our DiffMatte. When adapted to various matting encoders, DiffMatte outperforms the respective baseline methods on Composition-1k, with the adaptation using ViT-B as the encoder outperforming the previous best method by **8%** on the SAD metric (**18.63**) and by **15%** on the MSE metric (**2.54**). DiffMatte also obtains higher accuracy when generalizing to AIM-500 test sets (SAD **16.31**, MSE **3.3**), demonstrating the superior in-the-wild ability of our method.

## 2 Related Work

**Natural Image Matting.** Traditional methods are mainly divided into sampling-based [9, 17, 21, 60, 69] and propagation-based methods [3, 22, 37–39, 64], according to the way they make use of color features. These approaches lack the use of context and prone to producing artifacts. Benefiting from the rapid development of deep learning, learning-based methods can access high-level semantic information with the help of neural networks. [10, 43, 45, 46] design learnable modules to exploit contextual knowledge, and [11, 16, 54, 74, 76] introduces stronger backbones [14, 23, 48, 71] to improve matting accuracy. These methods have made

significant progress, but lack exploration of low-level texture features. This leads to matting models relying on shortcut connections to provide low-level features in the one-piece UNet-like structure. In contrast to the above, [74] proposes a decoder decoupled architecture to utilize a non-hierarchical backbone network ViT, indicating the unnecessary of previous coupled network design. [40, 75] incorporate [33] to extend matting to any instances.

**Diffusion Models.** Diffusion models have achieved significant breakthroughs in various modal generation tasks, owing to their delicate denoising processes. Denoising diffusion probabilistic models (DDPM) [25] accomplish the inverse diffusion process by training a noise predictor for fine-grained image generation. Denoising diffusion implicit models (DDIM) [62] use non-Markovian processes to speed up sampling. After some successful attempts [51, 57] to fuse textual information, a group of generative large models [53, 56] have achieved surprising results with wide applications in image editing [18, 47, 73]. Diffusion models have also been studied for the generation of a wide range of modalities, including video [24, 26, 31, 68], audio [29, 34, 35], biomedical image [59, 67] and text [19, 42].


**Diffusion models for perception tasks.** Diffusion methods attract extensive research interest due to the success of diffusion modeling in the generative field. Since the pioneering work [1] introduced diffusion methods to solve image segmentation, follow-up researchers use diffusion to attempt their respective tasks. [4] formulates object detection as a denoising process. [58] involves a diffusion pipeline into depth estimation approach. [20], [6], and [36] apply diffusion to instance segmentation, panoptic segmentation, and semantic segmentation respectively, where the diffusion denoising training technique used by Pix2Seq [6] is utilized by DDP [30] to solve diversified dense prediction tasks. These works have achieved good performance by introducing the diffusion process, but [4, 30, 58] observe that the performance decreases as the step increases, and this phenomenon is exacerbated on dense prediction tasks with high accuracy demand.

### 3 DiffMatte

#### 3.1 Constructing DiffMatte Framework

In this section, we introduce our iterative matting framework with the task-specific diffusion process. Each step of DiffMatte’s prediction is based on the fixed image and trimap  $c = \text{Concat}(\text{Image}, \text{Trimap})$ , and accepts noise sample  $X$  as additional guidance. The iteration process starts with pure Gaussian noise and ends up converging to a clean alpha matte. This gradual change enables the model to balance the guidance of  $c$  and  $X$  during the process.

**DiffMatte Framework.** To iteratively refine the alpha matte, we carefully design DiffMatte’s framework. The key point of this framework is the diffusion decoder  $\mathcal{D}$ , which is the only unit of the iterative process.  $\mathcal{D}$  receives  $X_t$  at each step  $t$  as an additional pixel-level prior. This prior information needs to be fused with the matting features provided by the matte encoder  $\mathcal{B}$  to generate the alpha matte prediction of the current step. We design  $\mathcal{D}$  as a symmetric



**Fig. 2: Proposed DiffMatte Framework.** DiffMatte provides noise sample  $X$  for training through the forward diffusion process  $q$ , and uses  $X$  as an iterative prior to supplement the guidance of  $c$  in the reverse process  $p_\theta$ . Diffusion decoder  $\mathcal{D}$  serves as an iterative component and is independent of matting encoder  $\mathcal{B}$ , receiving matting features  $F$  to predict the alpha matte at each step. As a universal decoder, it can cooperate with any matting encoder.  $\mathcal{D}$  consists of Down-Block and Up-Block, which only contain convolutional layers and linear layers to process image information and time embedding  $\tau$  respectively.

UNet-like network to complete this process. This iterative  $\mathcal{D}$  can cooperate with any matting encoder  $B$ , forming a general matting architecture.

**Providing Training Samples through Forward Process.** Different from the previous one-step matting method, DiffMatte needs to obtain noise samples during training for each iteration, which is completed through the forward process in the diffusion method. We adopt the following equation [5] to define the forward process:

$$X_t \sim q(X_t|X_0) = \sqrt{\gamma_t}(bX_0) + \sqrt{1 - \gamma_t}\epsilon \quad (2)$$

where  $\gamma_t \in (0, 1)$  is a mapping of  $t$  through the noisy schedule and represents the noise intensity.  $X_0$  indicates the ideal clean sample of alpha matte.  $\epsilon$  is a standard Gaussian noise. In the previous research, three noisy schedules are set up for image generation but also work for perception tasks, namely linear schedule [5, 25], cosine schedule [52], and sigmoid schedule [28]. Our experiments show that simple linear schedules are more suitable for matting.  $b \in (0, 1]$  denotes the input scaling factor, which amplifies the noises. A smaller  $b$  will bring more destruction of detailed information at the same  $\gamma_t$ , which can be interpreted as an increase to the signal-to-noise ratio (SNR) [5, 6].

In the current phase of training, we sample a single time step  $t$  from a uniform distribution  $U(0, 1)$  following the continuous time training paradigm [7, 32], and noise  $X_0$  to  $X_t$  for an iteration training according to Eq. 2.

**Table 1: Computational overhead and running time.** We use an image with a resolution of  $2048 \times 2048$  as input, with computational overhead in TFLOPs.  $SAD_1$  indicates one-step prediction results on Composition-1k benchmark. *first* and *sub* represents the computation and time consumption generated by the first step and subsequent iterations respectively.

Models	TFLOPs <sub>first</sub>	TFLOPs <sub>sub</sub>	Times <sub>first</sub>	Times <sub>sub</sub>	SAD <sub>1</sub>
GCA [43]	2.09	-	357 <sub>ms</sub>	-	35.3
DiffMatte-Res34	1.12	0.86 24%↓	156 <sub>ms</sub>	82 <sub>ms</sub> 47%↓	<b>31.3</b>
ViTMatte-S [74]	1.69	-	617 <sub>ms</sub>	-	21.46
DiffMatte-ViTTS	2.08	0.82 60%↓	578 <sub>ms</sub>	126 <sub>ms</sub> 78%↓	<b>20.61</b>

**Iterative Prediction in Reverse Process.** Given the noisy sample  $X_t$ , DiffMatte obtains the denoised sample  $X_{t-1}$  through estimating  $\mu_{t-1} = \hat{X}_{0|t}$  with trained  $f_\theta$ . After that a sampling technique proposed by DDIM [62] is used to sample  $X_{t-1}$ , which can be defined as:

$$X_{t-1} \sim \mathcal{N}(X_{t-1}; \sqrt{\gamma_{t-1}}\mu_{t-1}, (1 - \gamma_{t-1})\mathbf{I}) \quad (3)$$

Combining the estimation of  $\mu_{t-1}$  and DDIM sampling, we get one iteration of the reverse process:

$$p_\theta(X_{t-1}|X_t, c) = \mathcal{N}(X_{t-1}; \sqrt{\gamma_{t-1}}f_\theta(X_t, t, c), (1 - \gamma_{t-1})\mathbf{I}) \quad (4)$$

The complete reverse process starts with a standard Gaussian noise  $X_T$  and passes through  $T$ -step iteration to get the final estimation  $\hat{X}_0$ . During inference, as the time step goes from  $T$  to 0, the value of  $\gamma_t$  is mapped from 1 to 0 via the noisy function. Thus  $X_T \sim \mathcal{N}(0, I)$ , which is consistent with  $\epsilon$  at the beginning time step  $T$ .

### 3.2 Iterative Refinement

**Efficiency.** The network designed for the matting task is usually a tightly coupled one-piece structure [10, 11, 27, 43], joining the entire network to the iterative diffusion process without modification leads to excessive computational overhead. Inspired by structural designs of [74], we decouple the network into separate matting encoder  $\mathcal{B}$  and iterative diffusion decoder  $\mathcal{D}$ , and the connection between the two is limited to the top-level context feature of  $\mathcal{B}$ . We implement the DiffMatte model  $f_\theta$  as:

$$\begin{aligned} F &= \mathcal{B}(c) \\ f_\theta(X_t, t, c) &= \mathcal{D}(\text{cat}(X_t, c), t, F) = \mu_{t-1} \end{aligned} \quad (5)$$

As shown in Figure 2,  $\mathcal{B}$  encodes the image with the information of trimap to get the context knowledge  $F$ , which will be reused in the reverse process.  $\mathcal{D}$  is iteratively performed in the diffusion process, and its lightweight structure

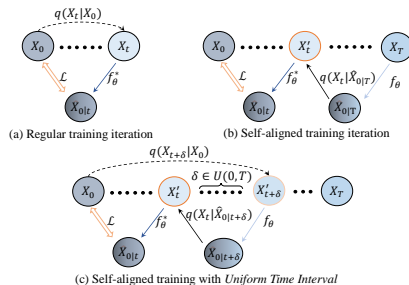
prevents excessive computational overheads. We construct the decoupled diffusion decoder  $\mathcal{D}$  with Down-Block and Up-Block. They are essentially residual convolution modules that incorporate time step information  $\tau$  encoded by linear mapping of time step  $t$ . Each module contains only 3 convolutional layers and a linear layer encoding temporal embedding. It uses only one pair of blocks per feature resolution instead of repeated stacking, making it very lightweight and fast (5.3M parameters and 126ms running time). More network details can be found in the appendix.

In this way, the diffusion process can be smoothly implemented using the lightweight diffusion decoder, and the heavy encoder only propagates once to avoid computational redundancy. As shown in Table 1, with this lightweight design, DiffMatte can save 24% of the computational overhead and 47% of the inference time in each subsequent iteration when using the ResNet34 encoder, and can save up to 60% of the computational overhead and 78% of the inference time when using the ViTS encoder. Compared to directly using the entire matting network for iteration, DiffMatte utilizes computing resources more efficiently and significantly reduces inference time.

**Self-aligned Strategy with Uniform Time Intervals.** The iterative process provided by the diffusion method is supposed to acquire performance gains with the step growth, but the results abnormally exhibit a continuous performance degradation (the evidence is displayed in Tab.5). This phenomenon can also be found in perceptual methods that introduce diffusion processes [4, 30, 58].

We attribute this phenomenon to the inconsistency of noise samples between training and inference. If only  $X$  formed by ground-truth alpha matte is used for guidance during training, the model will not be able to adapt to error-containing  $X$  during inference. Especially with the recursive nature of the iterative manner, errors in  $X$  will accumulate, exacerbating the loss of performance. This view is similarly mentioned in DDP [30], but the solution it proposes is limited as the ignorance of the accumulation of errors that occurs during the iterative process.

We propose a more refined strategy that can effectively correct this distribution variance by converting sampling targets to time intervals, obtaining a self-aligned strategy with Uniform Time Intervals (UTI). As shown in Fig.3, we sample a time interval  $\delta \in U(0, T)$ , which acts on the training time step  $t \in U(0, T - \delta)$  to obtain  $t' = t + \delta$ . In the subsequent procedures, we use an additional forward process  $q(X_{t'}|X_0)$  to calculate the previous step noise sample instead of using white noise like [30]. After that, we obtain the estimation  $\hat{X}_{0|t+\delta}$  of the current sample  $X'_{t+\delta}$  over  $X_0$  by a frozen diffusion model  $f_\theta$ . After



**Fig. 3:** Comparison of different training strategies. In contrast to the strategy proposed by [30], we take into account the drift caused by prediction errors at each time step and propose a self-aligned strategy with uniform time intervals to align the noisy sample  $X$  over all time steps.

replacing  $X_0$  with  $\hat{X}_{0|t+\delta}$  we perform a regular training iteration. The use of  $X_t$  helps the model to complete the alignment of the data distribution over the entire time domain and mitigate the accumulation of errors. Once  $\delta$  takes the value of 0, our method reverts to regular training, while when  $\delta$  takes the value of  $T - t$  it becomes the case used by [30]. We add our UTI self-aligned strategy after  $f_\theta$  is well-trained to prevent serious errors in estimation  $\hat{X}_{0|t}$  from causing the training failure.

**Algorithm 1:** DiffMatte Training

---

```
def train(cond, pha, flag, b):
    """ cond: [B, 4, H, W], pha: [B, 1, H, W] """
    """ flag: self-align entry, b: input scale """
    feat = mat_encoder(cond) # encode condition
    pha = (pha * 2) - 1 # normalize
    # forward process
    t, eps = uniform(0, 1), normal(0, 1)
    if flag == True:
        Xt = self_align(t, eps) # get aligned sample
    else:
        gamma = noisy_func(t)
        Xt = sqrt(gamma) * pha + sqrt(1-gamma) * eps
    # predict and backward
    X0 = diff_decoder(Xt, cond, feat, t)
    X0 = (X0 + 1) / 2
    loss = Losses(X0, pha)
    return loss
```

---

**Algorithm 2:** DiffMatte Inference

---

```
def inference(cond, T, b):
    """ cond: [B, 4, H, W], T: sampling steps """
    """ b: input scale """
    feat = mat_encoder(cond) # encode condition
    Xt = normal(0, 1) # noisy map of [B, 1, H, W]
    time_pairs = sample_timesteps(T)
    # reverse process
    for t, t_next in time_pairs:
        gamma, gamma_next = noisy_func(t, t_next)
        # normalize X_t by variance
        Xt = Xt / std(Xt)
        # predict X0
        X0 = diff_decoder(Xt, cond, feat, t)
        X0 = (X0 * 2) - 1 # normalize
        Xt = DDIM(X0, gamma, gamma_next, b)
    Xpred = Xt / b # rescaling
    return [Xpred + 1] / 2 # denormalize
```

---

### 3.3 Training and Inference.

Our training and inference algorithms are shown in Algorithm 1 and Algorithm 2. Our approach requires a selected noise schedule as well as input scaling, both of which parameterize the corresponding noise distribution, leading to different diffusion processes. DiffMatte involves the timing of the start of the UTI self-aligned strategy during training. We train to the 90th epoch to add it and continue until the end of training. At training time our  $f_\theta$  is supervised with the task-specific losses following the common practices [4, 6, 30]:

$$\mathcal{L} = E_{t \sim U(1, T), X_t \sim q(X_t | X_0, I)} \mathcal{L}_{mat}(f_\theta(X_t, t, c), X_0) \quad (6)$$

specifically using the combined matting loss with separate  $l_1$  loss [74],  $l_2$  loss, laplacian loss [27], and gradient penalty loss. We end up with the following objective:

$$\mathcal{L}_{mat} = \mathcal{L}_{sp} l_1 + \mathcal{L}_{l_2} + \mathcal{L}_{lap} + \mathcal{L}_{grad} \quad (7)$$

Thanks to continuous-time training, we are free to set the total number of iterations  $T$  during inference. As demonstrated in Fig.5, an increase in the number of sample steps from 1 to 10 is accompanied by an improvement in the accuracy of the final prediction.

## 4 Experiments

### 4.1 Datasets and Evaluation.

**Adobe Image Matting [72].** This dataset contains 431 unique training foreground images and 50 extra foregrounds for evaluation. The training set is constructed by compositing each foreground with 100 background images from the



COCO dataset [44]. Similarly, the validation set named Composition-1k is obtained by compositing 50 test foreground images with 20 background images from VOC2012 [15] to get a total of 1000 test images.

**Generalization.** We use the test sets of Distinctions-646 [55] (hereinafter D646) and Semantic Image Matting [65] (hereinafter SIMD) to verify the generalization performance of DiffMatte. D646 and SIMD contain 50 and 39 test foregrounds respectively, which are composited with the background in Pascol-VOC to obtain test images. SIMD provides the trimap of the foreground, while the trimap of D646 needs to be generated by ourselves.

**AIM-500 [41].** AIM-500 is the most comprehensive real image test set among the natural image matting benchmarks. It contains 500 real images with official trimap and detailed alpha matte annotations. We choose AIM-500 to evaluate DiffMatte’s in-the-wild performance.

We train our model on the Adobe Image Matting training set and perform inference on three composite image test sets to validate the matting performance as well as the generalization. We additionally use the training set in D646 for DiffMatte training of the ViT series and verify it on real-world images. The reason for using the D646 training set is that it has a wider data domain, containing more natural categories including flames and liquids. We use 4 common metrics to evaluate our DiffMatte: Sum of Absolute Differences (**SAD**), Mean Square Error (**MSE**), Gradient loss (**Grad**), and Connectivity loss (**Conn**). A lower value indicates better quality.

## 4.2 Main Results

In this section, We deploy DiffMatte’s general decoder on four popular matting encoders [43, 54, 75] and present our quantitative and qualitative comparison results with previous methods on various benchmarks. Then we discuss the performance improvements brought by DiffMatte’s unique iteration manner.

**Table 2:** Quantitative results on Composition-1k [72]. † indicates using the perturbation mask as guidance. The best results are shown in bold. S1 and S10 denote 1 and 10 steps refinement.

Methods	SAD	MSE( $10^3$ )	Grad	Conn
DIM [72]	50.4	14.0	31.0	50.8
IndexNet [49]	45.8	13.0	25.9	43.7
SampleNet [66]	40.4	9.9	-	-
Context-Aware [27]	35.8	8.2	17.3	33.2
A <sup>2</sup> U [10]	32.2	8.2	16.4	29.3
MG† [77]	31.5	6.8	13.5	27.3
SIM [65]	28.0	5.8	10.8	24.8
FBA [16]	25.8	5.2	10.6	20.8
TransMatting [2]	24.96	4.58	9.72	20.16
RMat [11]	22.87	3.9	7.74	17.84
GCA [43]	35.3	9.1	16.9	32.5
DiffMatte-Res34 (S1)	31.28	6.38	11.60	28.07
DiffMatte-Res34 (S10)	<b>29.20</b>	<b>6.04</b>	<b>11.31</b>	<b>25.48</b>
Matteformer [54]	23.80	4.03	8.68	18.90
DiffMatte-SwinT (S1)	22.05	3.54	6.67	17.03
DiffMatte-SwinT (S10)	<b>20.87</b>	<b>3.23</b>	<b>6.37</b>	<b>15.84</b>
ViTMatte-S [74]	21.46	3.3	7.24	16.21
DiffMatte-ViTS (S1)	20.61	3.08	7.14	14.98
DiffMatte-ViTS (S10)	<b>20.52</b>	<b>3.06</b>	<b>7.05</b>	<b>14.85</b>
ViTMatte-B [74]	20.33	3.0	6.74	14.78
DiffMatte-ViTB (S1)	18.84	2.56	5.86	13.23
DiffMatte-ViTB (S10)	<b>18.63</b>	<b>2.54</b>	<b>5.82</b>	<b>13.10</b>




Fig. 4: Qualitative results compared with previous SOTA methods on Composition-1k.

**Results on Composition-1k.** The quantitative results on Composition-1k are shown in Tab.2. With the ViT-B encoder, one-step DiffMatte achieves the best matting accuracy, improving the SAD metric by 1.49 (+7.3%) and the MSE metric by 0.44 (+14.7%) compared with the previous SOTA method . Fig.5 shows the refining process of each step on the Composition-1k with different matting encoders. The overall results improve as the number of iteration steps increases. Variants with Swin Tiny and ResNet-34 encoders enjoy more significant improvement. We will further explain this phenomenon in the following subsection. Fig.4 qualitatively comparing our approach to previous SOTA methods, represents the better performance in the challenging local regions, demonstrating the superiority of DiffMatte.

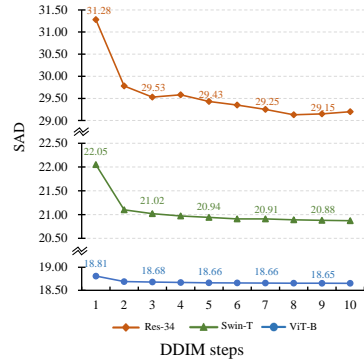


Fig. 5: The iteration results of each step of DiffMatte with different matting encoders.

**Generalization on D646 and SIMD.** The quantitative results on the D646 and SIMD test sets are shown in Tab.3. All baselines use the official weights trained on the Adobe Image Matting training set to test generalizability. The results indicate that DiffMatte outperforms the competitors on both the D646 and SIMD test sets. Notably, DiffMatte based on the ViT backbone shows a trend of performance degradation with increasing time steps on the D646 test set. We speculate that this is due to the low feature resolution provided by ViT (16x downsampling), resulting in a shallow supporting decoder that is insufficient for handling the D646 test set, which significantly differs from the training domain. We additionally trained the ViTS and ViTB models on the D646 training set,

**Table 3:** Quantitative results on D646 [55] and SIMD [65]. All methods are trained only on the Adobe Image Matting dataset. S1 and S10 denote 1 and 10 steps refinement.

Dataset Method	Distinctions-646				Semantic Image Matting Dataset				Params
	SAD	MSE	Grad	Conn	SAD	MSE	Grad	Conn	
GCA [43]	35.33	18.4	28.78	34.29	68.23	25.74	33.19	67.67	25.3M
DiffMatte-Res34 (S1)	31.53	11.87	17.52	30.86	51.49	14.10	18.32	48.70	23.9M
DiffMatte-Res34 (S10)	<b>29.38</b>	<b>11.31</b>	<b>16.19</b>	<b>28.26</b>	<b>47.75</b>	<b>14.10</b>	<b>17.19</b>	<b>44.53</b>	
Matteformer [54]	23.90	8.16	12.65	<b>18.90</b>	29.66	5.91	12.52	24.19	48.8M
DiffMatte-SwinT (S1)	23.46	6.71	10.20	21.23	30.26	5.64	9.45	24.64	38.8M
DiffMatte-SwinT (S10)	<b>23.17</b>	<b>6.58</b>	<b>10.04</b>	20.03	<b>27.51</b>	<b>5.20</b>	<b>9.12</b>	<b>22.04</b>	
ViTMatte-S [74]	23.18	7.14	13.97	19.65	27.96	5.02	10.68	22.38	25.8M
DiffMatte-ViTS (S1)	<b>22.56</b>	<b>7.09</b>	13.21	<b>19.23</b>	27.55	4.78	<b>10.26</b>	21.24	29.0M
DiffMatte-ViTS (S10)	22.96	7.22	<b>13.06</b>	19.66	<b>27.38</b>	<b>4.71</b>	10.31	<b>21.03</b>	
ViTMatte-B [74]	20.36	5.58	9.34	17.19	27.15	5.45	9.67	21.51	96.7M
DiffMatte-ViTB (S1)	<b>19.07</b>	<b>5.23</b>	9.26	<b>15.99</b>	26.83	4.92	8.26	20.72	101.4M
DiffMatte-ViTB (S10)	19.19	5.34	<b>9.26</b>	16.17	<b>25.60</b>	<b>4.69</b>	<b>8.20</b>	<b>19.84</b>	

and both recovered their refinement capabilities after training on a larger data domain. (Results shown in the Appendix)

**Results on AIM-500.** We compare the results of training on the Adobe Image Matting training set and Distinction-646 training set based on ViT-B encoder with the previous strongest model, ViTMatte, on the in-the-wild benchmark AIM-500. The results are shown in the Tab.4 . We still provide the results of

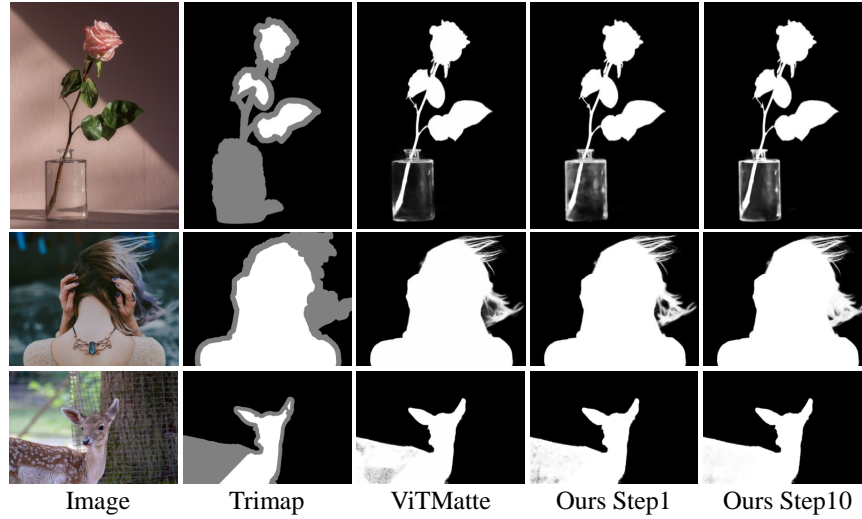
running one step and iterating ten steps respectively. We find that the effect of DiffMatte in Distinction-646 training is significantly better than Adobe Image Matting, and the effect will not deteriorate with iteration. This phenomenon shows that training on synthetic images using a wider data domain can help improve DiffMatte’s generalization ability in the real world. DiffMatte’s performance is better than ViTMatte in all four metrics, reflecting its effectiveness on real images. We further provide visualization results in Fig.6, from which we can see the quality improvements through multi-step iteration.

### 4.3 Discussion

**Promising one-step results.** When DiffMatte is applied to various matting encoders, its initial predictions outperform the corresponding baselines. This result is surprising, as the model uses only white noise as guidance in the first step, similar to one-step methods. We attribute this improvement to our UTI

**Table 4:** Quantitative results on AIM-500 [41]. ‡ indicates training on Distinctions-646. The best results are shown in bold. S1 and S10 denote 1 and 10 steps refinement.

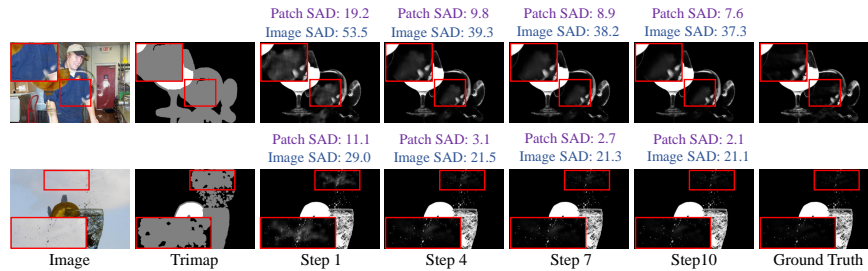
Methods	SAD	MSE( $10^3$ )	Grad	Conn
ViTMatte-B [74]	17.93	1.88	15.52	17.2
DiffMatte-ViTB (S1)	22.32	1.9	14.63	22
DiffMatte-ViTB (S10)	23.57	2.14	<b>14.58</b>	23.23
DiffMatte-ViTB <sup>‡</sup> (S1)	17.06	1.73	15.39	16.87
DiffMatte-ViTB <sup>‡</sup> (S10)	<b>16.73</b>	<b>1.7</b>	14.78	<b>16.35</b>



**Fig. 6:** We select examples to represent three difficult natural image matting scenarios. The first row is the scene with complex foregrounds containing fully transparent objects. The second row is a scene with a complex background, including strong light and blur. The third row is the scene with only a rough trimap for guidance.

training strategy. By forcing the model to overcome the errors in the guidance, the model undergoes more thorough training.

**Performance gain with iteration.** The results in 4.2 demonstrate that DiffMatte can progressively enhance the alpha matte. As shown in Fig.7, by examining the results of each iteration, we believe DiffMatte achieves this by focusing on **error-intensive** areas in the predictions. Although noise samples do not explicitly guide the model to optimize specific areas, DiffMatte’s iterative process allows the model to autonomously correct errors. Consequently, DiffMatte helps improve performance on weaker encoders like ResNet-34 and Swin-Tiny. For stronger encoders in the ViT series, the performance gains from local error correction are averaged across the test set, thus less noticeable in the curve.



**Fig. 7:** Visualization of the iterative refinement. DiffMatte tends to correct error-intensive local areas. Patch indicates the area within the red box.


**Table 5:** Ablation on the role of noises and self-aligned strategy. **iter** denotes the iterative modification.  $i$  in  $SAD_i/MSE_i$  represents the number of iteration steps. The original Self-aligned shows performance degradation after 10 steps.

Method	$SAD_1/MSE_1$	$SAD_5/MSE_5$	$SAD_{10}/MSE_{10}$	$SAD_{15}/MSE_{15}$
ViTMatte [74]	21.46/3.30	-/-	-/-	-/-
ViTMatte <b>iter</b>	23.87/4.21	24.19/4.38	24.84/4.79	25.12/5.08
<i>Diffusion</i>				
w/o Self-aligned	21.05/3.16	22.51/3.37	23.04/4.07	23.68/4.21
Self-aligned [30]	21.11/3.21	21.08/3.17	21.04/3.19	21.06/3.22
UTI Self-aligned	<b>20.61/3.08</b>	<b>20.54/3.07</b>	<b>20.52/3.06</b>	<b>20.50/3.04</b>

#### 4.4 Ablation Study

We first perform ablation experiments on our denoising diffusion approach on the Composition-1k test set. All models are trained using ViT-S as the encoder. **Effect of Diffusion Process.** We ablate the noises in our diffusion framework to explore its effects. We first remove the noise term of DiffMatte, *i.e.*,  $\epsilon$  in the Eq.2, and predict alpha matte starting from an empty image during inference. This practice makes the model overly dependent on the direct guidance of clean alpha matte, causing catastrophic inference failures that cannot be fixed even with self-aligned training. We also try to fix this issue with iterative training on the existing one-step matting method. We modified ViTMatte to accept an additional 1-channel alpha matte and constrained two consecutive predictions during retraining, using an empty image and the first prediction for guidance. The results in the second row of Tab.5 show that iterative ViTMatte is still affected by errors in previous predictions, leading to performance degradation. Additionally, the first step results are weaker than the non-iterative manner, likely due to the large gap in guidance information. In summary, noise in iterative matting disrupts previous predictions and prevents over-reliance, allowing the model to correct errors through denoising and improving matting quality.

**Effect of Self-aligned Strategies.** As shown in Tab.5, the matting performance will gradually decrease without using any self-aligned strategy. This phenomenon has been explained in Sec.3.2. When using the self-aligned strategy, our UTI has higher single-step performance and more sustained iterative improvement compared with the approach proposed in [30]. This is because the noise samples provided by UTI are more abundant, allowing the model to cope with different levels of errors, and thus better adapt the guidance provided by the noise samples. Fig.8 intuitively shows the improvement brought by the self-aligned strategy and the superiority of UTI.



**Fig. 8:** Performance gain compared with no self-aligned adopted.

**Table 6:** Ablation experiments with ViT-S [74] on Composition-1k test set [72]. We report the performance with 10 steps. Default settings are marked in gray.

Noise Schedule	SAD	MSE	Input Scaling	SAD	MSE	$N_d$	SAD	MSE	Params	TFLOPs
sigmoid	21.46	3.31	0.05	21.0	3.32	16	21.63	3.55	1.9M	0.2
cosine	21.58	3.35	0.1	21.30	3.24	24	21.23	3.27	3.4M	0.5
<b>linear</b>	<b>20.52</b>	<b>3.06</b>	<b>0.2</b>	<b>20.52</b>	<b>3.06</b>	32	20.52	3.06	5.3M	0.8
			0.5	20.88	3.24	48	20.29	3.05	10.0M	1.8
			1.0	22.19	3.75					

**(a) Noise Schedule.** We find linear schedule works best for matting task.**(b) Input Scaling.** The best input scaling factor for DiffMatte is 0.2.**(c) Diffusion Decoder Channels.**  $N_d$  set to 32 can best balance the parameters and performance.

**Effect of Noisy Schedule.** We study the effect of the noisy schedule in Tab.7a and observe that the linear schedule is best suited for matting tasks compared to cosine and sigmoid schedules. This conclusion differs from the practice of diffusion methods on generative [5] and other vision tasks [30]. This may attributed to the need for fine area guidance, as the linear schedule’s relatively stable SNR variation ratio facilitates the model for detail perception.

**Effect of Input Scaling.** We conduct the ablations on the input scaling and the results are shown in Tab.7b. As the factor decreases from 1 to 0.01, the performance of DiffMatte shows a trend of first increasing and then decreasing, finally reaching the optimum at 0.2. We believe that appropriately reducing the scaling factor can increase the difficulty of the model extracting information from noisy samples and help the model learning. However, a large scaling factor destroys too much information and reduces performance.

**Effect of Diffusion Decoder Channels.** As shown in Tab.7c, we study the role of the number of decoder channels  $N_d$ . The hyperparameter  $N_d$ , which controls the number of  $C_{out}$  channels of the convolutional layer in Up-Block and Down-Block pairs, can be varied to meet accuracy or recourses requirements. As  $N_d$  increases, the matting accuracy of the model will be improved, but the number of parameters and computational overhead will also increase. We find that when  $N_d$  is set to 32, it can best balance the performance and costs.

## 5 Conclusion

In this paper, we proposes an iterative natural image matting framework by introducing the denoising diffusion process. Our DiffMatte achieves new state-of-the-art performances on the Composition-1k, and beats the corresponding baseline on the generalization test and AIM-500 benchmark. By iteratively revising alpha matte, DiffMatte can improve prediction quality with low computational overhead and running time. This contributes to our decoupled decoder design and UTI self-aligned training strategy. We hope that DiffMatte can promote research on interactive matting and lead to practical image editing applications.

## Acknowledgements

This research was funded by the Fundamental Research Funds for the Central Universities (No. 2023JBZD003, No. 2022XKRC015), and supported by the National NSF of China (No.U23A20314).

## References

1. Amit, T., Shaharbany, T., Nachmani, E., Wolf, L.: Segdiff: Image segmentation with diffusion probabilistic models. arXiv preprint arXiv:2112.00390 (2021) [4](#)
2. Cai, H., Xue, F., Xu, L., Guo, L.: Transmatting: Enhancing transparent objects matting with transformers. In: European Conference on Computer Vision. pp. 253–269. Springer (2022) [9](#)
3. Chen, Q., Li, D., Tang, C.K.: Knn matting. IEEE transactions on pattern analysis and machine intelligence **35**(9), 2175–2188 (2013) [3](#)
4. Chen, S., Sun, P., Song, Y., Luo, P.: Diffusiondet: Diffusion model for object detection. In: Proceedings of the IEEE/CVF International Conference on Computer Vision. pp. 19830–19843 (2023) [4](#), [7](#), [8](#)
5. Chen, T.: On the importance of noise scheduling for diffusion models. arXiv preprint arXiv:2301.10972 (2023) [5](#), [14](#)
6. Chen, T., Li, L., Saxena, S., Hinton, G., Fleet, D.J.: A generalist framework for panoptic segmentation of images and videos. In: Proceedings of the IEEE/CVF International Conference on Computer Vision. pp. 909–919 (2023) [4](#), [5](#), [8](#)
7. Chen, T., Zhang, R., Hinton, G.: Analog bits: Generating discrete data using diffusion models with self-conditioning. arXiv preprint arXiv:2208.04202 (2022) [5](#)
8. Cheng, H.K., Chung, J., Tai, Y.W., Tang, C.K.: 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. pp. 8890–8899 (2020) [2](#)
9. 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) [3](#)
10. Dai, Y., Lu, H., Shen, C.: Learning affinity-aware upsampling for deep image matting. In: Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition. pp. 6841–6850 (2021) [3](#), [6](#), [9](#)
11. Dai, Y., Price, B., Zhang, H., Shen, C.: Boosting robustness of image matting with context assembling and strong data augmentation. In: Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition. pp. 11707–11716 (2022) [3](#), [6](#), [9](#)
12. Daras, G., Dagan, Y., Dimakis, A.G., Daskalakis, C.: Consistent diffusion models: Mitigating sampling drift by learning to be consistent. arXiv preprint arXiv:2302.09057 (2023) [3](#)
13. Dhariwal, P., Nichol, A.: Diffusion models beat gans on image synthesis. Advances in neural information processing systems **34**, 8780–8794 (2021) [2](#)
14. 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) [3](#)

15. 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**, 303–338 (2010) [9](#)
16. Forte, M., Pitié, F.: *f, b*, alpha matting. arXiv preprint arXiv:2003.07711 (2020) [3](#), [9](#)
17. 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) [3](#)
18. Goel, V., Peruzzo, E., Jiang, Y., Xu, D., Sebe, N., Darrell, T., Wang, Z., Shi, H.: Pair-diffusion: Object-level image editing with structure-and-appearance paired diffusion models. arXiv preprint arXiv:2303.17546 (2023) [4](#)
19. Gong, S., Li, M., Feng, J., Wu, Z., Kong, L.: Diffuseq: Sequence to sequence text generation with diffusion models. arXiv preprint arXiv:2210.08933 (2022) [4](#)
20. Gu, Z., Chen, H., Xu, Z., Lan, J., Meng, C., Wang, W.: Diffusioninst: Diffusion model for instance segmentation. arXiv preprint arXiv:2212.02773 (2022) [4](#)
21. 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) [3](#)
22. He, K., Sun, J., Tang, X.: Fast matting using large kernel matting laplacian matrices. In: *2010 IEEE Computer Society Conference on Computer Vision and Pattern Recognition*. pp. 2165–2172. IEEE (2010) [3](#)
23. He, K., Zhang, X., Ren, S., Sun, J.: Deep residual learning for image recognition. In: *Proceedings of the IEEE conference on computer vision and pattern recognition*. pp. 770–778 (2016) [3](#)
24. Ho, J., Chan, W., Saharia, C., Whang, J., Gao, R., Gritsenko, A., Kingma, D.P., Poole, B., Norouzi, M., Fleet, D.J., et al.: Imagen video: High definition video generation with diffusion models. arXiv preprint arXiv:2210.02303 (2022) [4](#)
25. Ho, J., Jain, A., Abbeel, P.: Denoising diffusion probabilistic models. *Advances in neural information processing systems* **33**, 6840–6851 (2020) [2](#), [4](#), [5](#)
26. Hong, W., Ding, M., Zheng, W., Liu, X., Tang, J.: Cogvideo: Large-scale pretraining for text-to-video generation via transformers. arXiv preprint arXiv:2205.15868 (2022) [4](#)
27. 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) [2](#), [6](#), [8](#), [9](#)
28. Jabri, A., Fleet, D., Chen, T.: Scalable adaptive computation for iterative generation. arXiv preprint arXiv:2212.11972 (2022) [5](#)
29. Jeong, Y., Ryoo, W., Lee, S., Seo, D., Byeon, W., Kim, S., Kim, J.: The power of sound (tpos): Audio reactive video generation with stable diffusion. In: *Proceedings of the IEEE/CVF International Conference on Computer Vision*. pp. 7822–7832 (2023) [4](#)
30. Ji, Y., Chen, Z., Xie, E., Hong, L., Liu, X., Liu, Z., Lu, T., Li, Z., Luo, P.: Ddp: Diffusion model for dense visual prediction. arXiv preprint arXiv:2303.17559 (2023) [4](#), [7](#), [8](#), [13](#), [14](#)
31. Khachatryan, L., Movsisyan, A., Tadevosyan, V., Henschel, R., Wang, Z., Navasardyan, S., Shi, H.: Text2video-zero: Text-to-image diffusion models are zero-shot video generators. In: *Proceedings of the IEEE/CVF International Conference on Computer Vision (ICCV)* (2023) [4](#)
32. Kingma, D., Salimans, T., Poole, B., Ho, J.: Variational diffusion models. *Advances in neural information processing systems* **34**, 21696–21707 (2021) [5](#)
33. Kirillov, A., Mintun, E., Ravi, N., Mao, H., Rolland, C., Gustafson, L., Xiao, T., Whitehead, S., Berg, A.C., Lo, W.Y., et al.: Segment anything. arXiv preprint arXiv:2304.02643 (2023) [4](#)



34. Kolesnikov, A., Beyer, L., Zhai, X., Puigcerver, J., Yung, J., Gelly, S., Houlsby, N.: Big transfer (bit): General visual representation learning. In: Computer Vision–ECCV 2020: 16th European Conference, Glasgow, UK, August 23–28, 2020, Proceedings, Part V 16. pp. 491–507. Springer (2020) [4](#)
35. Kong, Z., Ping, W., Huang, J., Zhao, K., Catanzaro, B.: Diffwave: A versatile diffusion model for audio synthesis. arXiv preprint arXiv:2009.09761 (2020) [4](#)
36. Lai, Z., Duan, Y., Dai, J., Li, Z., Fu, Y., Li, H., Qiao, Y., Wang, W.: Denoising diffusion semantic segmentation with mask prior modeling. arXiv preprint arXiv:2306.01721 (2023) [4](#)
37. Lee, P., Wu, Y.: Nonlocal matting. In: CVPR 2011. pp. 2193–2200. IEEE (2011) [3](#)
38. 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) [3](#)
39. Levin, A., Rav-Acha, A., Lischinski, D.: Spectral matting. IEEE transactions on pattern analysis and machine intelligence **30**(10), 1699–1712 (2008) [3](#)
40. Li, J., Jain, J., Shi, H.: Matting anything. arXiv preprint arXiv:2306.05399 (2023) [4](#)
41. Li, J., Zhang, J., Tao, D.: Deep automatic natural image matting. arXiv preprint arXiv:2107.07235 (2021) [3](#), [9](#), [11](#)
42. Li, X., Thickstun, J., Gulrajani, I., Liang, P.S., Hashimoto, T.B.: Diffusion-lm improves controllable text generation. Advances in Neural Information Processing Systems **35**, 4328–4343 (2022) [4](#)
43. Li, Y., Lu, H.: Natural image matting via guided contextual attention. In: AACL (2020) [3](#), [6](#), [9](#), [11](#)
44. 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: Computer Vision–ECCV 2014: 13th European Conference, Zurich, Switzerland, September 6–12, 2014, Proceedings, Part V 13. pp. 740–755. Springer (2014) [9](#)
45. 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) [3](#)
46. Liu, Q., Zhang, S., Meng, Q., Li, R., Zhong, B., Nie, L.: Rethinking context aggregation in natural image matting. arXiv preprint arXiv:2304.01171 (2023) [3](#)
47. Liu, X., Park, D.H., Azadi, S., Zhang, G., Chopikyan, A., Hu, Y., Shi, H., Rohrbach, A., Darrell, T.: More control for free! image synthesis with semantic diffusion guidance. In: Proceedings of the IEEE/CVF Winter Conference on Applications of Computer Vision. pp. 289–299 (2023) [4](#)
48. 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) [3](#)
49. 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) [9](#)
50. Lutz, S., Amliantitis, K., Smolic, A.: Alphagan: Generative adversarial networks for natural image matting. arXiv preprint arXiv:1807.10088 (2018) [2](#)
51. Nichol, A., Dhariwal, P., Ramesh, A., Shyam, P., Mishkin, P., McGrew, B., Sutskever, I., Chen, M.: Glide: Towards photorealistic image generation and editing with text-guided diffusion models. arXiv preprint arXiv:2112.10741 (2021) [4](#)
52. Nichol, A.Q., Dhariwal, P.: Improved denoising diffusion probabilistic models. In: International Conference on Machine Learning. pp. 8162–8171. PMLR (2021) [2](#), [5](#)

53. Oppenlaender, J.: The creativity of text-to-image generation. In: Proceedings of the 25th International Academic Mindtrek Conference. pp. 192–202 (2022) [4](#)
54. Park, G., Son, S., Yoo, J., Kim, S., Kwak, N.: Matteformer: Transformer-based image matting via prior-tokens. In: Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition. pp. 11696–11706 (2022) [3](#), [9](#), [11](#)
55. 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) [3](#), [9](#), [11](#)
56. Ramesh, A., Dhariwal, P., Nichol, A., Chu, C., Chen, M.: Hierarchical text-conditional image generation with clip latents, 2022. URL <https://arxiv.org/abs/2204.06125> [7](#) (2022) [4](#)
57. Rombach, R., Blattmann, A., Lorenz, D., Esser, P., Ommer, B.: High-resolution image synthesis with latent diffusion models. In: Proceedings of the IEEE/CVF conference on computer vision and pattern recognition. pp. 10684–10695 (2022) [4](#)
58. Saxena, S., Kar, A., Norouzi, M., Fleet, D.J.: Monocular depth estimation using diffusion models. arXiv preprint arXiv:2302.14816 (2023) [4](#), [7](#)
59. Schneuing, A., Du, Y., Harris, C., Jamasb, A., Igashov, I., Du, W., Blundell, T., Lió, P., Gomes, C., Welling, M., et al.: Structure-based drug design with equivariant diffusion models. arXiv preprint arXiv:2210.13695 (2022) [4](#)
60. Shahriari, 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) [3](#)
61. Shen, T., Zhang, Y., Qi, L., Kuen, J., Xie, X., Wu, J., Lin, Z., Jia, J.: High quality segmentation for ultra high-resolution images. In: Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition. pp. 1310–1319 (2022) [2](#)
62. Song, J., Meng, C., Ermon, S.: Denoising diffusion implicit models. arXiv preprint arXiv:2010.02502 (2020) [2](#), [4](#), [6](#)
63. Song, Y., Ermon, S.: Generative modeling by estimating gradients of the data distribution. *Advances in neural information processing systems* **32** (2019) [2](#)
64. Sun, J., Jia, J., Tang, C.K., Shum, H.Y.: Poisson matting. In: ACM SIGGRAPH 2004 Papers, pp. 315–321 (2004) [3](#)
65. 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) [3](#), [9](#), [11](#)
66. 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) [2](#), [9](#)
67. Trippe, B.L., Yim, J., Tischer, D., Baker, D., Broderick, T., Barzilay, R., Jaakkola, T.: Diffusion probabilistic modeling of protein backbones in 3d for the motif-scaffolding problem. arXiv preprint arXiv:2206.04119 (2022) [4](#)
68. Villegas, R., Babaeizadeh, M., Kindermans, P.J., Moraldo, H., Zhang, H., Saffar, M.T., Castro, S., Kunze, J., Erhan, D.: Phenaki: Variable length video generation from open domain textual description. arXiv preprint arXiv:2210.02399 (2022) [4](#)
69. 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) [3](#)
70. Wang, M., Ding, H., Liew, J.H., Liu, J., Zhao, Y., Wei, Y.: Segrefiner: Towards model-agnostic segmentation refinement with discrete diffusion process. arXiv preprint arXiv:2312.12425 (2023) [2](#)

71. 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) [3](#)
72. 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) [2](#), [3](#), [8](#), [9](#), [14](#)
73. Xu, X., Guo, J., Wang, Z., Huang, G., Essa, I., Shi, H.: Prompt-free diffusion: Taking "text" out of text-to-image diffusion models. arXiv preprint arXiv:2305.16223 (2023) [4](#)
74. Yao, J., Wang, X., Yang, S., Wang, B.: Vitmatte: Boosting image matting with pretrained plain vision transformers. arXiv preprint arXiv:2305.15272 (2023) [3](#), [4](#), [6](#), [8](#), [9](#), [11](#), [13](#), [14](#)
75. Yao, J., Wang, X., Ye, L., Liu, W.: Matte anything: Interactive natural image matting with segment anything models. arXiv preprint arXiv:2306.04121 (2023) [4](#), [9](#)
76. Yu, H., Xu, N., Huang, Z., Zhou, Y., Shi, H.: High-resolution deep image matting. In: Proceedings of the AAAI Conference on Artificial Intelligence. vol. 35, pp. 3217–3224 (2021) [3](#)
77. 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) [9](#)