

Fast Soft Color Segmentation

Naofumi Akimoto¹ *Huachun Zhu²Yanghua Jin²Yoshimitsu Aoki¹¹ Keio University² Preferred Networks

nakimoto@aoki-medialab.jp

{zhu, jinhy}@preferred.jp

aoki@elec.keio.ac.jp

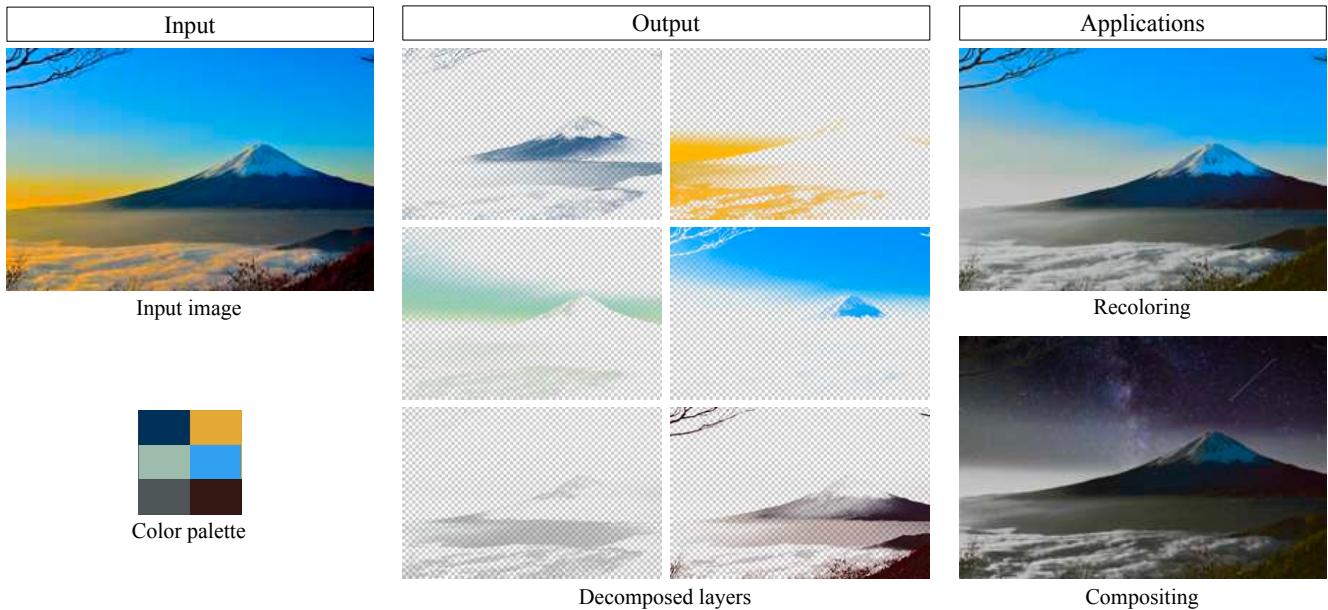


Figure 1: We propose a neural network based framework that, by using a single RGB image and a specified color palette, decomposes the image into multiple RGBA layers, each of which contains homogeneous colors. The decomposed layers can be created instantaneously and are useful for image and video editing, such as recoloring and compositing.

Abstract

We address the problem of soft color segmentation, defined as decomposing a given image into several RGBA layers, each containing only homogeneous color regions. The resulting layers from decomposition pave the way for applications that benefit from layer-based editing, such as recoloring and compositing of images and videos. The current state-of-the-art approach for this problem is hindered by slow processing time due to its iterative nature, and consequently does not scale to certain real-world scenarios. To address this issue, we propose a neural network based method for this task that decomposes a given image into multiple layers in a single forward pass. Furthermore, our method separately decomposes the color layers and the al-

pha channel layers. By leveraging a novel training objective, our method achieves proper assignment of colors amongst layers. As a consequence, our method achieves promising quality without existing issue of inference speed for iterative approaches. Our thorough experimental analysis shows that our method produces qualitative and quantitative results comparable to previous methods while achieving a 300,000x speed improvement. Finally, we utilize our proposed method on several applications, and demonstrate its speed advantage, especially in video editing.

1. Introduction

Image segmentation is the task of decomposing an image into meaningful regions. A typical approach to this problem assigns a single class to each pixel in an image. However, such hard segmentation is far from ideal when the distinc-

*A part of this work was done while the first author worked at Preferred Networks as an intern.

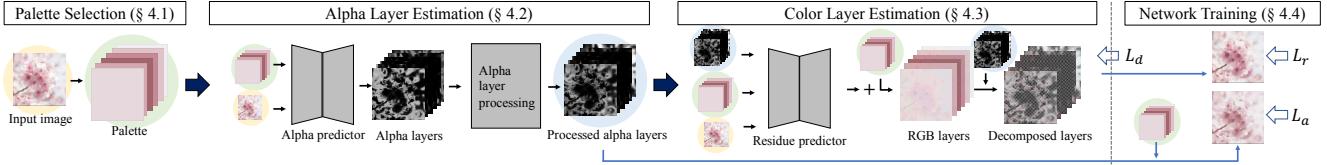


Figure 2: Our framework consists of three stages. First, palette colors are selected manually or automatically based on an input image. Second, an alpha predictor (U-Net) predicts alpha layers using the input image and the palette. Then, the alpha layers are normalized to satisfy an alpha-add condition (Eq. (2)), or undergo an alpha layer processing, such as smooth filtering or mask operation. Finally, with the input image, the palette, and the processed alpha layers, a residue predictor (U-Net) estimates the difference between the palette color and the ground-truth image at each layer in order to restore the input image from the decomposed layers. These two networks are trained jointly.

tion of meaningful regions is ambiguous, such as objects with motion blur, transparency or illumination. A flexible alternative is soft segmentation, which allow a single pixel to be assigned to multiple classes. This paper addresses a subtask of soft segmentation: soft color segmentation. the goal of this task is to decompose an RGB image into several color-considering RGBA layers, each of which is made up of homogeneous colors. The decomposed soft color layers allow a user to target specific color regions in editing tasks such as recoloring and compositing. In an approach to solve the problem, existing methods adopt either an optimization-based [2, 3, 12] or a geometric [23, 24] approach. However, it takes a considerable amount of time for existing methods to decompose a high-resolution image or a series of images (e.g. frames in a video).

In this work, we propose a neural network based approach that significantly speed up soft color segmentation while retaining the quality of the soft color layers produced. Our system consists of three stages, as shown in Figure 2. In 1) **palette selection**, we automatically or manually select palette colors, each of which indicates the mean color of a target color layer. In 2) **alpha layer estimation**, the alpha predictor estimates alpha layers corresponding to the selected palette colors. After processing the alpha layers, in 3) **color layer estimation**, the residue predictor estimates color residues that indicate the displacement of the colors from the palette color. For each pixel in a final RGBA layer, the color value (the RGB part) is the sum of the palette color and its residue at that pixel, and the alpha value (the A part) is taken from the corresponding processed alpha layer.

Compared to existing methods, the proposed method reduces the decomposition time by a factor of 300,000. We achieve this through optimization of the objective function on the training dataset in advance instead of minimizing an energy function on the input image on the fly, which is common in optimization-base approaches [2, 3, 12]. The training objective of the networks consists of a reconstruction loss, a regularization loss, and a distance loss. We jointly

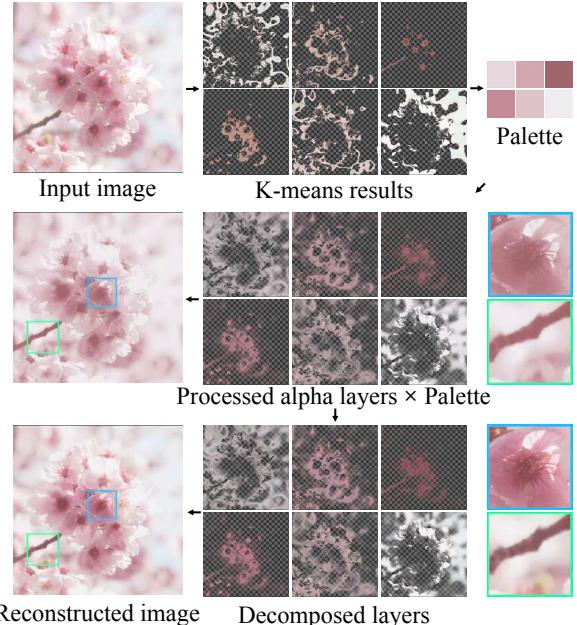


Figure 3: The reconstruction process from Figure 2. Note that if we do a weighted sum of the palette colors and the corresponding alpha layers, we can also obtain an image (middle left) similar to the input image, but residues are required to recover the input image with various colors and enhanced details (bottom left).

train the alpha predictor and the residue predictor in a self-supervised manner, and thus no extra manual labeling or annotation are required. After training, the networks decompose an image in a feed-forward pass. The speed gain paves the way for practical applications including real-time editing and frame-by-frame video decomposition.

To summarize our main contributions,

- We propose the first neural network based approach and a novel training objective for soft color segmentation.

- We conduct qualitative and quantitative experiments to demonstrate that our method significantly outpaces state-of-the-art methods while maintaining comparable visual quality.
- We apply our method to several practical applications, especially video editing, with unprecedented efficiency.

2. Related Work

Layer decomposition, actively studied in both computer vision and computer graphics, is the task of decomposing a single image into multiple RGB or RGBA images. For example, reflection removal [7, 27, 33], haze removal [10, 19, 32, 37], and rain removal [9, 31, 14] estimate the foreground and background layers and the mixture ratios for each pixel. Segmentation is also one of the layer decomposition tasks. Part segmentation includes human parsing [15, 30] or vector graphics layer generation [8, 21]. Semantic segmentation has been actively addressed [5, 16, 17, 20, 35]. The segmentation tasks mentioned above classify each pixel into a single class. In contrast to these hard segmentation tasks, *soft* segmentation are studied [2, 18, 4, 22] to express transitions between regions of classes using alpha value. For instance, matting [6, 1, 13, 25, 26, 29, 34] obtains transitions between a foreground and a background. Soft segmentation is more suitable than hard segmentation for handling motion blur, transparency, illumination, and so on.

Soft color segmentation is a specific soft segmentation task of decomposing a single image into multiple RGB or RGBA color layers, each of which contains homogeneous colors. The state-of-the-art method proposed by Aksoy *et al.* [3] improved the method in Aksoy *et al.* [2] by adding a sparsity term to the color unmixing formulation. Research similar to Aksoy *et al.* [3] includes Koyama *et al.* [12] and Tan *et al.* [23]. Koyama *et al.* [12] generalized Aksoy *et al.* [3] to enable advanced color blending. Tan *et al.* [23] are a geometric approach that finds a RGBXY convex hull, extending Tan *et al.* [24]. Whereas Koyama *et al.* [12] and Tan *et al.* [24] consider the order of layer stacking, Aksoy *et al.* [3] and Tan *et al.* [23] deal with a *linear additive model*, which does not consider the order. Aksoy *et al.* [3] have shown that color unmixing based segmentation can be achieved by solving the constrained minimization problem, as explained in Section 3.

3. Motivation

Several works [3, 12, 24] adopt an optimization-based approach to decompose an image into multiple color-considering RGBA layers. In particular, the state-of-the-art method from Aksoy *et al.* [3] proposes the color unmixing formulation, in which the objective function consists of a

color constraint, an alpha constraint, a box constraint, and a sparse color unmixing energy function.

To be specific, the color constraint requires that the decomposed layers restore the original image. For a pixel p in the i^{th} layer,

$$\sum_i \alpha_i^p \mathbf{u}_i^p = \mathbf{c}^p \quad \forall p, \quad (1)$$

where \mathbf{c}^p denotes an original color at a pixel p , and \mathbf{u}_i^p denotes a layer color at p . We omit the superscript p in the remainder of the paper for convenience.

The alpha constraint enforces the sum of the decomposed layers to be an opaque image:

$$\sum_i \alpha_i = 1. \quad (2)$$

The box constraint requires that alpha and color values should be in bound:

$$\alpha_i, \mathbf{u}_i \in [0, 1] \quad \forall i. \quad (3)$$

The sparse color unmixing energy function is the weighted sum of the distances between colors and the color distribution in each layer, plus a sparsity term:

$$\mathcal{F}_S = \sum_i \alpha_i \mathcal{D}_i(\mathbf{u}_i) + \sigma \left(\frac{\sum_i \alpha_i}{\sum_i \alpha_i^2} - 1 \right), \quad (4)$$

where the layer cost $\mathcal{D}_i(\mathbf{u}_i)$ is defined as the squared Mahalanobis distance of the layer color \mathbf{u}_i to the layer distribution $\mathcal{N}(\mathbf{u}_i, \Sigma_i)$, and σ is the sparsity weight.

There is space for improvement of the above approach [3] in terms of inference speed. Based on an input image, the iterative optimization of the energy function tends to be slow and running time scales up linearly with the number of pixels, as shown in Section 5.2.

The success of optimization-based methods [3, 12, 24] and its disadvantage in speed inspire us to train neural networks on the dataset. Without any on-the-fly iteration, the networks decompose the original image into soft color layers in a significant lower inference time. Partly inspired by the minimization of the energy function, we come up with an objective function that is designed for training our neural networks system, as detailed in the next section.

4. Methods

Our proposed method consists of three stages: palette color selection, alpha layer estimation and color layer estimation. The input of the proposed system is a single RGB image, and the outputs are RGBA soft color layers.

4.1. Palette Color Selection

The inputs of the methods of both Aksoy *et al.* and Koyama *et al.* [3, 12] are a set of color models that represent the means and the (co)variances of colors of the desired

output layers. Although covariances provide additional controllability, a user has to understand the definition of the covariance of colors, how to adjust it, and how it interacts with the system to produce the final color layers. Consequently, the user may not find it intuitive to manipulate covariances. Aiming for an easy-to-use user experience, we believe that a user should not be exposed to more parameters than necessary. Therefore, we make the design choice to take palette colors (means only) instead of color models (means + covariances) as inputs. A palette color has the simple interpretation as the mean value of the colors that should be included in a color layer.

During training, we use K-means algorithm to partition pixels in an input image into K clusters in the 3-dimensional RGB space. We pick the center RGB values of the clusters as the palette colors. The number of palette colors K is fixed throughout the training of the networks. During inference, the palette colors can be specified manually, in addition to automatic selection using K-means.

4.2. Alpha Layer Estimation

We adopt the U-Net architecture [20] for our *alpha predictor*. The inputs of the alpha predictor are the original RGB image and K single-color images. Each of the single-color image is simply an RGB image filled with a single palette color. The outputs are K alpha layers, which are single-channel images of alpha values. In the subsequent *alpha layer processing* in Figure 2, the outputs of the network are normalized by:

$$\alpha_i = \frac{\alpha_i}{\sum_k \alpha_k}, \quad (5)$$

where α_i is the alpha value (opacity) at a certain pixel position. This normalization step ensures that the output satisfies the alpha-add condition (Eq. (2)). For inference, we can add various kinds of alpha layer processing in addition to normalization, which we detail in Section 4.4.

4.3. Color Layer Estimation

The palette colors and the alpha layers are not sufficient to reconstruct the original image, as shown in Figure (3). Although a user only needs to specify the palette colors, each color layer should contain a more variety of colors than a single palette color.

To introduce color variations, we add a *residue predictor* to estimate color residues from the palette colors. The inputs of the residue predictor are the original image, K single-color images, and K normalized alpha layers. The residue predictor network has an identical architecture as the alpha predictor, except for the number of input and output channels. We add the output residues to palette colors to compute K RGB layers. The final RGBA layers can be obtained by concatenating RGB layers and normalized alpha layers along the channel axis.

4.4. Network Training

The above two networks are trained end-to-end with an objective composed of several losses. The main training objective, inspired by the color constraint (Eq. (1)), is to minimize a self-supervised reconstruction loss between the input and the output:

$$L_r = \left\| \sum_i \alpha_i \mathbf{u}_i - \mathbf{c} \right\|_1. \quad (6)$$

To regularize the training of the alpha predictor, we propose a novel loss for regularization, formulated as the reconstruction loss between the original image and the reconstructed image *without* color residues:

$$L_a = \left\| \sum_i \alpha_i \mathbf{p}_i - \mathbf{c} \right\|_1, \quad (7)$$

where \mathbf{p}_i denotes the palette color of the i^{th} layer. In other words, $\sum_i \alpha_i \mathbf{p}_i$ is an image that is reconstructed only using the normalized alpha layers and the palette colors, as shown in the middle left of Figure (3).

To gather only homogeneous colors in each color layer, we propose a novel distance loss, in reminiscence of Eq. (4), formulated as

$$L_d = \sum_i \alpha_i \|\mathbf{p}_i - \mathbf{u}_i\|_2. \quad (8)$$

We use Euclidean distance in RGB space because our inputs are simply K palette colors.

Now we are ready to formulate our total loss as follows:

$$L_{\text{total}} = L_r + \lambda_a L_a + \lambda_d L_d, \quad (9)$$

where λ_a and λ_d are coefficients for regularization loss and distance loss, respectively. In comparison with the method of Aksoy *et al.* [3] which minimizes the proposed energy function for a given input image, our method trains neural networks to minimize the total loss on a training dataset.

We note that 1) the outputs of the networks automatically satisfy the box constraint (Eq. (3)) because of the sigmoid functions and clip operations appended to the networks. And 2) we do not enforce sparsity like the sparsity term in sparse color unmixing energy function (Eq. (4)). In preliminary experiments we introduced such a sparsity loss with a coefficient to control its weight in the total loss. When its weight is high, there are no soft transition on the boundary of regions, resulting in nearly hard segmentation. When we decrease the weight, however, the sparsity loss no longer promotes sparsity. We believe that the novel regularization loss L_a and the normalization step in alpha layer processing (Eq. (5)) have collectively encouraged the sparsity of alpha layers, and therefore it is redundant to introduce an extra sparsity loss.

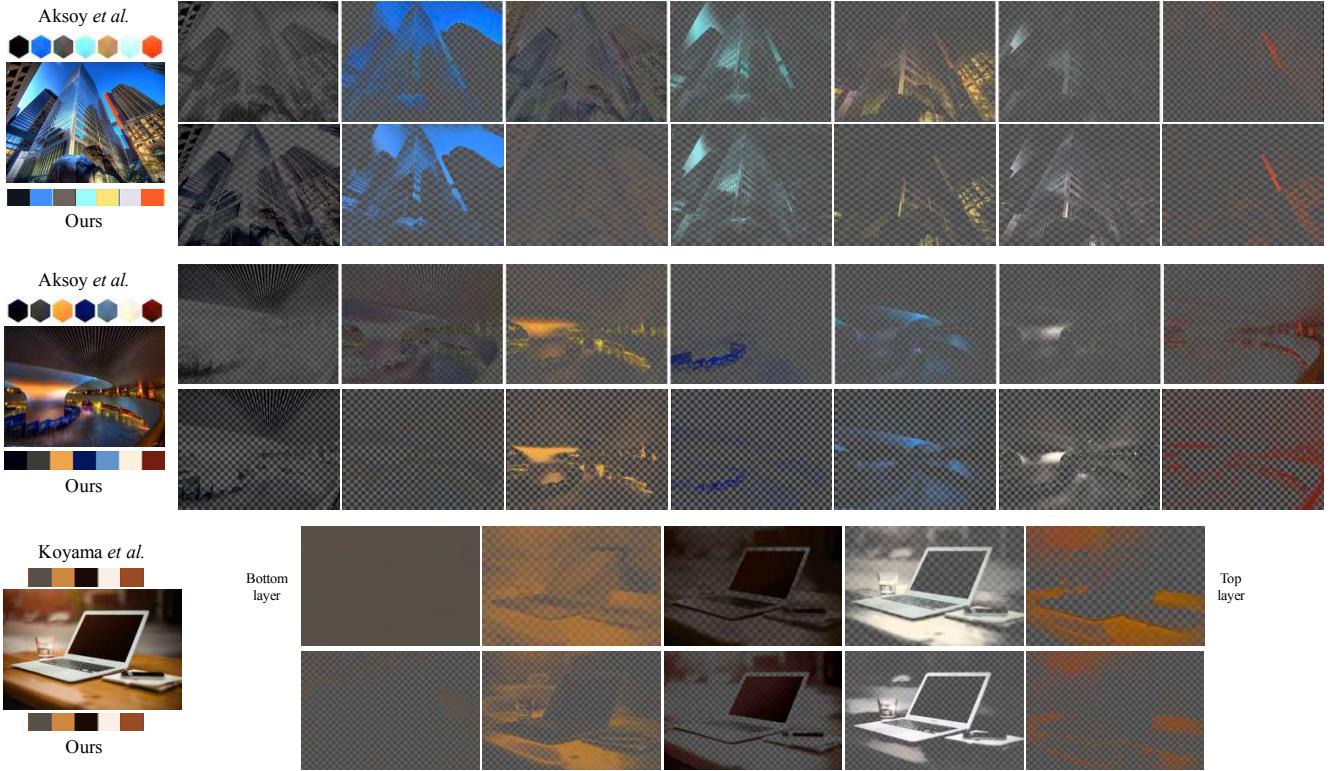


Figure 4: Qualitative comparison with the state-of-the-art baselines [3, 12]. Our method performs comparably to the state-of-the-art methods [3] in producing meaningful soft color layers. We note that Koyama *et al.* [12] deals with order-dependent blending while we are order-independent.

5. Experiments

For training, we use Places365-Standard validation images [36]. All images in the training dataset have a resolution of 256×256 . We use the Adam optimizer with $lr = 0.2$, $\beta_1 = 0.0$, and $\beta_2 = 0.99$. For all experiments, we set $\lambda_a = 1$ and $\lambda_d = 0.5$.

For inference, since our networks are fully-convolutional, we can apply the model to decompose images of various resolutions. Both of the networks are U-Net with five convolutions, three transposed convolutions, and skip connections at each scale. We discuss the network structures in detail in the supplementary material.

5.1. Qualitative Evaluation

Comparison with state-of-the-art methods. Figure 4 shows a qualitative comparison between our results and results of Aksoy *et al.* [3] or Koyama *et al.* [12]. Note that the inputs are different: their inputs are distributions of colors, while ours are simply colors.

We would like to mention that there is no optimal solution for soft color segmentation, because an image can be decomposed into meaningful and high-quality color layers

in various ways. As an attempt to make the comparison as intuitive as possible, we manually select palette colors for our method so that the result color layers look similar to the result in Aksoy *et al.* and Koyama *et al.* [3, 12].

5.2. Quantitative Evaluation

Speed. Figure 5 shows the running time of our algorithm, algorithm of Aksoy *et al.* and Tan *et al.* [3, 23], depending on the size of the input image. Both Aksoy *et al.* [3] and our method use a palette size of 7, and Tan *et al.* [23] use an average palette size of 6.95, and the median palette size is 7. It takes 2.9 ms for the proposed method to decompose a 1080p resolution (about 2 MP) image, 3.0 ms for a 5 MP image and 3.4 ms for a 4K resolution (about 8 MP) image, averaged from 20 experiments at each resolution. Aksoy *et al.* [3] reports that their algorithm takes 1,000 s to decompose a 5 MP image. Their running time scales up linearly as the number of pixels grows due to the per-pixel optimization. Tan *et al.* [23] improve the execution speed of Aksoy *et al.*, utilizing RGBXY convex hull computation and layer updating. We collect their running time (about 50 s to 100 s) in Figure 9 in their paper [23] because exact values are not reported. See supplementary

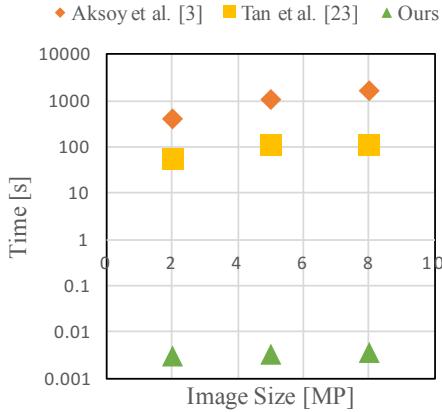


Figure 5: Speed comparison with the state-of-the-art methods. Note that the time axis is in logarithmic scale.

material for detail experimental settings.

The comparison shows that our neural network based method has a significant speed improvement over the state-of-the-art methods. Furthermore, only our method can decompose a video in a practical amount of time. Specifically, when the video consists of 450 frames (1080p, 30 fps, 15 s), our method takes 1.35 seconds, while the method of Aksoy *et al.* [3] takes around 50 hours.

Reconstruction error. In Table 1, We use pixel-level mean squared error to evaluate the difference between an input image and the reconstructed image. Although we minimize the reconstruction loss (Eq. (6)) on a training dataset instead of the input image as Aksoy *et al.* [3] do, our reconstruction errors are sufficiently low, indicating that our networks generalize well on the test images. We note that the reconstruction error also depends on the palette colors selected. In particular, the error increases if some palette colors never appear in the original image at all.

5.3. Ablation Study

In this section, we validate the losses and architectures of the neural networks. Figure 6 shows a sample for qualitative comparison. The reconstruction error and the sparsity scores shown in Table 1 are the averaged scores of 100 images of 1 MP or more. The sparsity score is calculated as

$$L_s = \frac{\sum_i \alpha_i}{\sum_i \alpha_i^2} - 1. \quad (10)$$

A lower value of L_s means the decomposed layers are sparser. The color variance score in Table 2 represents the score corresponding to the results in Figure 6. This score is the sum of individual variances of the RGB channel averaged over all decomposed layers. For fair comparison, we used same palette colors to decompose an input for each setting.

Method	Reconst. \downarrow	PSNR \uparrow	SSIM \uparrow	Sparsity \downarrow
Aksoy <i>et al.</i> [3]	0.00050	-	-	-
Ours	0.00088	31.07	0.9740	1.456
w/o L_r	0.00308	25.70	0.9158	1.279
w/o L_a	0.00090	31.17	0.9750	1.959
w/o L_d	0.00076	31.72	0.9743	1.640
w/o Skip	0.00350	27.37	0.9366	1.149
w/o Zero-centered	0.00073	31.82	0.9710	1.450
Single network	0.00104	30.71	0.9633	NaN

Table 1: Quantitative comparison. Our method is the only setting in the ablation study that achieves similar reconstruction error as Aksoy *et al.* [3] with both low sparsity score and high image quality. See 5.2 for details of the quantitative evaluation, and 5.3 for details of the ablation study.

Ours versus ours without L_r . Table 1 and Figure 6 show that our method without L_r cannot properly reconstruct an input image. L_a trains the only *alpha predictor*, so the *residue predictor* cannot function properly.

Ours versus ours without L_a . Although the SSIM score is marginally better without L_a , the sparsity score is significantly higher, which suggests that excessive overlapping exists between alpha layers. On the third row in Figure 6, the blue plate on the person’s hand wrongly shows up in the white layer (highlighted in red), causing overlapping between white and blue layers. Overlapping is intended to occur only sparsely, e.g. at the boundary of a region, because excessive overlapping is not suitable for application to image editing.

Moreover, the color variance score is higher without L_a , indicating that some of the layer might be contaminated, i.e. containing colors that is much different from the corresponding palette color. We can observe such contamination on the top-left corner of the blue layer. We believe this is because L_a improves the performance alpha predictor.

Ours versus ours without L_d . Without L_d suppressing the variance of colors in each layer, the reconstruction error decreases. However, a large color variance causes the same color to spread across multiple layers, which is not desirable for color-based editing, as shown in Table 2 and Figure 6.

Ours versus ours without skip connection. Without skip connections in the *alpha predictor* and the *residue predictor*, alpha layers are not accurate, leading to a higher the reconstruction error.

Ours versus ours without zero-centered residues. It is easier to train a neural network with zero-centered output, and we can make sure the palette color is the mean value of each layer. Without zero-centering, the PSNR increases, but the color variance increases even more.

Ours versus a single network only. If we use a plain single network, the reconstruction error increases. We be-

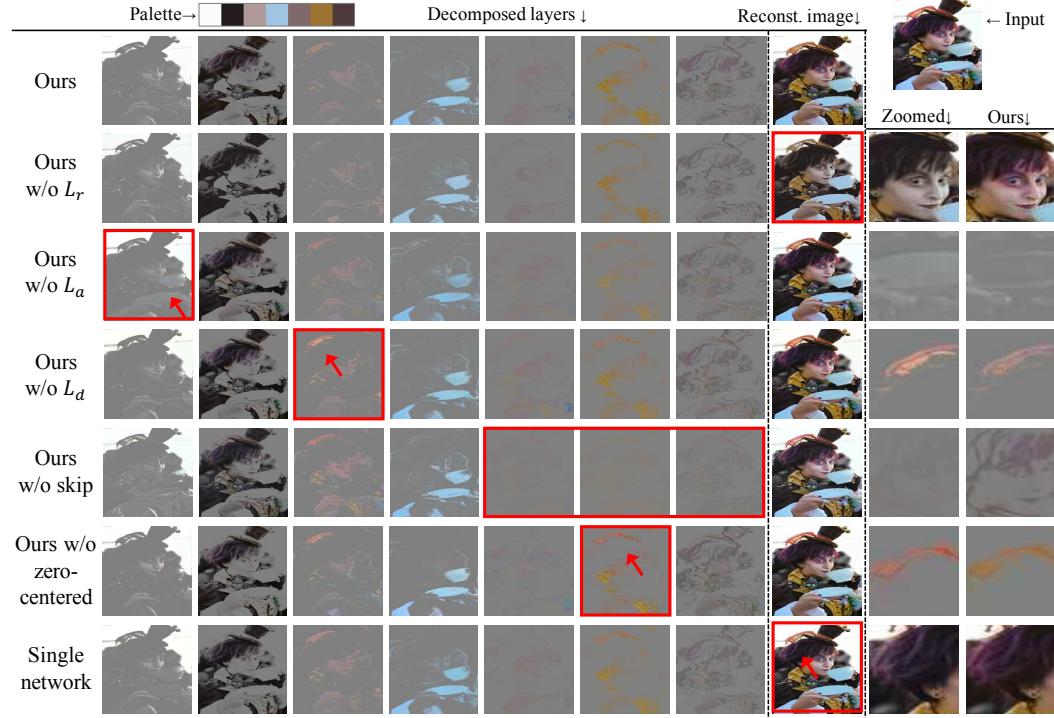


Figure 6: Visual comparison for ablation study. Removing any component of our method results in higher reconstruction error, color contamination or unnecessary overlapping of color layers. See section 5.3 for details.

Method	Aksøy <i>et al.</i> [3]	Ours	w/o L_a	w/o L_d	w/o Zero-centered
Color Var. ↓	0.005	0.003	0.007	0.006	0.162

Table 2: Quantitative comparison of color variance. The values of our methods correspond to the images in Figure 6. We obtain the color variance of the method of Aksøy *et al.* from their paper [3].

lieve that accurate alpha layers as inputs enhance the performance of the residue predictor. Specifically, in our method, we apply smoothing filters to remove checkerboard artifacts from alpha layers, as shown in Figure 11, and predict the RGB channels based on the processed alpha layers. We doubt that there is a way to incorporate smoothing filter processing in a neural network that predicts both alpha and RGB channels simultaneously.

5.4. Applications

Decomposition of images with blurry boundary. Soft color segmentation is useful not only for decomposing an image into soft color layers, but for any case where the resulting layers are preferred to have blurry boundary. Figure 7 shows the results of our decomposition and the benefits of soft color segmentation. In the case of an image with a foreground object with motion blur, although both the foreground and the background are visible at a single



Figure 7: Comparison with hard segmentation and color-based soft segmentation on images with ambiguous object boundaries. In (a), hard segmentation produces a background layer that is tainted by the moving object in the foreground. In (b), hard semantic segmentation divides the image into semantic areas, but soft color segmentation enable us to change simultaneously the reflected light on the rock along with the lighting condition of the sky.

pixel, hard segmentation has to assign that pixel into either of the classes. In such a case, soft segmentation can recognize the pixel as a mixture of fore/background and thus has

an advantage over conventional hard segmentation.

Natural image editing. Figure 1 and Figure 8 show examples of recoloring and compositing. These editing results are created by editing each decomposed layer with the *alpha add* mode in Adobe After Effects.

Video decomposition. Figure 9 shows our method performing video soft color segmentation. We decompose the video frame-by-frame, without any constraints for temporal consistency. Nevertheless, the decomposed layers do not flicker. It can be partly attributed to applying smoothing filter and fixing color palettes, both of which encourage consistent alpha layers, and regularization loss L_d , which encourages consistent color layers. Compared to other methods, only our method can decompose video in practical time. See the video material for detailed results.

Alpha layer processing. In the *alpha layer processing* stage, a user can edit the predicted alpha layers, and subsequently use these edited alpha layers for color estimation, thanks to the fact that the estimation of alpha layers is independent from color layers. As shown in Figure 10, we can use the guided filter [11, 28] to smooth the image, or manipulate a mask to change the alpha region. We can prepare the mask manually or automatically, capitalizing on state-of-the-art methods (e.g. semantic segmentation and depth estimation). Therefore, our method can be complementary to various image editing techniques.

5.5. Limitations

Memory limit. Because we use a GPU for inference, we cannot handle high-resolution images that exceed the GPU memory limit. Also, sufficient GPU memory is needed to hold the intermediate features of the encoder-decoder networks with skip connections. Concretely, when an input is 1080p (4K) resolution and computation is based on 32-bit floating point, about 12 (21) GB is consumed. Countermeasures include using 16-bit floating point and discarding unnecessary intermediate features.

Fixed palette size. In our approach, the number of decomposed layers for each trained model is fixed. To handle various numbers of layers, one solution is to train a model that decomposes an image into the sufficient number of layers, and use palettes with duplicated colors, as shown in Figure 11. We merge layers of duplicated colors after decomposition.

6. Conclusion

In this paper, we tackled the problem of soft color segmentation using the first neural network based approach. We showed that the proposed method enabled faster decomposition relative to prior state-of-the-art methods. We also demonstrated that our fast and simple approach enabled new important applications, such as soft color segmentation of videos.



Figure 8: Examples of image editing using the layers decomposed by our method.

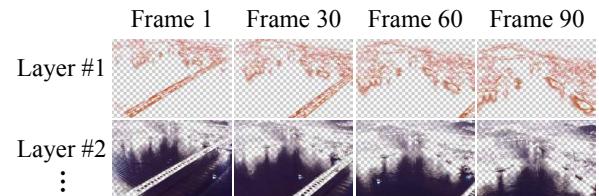


Figure 9: Examples of video decomposition. Our method decomposes all the frames in the video in a few seconds.

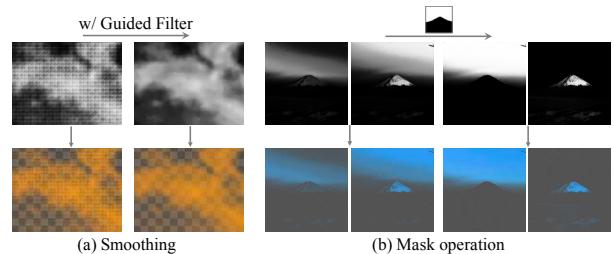


Figure 10: Examples of alpha layer processing before predicting colors. (a) Comparison between with/without smooth filtering. (b) An example of using a mask to manipulate the decomposed layers.

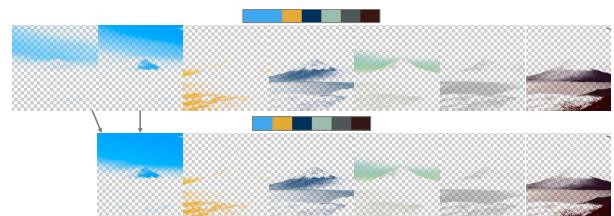


Figure 11: An example when the palette contains a duplicated color. It shows that, although the number of layers is fixed during training, our model tolerates the duplication of colors and is thus capable of decomposing an image into fewer layers.

Acknowledgement

We are grateful for Yingtao Tian and Prabhat Nagarajan for helpful advice on writing the paper. We thank Edgar Simo-Serra and all reviewers for helpful discussions and comments.

References

[1] Yağız Aksoy, Tunç Ozan Aydin, and Marc Pollefeys. Designing effective inter-pixel information flow for natural image matting. In *The IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, 2017.

[2] Yağız Aksoy, Tunç Ozan Aydin, Marc Pollefeys, and Aljoša Smolić. Interactive high-quality green-screen keying via color unmixing. *ACM Trans. Graph.*, 35(5):152:1–152:12, 2016.

[3] Yağız Aksoy, Tunç Ozan Aydin, Aljoša Smolić, and Marc Pollefeys. Unmixing-based soft color segmentation for image manipulation. *ACM Trans. Graph.*, 36(2):19:1–19:19, 2017.

[4] Yağız Aksoy, Tae-Hyun Oh, Sylvain Paris, Marc Pollefeys, and Wojciech Matusik. Semantic soft segmentation. *ACM Trans. Graph. (Proc. SIGGRAPH)*, 37(4):72:1–72:13, 2018.

[5] Liang-Chieh Chen, George Papandreou, Iasonas Kokkinos, Kevin Murphy, and Alan L Yuille. Deeplab: Semantic image segmentation with deep convolutional nets, atrous convolution, and fully connected crfs. *IEEE Transactions on Pattern Analysis and Machine Intelligence*, 40(4):834–848, 2017.

[6] Qifeng Chen, Dingzeyu Li, and Chi-Keung Tang. Knm matting. *IEEE transactions on pattern analysis and machine intelligence*, 35(9):2175–2188, 2013.

[7] Qingnan Fan, Jiaolong Yang, Gang Hua, Baoquan Chen, and David Wipf. A generic deep architecture for single image reflection removal and image smoothing. In *IEEE International Conference on Computer Vision (ICCV)*, 2017.

[8] Jean-Dominique Favreau, Florent Lafarge, and Adrien Bousseau. Photo2clipart: Image abstraction and vectorization using layered linear gradients. *ACM Transactions on Graphics (SIGGRAPH Asia Conference Proceedings)*, 36(6), November 2017.

[9] Xueyang Fu, Jiabin Huang, Delu Zeng, Yue Huang, Xinghao Ding, and John Paisley. Removing rain from single images via a deep detail network. In *The IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, pages 1715–1723, July 2017.

[10] Yosef Ganelsman, Assaf Shocher, and Michal Irani. “double-dip”: Unsupervised image decomposition via coupled deep-image-priors. In *The IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, June 2019.

[11] Kaiming He, Jian Sun, and Xiaou Tang. Guided image filtering. In *European conference on computer vision*, pages 1–14. Springer, 2010.

[12] Yuki Koyama and Masataka Goto. Decomposing images into layers with advanced color blending. *Comput. Graph. Forum*, 7(37):pp.397–407, 2018.

[13] Anat Levin, Alex Rav-Acha, and Dani Lischinski. Spectral matting. *IEEE Transactions on Pattern Analysis and Machine Intelligence*, 30(10):1699–1712, 2008.

[14] Ruoteng Li, Loong-Fah Cheong, and Robby T. Tan. Heavy rain image restoration: Integrating physics model and conditional adversarial learning. In *The IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, June 2019.

[15] Xiaodan Liang, Si Liu, Xiaohui Shen, Jianchao Yang, Luoqi Liu, Jian Dong, Liang Lin, and Shuicheng Yan. Deep human parsing with active template regression. *IEEE Transactions on Pattern Analysis and Machine Intelligence*, 37(12):2402–2414, Dec 2015.

[16] Chenxi Liu, Liang-Chieh Chen, Florian Schroff, Hartwig Adam, Wei Hua, Alan L Yuille, and Li Fei-Fei. Auto-deeplab: Hierarchical neural architecture search for semantic image segmentation. In *The IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, pages 82–92, 2019.

[17] Jonathan Long, Evan Shelhamer, and Trevor Darrell. Fully convolutional networks for semantic segmentation. In *The IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, pages 3431–3440, 2015.

[18] Jinshan Pan, Zhe Hu, Zhixun Su, Hsin-Ying Lee, and Ming-Hsuan Yang. Soft-segmentation guided object motion deblurring. In *The IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, pages 459–468, 2016.

[19] Wenqi Ren, Si Liu, Hua Zhang, Jinshan Pan, Xiaochun Cao, and Ming-Hsuan Yang. Single image dehazing via multi-scale convolutional neural networks. In *European Conference on Computer Vision (ECCV)*, 2016.

[20] Olaf Ronneberger, Philipp Fischer, and Thomas Brox. U-net: Convolutional networks for biomedical image segmentation. In *International Conference on Medical image computing and computer-assisted intervention*, pages 234–241. Springer, 2015.

[21] Othman Sbai, Camille Couprie, and Mathieu Aubry. Vector image generation by learning parametric layer decomposition. *arXiv preprint arXiv:1812.05484*, 2018.

[22] Yu-Wing Tai, Jiaya Jia, and Chi-Keung Tang. Soft color segmentation and its applications. *IEEE Transactions on Pattern Analysis and Machine Intelligence*, 29(9):1520–1537, Sep. 2007.

[23] Jianchao Tan, Jose Echevarria, and Yotam Gingold. Efficient palette-based decomposition and recoloring of images via rgbxy-space geometry. *ACM Transactions on Graphics (TOG)*, 37(6):262:1–262:10, Dec. 2018.

[24] Jianchao Tan, Jyh-Ming Lien, and Yotam Gingold. Decomposing images into layers via RGB-space geometry. *ACM Transactions on Graphics (TOG)*, 36(1):7:1–7:14, Nov. 2016.

[25] Jingwei Tang, Yağız Aksoy, Cengiz Öztireli, Markus Gross, and Tunç Ozan Aydin. Learning-based sampling for natural image matting. In *The IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, 2019.

[26] Yu Wang, Yi Niu, Peiyong Duan, Jianwei Lin, and Yuanjie Zheng. Deep propagation based image matting. In *IJCAI*, pages 999–1006, 2018.

[27] Qiang Wen, Yinjie Tan, Jing Qin, Wenxi Liu, Guoqiang Han, and Shengfeng He. Single image reflection removal beyond linearity. In *The IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, June 2019.

[28] Huikai Wu, Shuai Zheng, Junge Zhang, and Kaiqi Huang. Fast end-to-end trainable guided filter. *The IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, pages 1838–1847, 2018.

[29] Ning Xu, Brian Price, Scott Cohen, and Thomas Huang. Deep image matting. In *The IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, July 2017.

- [30] Lu Yang, Qing Song, Zhihui Wang, and Ming Jiang. Parsing r-cnn for instance-level human analysis. In *The IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, June 2019.
- [31] He Zhang and Vishal M. Patel. Density-aware single image de-raining using a multi-stream dense network. In *The IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, 2018.
- [32] Jing Zhang, Yang Cao, Shuai Fang, Yu Kang, and Chang Wen Chen. Fast haze removal for nighttime image using maximum reflectance prior. In *The IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, July 2017.
- [33] Xuaner Zhang, Ren Ng, and Qifeng Chen. Single image reflection separation with perceptual losses. In *The IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, 2018.
- [34] Yunke Zhang, Lixue Gong, Lubin Fan, Peiran Ren, Qixing Huang, Hujun Bao, and Weiwei Xu. A late fusion cnn for digital matting. In *The IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, pages 7469–7478, 2019.
- [35] Hengshuang Zhao, Jianping Shi, Xiaojuan Qi, Xiaogang Wang, and Jiaya Jia. Pyramid scene parsing network. In *The IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, pages 2881–2890, 2017.
- [36] Bolei Zhou, Agata Lapedriza, Aditya Khosla, Aude Oliva, and Antonio Torralba. Places: A 10 million image database for scene recognition. *IEEE Transactions on Pattern Analysis and Machine Intelligence*, 2017.
- [37] Qingsong Zhu, Jiaming Mai, and Ling Shao. A fast single image haze removal algorithm using color attenuation prior. *IEEE Transactions on Image Processing*, 24(11):3522–3533, Nov 2015.