

Separate in Latent Space: Unsupervised Single Image Layer Separation

Yunfei Liu,¹ Feng Lu^{1,2*}

¹State Key Laboratory of VR Technology and Systems, School of CSE, Beihang University, Beijing, China

²Peng Cheng Laboratory, Shenzhen, China
{lyunfei, lufeng}@buaa.edu.cn

Abstract

Many real world vision tasks, such as reflection removal from a transparent surface and intrinsic image decomposition, can be modeled as single image layer separation. However, this problem is highly ill-posed, requiring accurately aligned and hard to collect triplet data to train the CNN models. To address this problem, this paper proposes an unsupervised method that requires no ground truth data triplet in training. At the core of the method are two assumptions about data distributions in the latent spaces of different layers, based on which a novel unsupervised layer separation pipeline can be derived. Then the method can be constructed based on the GANs framework with self-supervision and cycle consistency constraints, *etc.* Experimental results demonstrate its successfulness in outperforming existing unsupervised methods in both synthetic and real world tasks. The method also shows its ability to solve a more challenging multi-layer separation task.

1 Introduction

Real imaging process is influenced by different factors, *e.g.*, shape, illumination, reflection and refraction, making the captured images highly complex. Such images, when used as input data to an AI system, may heavily affect the system performance on various vision-based tasks, such as object detection (Girshick 2015), makeup removal (Cao et al. 2019), rain drop removal (Hao et al. 2019) and 3D reconstruction (Tatarchenko et al. 2019). As a result, separating different visual attributes from a single image has been considered an important research goal for decades.

This paper considers visual attributes separation via Single Image Layer Separation (SILS). It aims at decomposing an image into several independent layers with specific physical meanings. In general, this problem can be formulated as:

$$I = L_1 \oplus L_2 \oplus \dots, \quad (1)$$

where L_i is the i -th layer and I is the input image that can be reconstructed by a pixel-wise blend of all the layers. Plenty of real world problems can be represented in this

form. As shown in Fig. 1, for reflection separation, the input image I can be decomposed into a reflection layer L_R and a background layer L_B in the form of $I = L_R + L_B$. While for intrinsic decomposition, the input image I can be represented by the pixel-wise product of an albedo image L_A and a shading image L_S , or alternatively formulated as $\log(I) = \log(L_A) + \log(L_S)$.

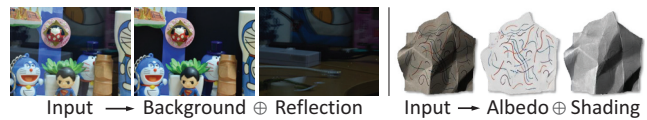


Figure 1: Examples of single image layer separation: reflection separation and intrinsic image decomposition.

Despite its wide applicability, single image layer separation has been found fundamentally ill-posed, due to the existence of countless feasible solutions. To constrain the solution space, conventional methods use different priors such as low rank (Han and Sim 2017), sparsity (Li and Brown 2013) and relative smoothness (Li and Brown 2014) in optimization. Such hand-crafted constraints may work well in specific cases but fail in more complex cases. Recently, the deep convolutional neural network (CNN) has been widely used to handle similar problems (Fan et al. 2018; Yang et al. 2018; Zhang, Ng, and Chen 2018). However, CNN-based methods face major difficulties in high-quality data acquisition: 1) it is difficult or even impossible to obtain the ground truth of each layer for real images, and 2) different datasets vary greatly and cannot be used jointly. Since the performance of CNNs relies heavily on the training data, existing methods have to either use synthetic data with limited reality or train their networks with imperfect data.

In order to overcome such difficulties, this paper suggests *not requiring the ground truth data triplet, i.e.*, one input image and its exact two layers as training data, but simply collecting images from the three domains independently. In this manner, sufficient training data can be easily obtained without the acquisition problem. The absence of ground truth triplet, however, makes it an *unsupervised learning problem*. Therefore, we propose an unsupervised single image layer separation (unsupervised SILS) method in this paper.

*Corresponding Author: Feng Lu.

This work is partially supported by the National Natural Science Foundation of China (NSFC) under Grant 61972012 and Grant 61732016.

Copyright © 2020, Association for the Advancement of Artificial Intelligence (www.aaai.org). All rights reserved.

The proposed unsupervised SILS method does not rely on training data triplets from the three domains. Instead, we introduce two key assumptions, namely the shared information consistency and the layer independence, to help link the three different domains and construct the latent spaces. Based on the assumptions, we design the method by using techniques such as self-supervision, cycle consistency and GANs. The full method is able to perform unsupervised SILS without requiring ground truth triplets in training.

We test our method on two classical tasks, *i.e.*, single image reflection separation and intrinsic image decomposition, and it outperforms state-of-the-art unsupervised methods. As our method does not need ground truth triplets, it can be quickly applied to other layer separation tasks with little data collection burden. Finally, our method can be extended to solve the more challenging multi-layer separation task.

1.1 Related Works

The problem of image layer separation can be considered as a special case of image domain translation. Previous domain translation can be formulated as a mapping function between a source and a target domain (Radford, Metz, and Chintala 2016; Yi et al. 2017; Zhu et al. 2017; Liu, Breuel, and Kautz 2017), while single image separation is a mapping from the source domain to two other domains physically related to the source domain. There are various image separation problems in the field of computer vision where different physical models are applied as priors for separation tasks. For instance, relative smoothness (Li and Brown 2014), ghost cues (Shih et al. 2015) and layer independence priors (Li and Brown 2013) are introduced for separating and removing the reflection on the glass surface from the background. Bi *et al* (Bi, Han, and Yu 2015) applied sparse prior to tackle image intrinsic decomposition. However, the above mentioned methods can only handle simple cases with image gradient or color changes, and can not be adopted to more complex situations that may require other techniques such as photometric stereo (Lu et al. 2015; 2017).

In recent years, supervised deep learning methods are applied to many image separation tasks. For example, fully convolutional networks with various guidance, *e.g.*, image gradient information (Fan et al. 2017), face structure priors (Wan et al. 2019), feature discriminative prior (Wang and Lu 2019) and perceptual losses (Zhang, Ng, and Chen 2018), are used for single image reflection removal. Besides, Ronneberger *et al.* used U-net-like CNNs for intrinsic image decomposition (Ronneberger, Fischer, and Brox 2015).

Since real shot images with layer separation results are hard to collect, the development of unsupervised algorithms is in great need. Michael *et al* (UII 2017) proposed a self-supervised intrinsic image decomposition method. It can be trained with only a few images with ground-truth, and then transferred to other unpaired images. However, the albedo layer should be the same among different training images. Li *et al* (Li and Snavely 2018) and Lettry *et al* (Lettry, Vanhoey, and Van Gool 2018) proposed unsupervised intrinsic image decomposition methods, but these methods need multiple inputs with the same albedo layer for training. More recently,

Hoshen *et al* (Hoshen) proposed an unsupervised single-channel blind source separation method, while the method requires a pair of input images which limits its applicability. Overall, unsupervised single image layer separation remains an open problem to solve.

2 Problem Formulation and Analysis

Single Image Layer Separation can be considered as a one-to-two image domain translation problem. In particular, let \mathcal{X} , \mathcal{Y} and \mathcal{Z} denote three image domains, then we try to find out a proper separation $x = y \oplus z$. In the supervised case, we have training samples (x, y, z) drawn from a joint distribution $P_{\mathcal{X}, \mathcal{Y}, \mathcal{Z}}(x, y, z)$. While in the unsupervised case, we only have samples drawn from the marginal distributions $P_{\mathcal{X}}(x)$, $P_{\mathcal{Y}}(y)$ and $P_{\mathcal{Z}}(z)$. Since we can infer nothing about the joint distribution only from the marginal samples without additional assumptions, the unsupervised problem is fundamentally ill-posed, as mentioned in Sec. 1. To address this problem, we introduce the following assumptions.

Assumption A: Shared information consistency. Given that $x = y \oplus z$, we assume that there exists a shared latent space \mathcal{F}_y for x and y , *i.e.*, $\forall x, y, E_{\mathcal{Y}}(x) \in \mathcal{F}_y$ and $E_{\mathcal{Y}}(y) \in \mathcal{F}_y$, where $E_{\mathcal{Y}}$ is a function that maps data from the color image space to the latent space \mathcal{F}_y . Similarly, we have $E_{\mathcal{Z}}(x)$ and $E_{\mathcal{Z}}(z)$ both belong to the same latent space \mathcal{F}_z .

Following this assumption (also see Fig. 2 (a)), the original image layer separation pipeline can be reformed as:

$$x \rightarrow \{y, z\} \Rightarrow x \begin{cases} \nearrow f_y \xrightarrow{G_{\mathcal{Y}}} y \\ \searrow f_z \xrightarrow{G_{\mathcal{Z}}} z \end{cases}, \quad (2)$$

where $f_y = E_{\mathcal{Y}}(x) \in \mathcal{F}_y$ and $f_z = E_{\mathcal{Z}}(x) \in \mathcal{F}_z$. Such an f_y can be understood as the extracted information, which only belongs to the domain \mathcal{Y} , from the blended input x . Then, $G_{\mathcal{Y}}$ is another mapping function to project f_y back to the color image space and generate the layer image y .

In conclusion, this assumption allows producing y/z from a single x if we find out correct $E_{\mathcal{Y}}/E_{\mathcal{Z}}$ and $G_{\mathcal{Y}}/G_{\mathcal{Z}}$.

Assumption B: Layer independence. We assume that the separated y and z should be independent to each other in the latent space, *i.e.*, samples belonging to the same domain have more similar features in the latent space.

Computationally, for any two samples f_y^a, f_y^b from \mathcal{F}_y , and any other two samples f_z^a, f_z^b from \mathcal{F}_z , we assume:

$$d_{\psi} \gg d_{\phi}, \forall d_{\psi} \in \{|f_y^a - f_z^a|, |f_y^a - f_z^b|, |f_y^b - f_z^a|, |f_y^b - f_z^b|\} \text{ and } \forall d_{\phi} \in \{|f_y^a - f_y^b|, |f_z^a - f_z^b|\}, \quad (3)$$

where d_{ψ} indicates an inter latent space distance and d_{ϕ} indicates an inner latent space distance. This assumption can be satisfied by minimizing d_{ϕ} while maximizing d_{ψ} .

3 Learning for Unsupervised SILS

Based on the two assumptions in Sec. 2, we propose an unsupervised single image layer separation method. In particular, our method implements self-supervision, cycle consistency,

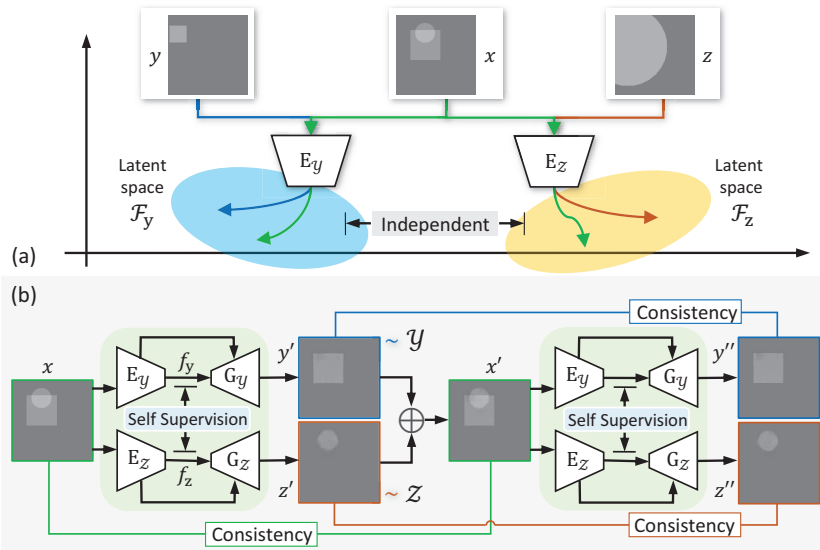


Figure 2: (a) Illustrations of the shared information consistency and layer independence assumptions. E_y , E_z are two encoding functions that map images into different and independent latent spaces (\mathcal{F}_y for square and \mathcal{F}_z for circle). Intuitively speaking, the square features in both x and y go to the latent space \mathcal{F}_y , and the circle features in both x and z go to the latent space \mathcal{F}_z . In this manner, the method learns how to separate information from x . (b) The overall framework of our unsupervised SILS method. The encoder/decoder networks E_y , E_z , G_y and G_z are implemented by using CNNs. Following our two assumptions, we utilize the cycle consistency constraint, and also the self-supervision learning. The \sim indicates where we add adversarial discriminators for evaluating whether the generated images are realistic.

GANs, *etc.* to achieve efficient unsupervised learning. The overall framework of our method can be seen in Fig. 2 (b).

Self Supervision (SS). According to assumption A, x contains all the information to produce f_y and f_z , respectively. Meanwhile, assumption B requires that f_y and f_z should be as different as possible. Therefore, we need to carefully design the encoders $E_y(x)$ and $E_z(x)$ so that their outputs have maximized distances.

To this end, we propose a novel self-supervision framework to optimize both encoders and constrain the feature distributions in both latent spaces. In particular, we maximize the distance between output features from E_y and E_z ; for E_y , we minimize the $L1$ distance between $E_y(x_i)$ and $E_y(y_i)$, where x_i and y_i are unpaired samples randomly selected from \mathcal{X} and \mathcal{Y} ; for E_z , we minimize the $L1$ distance between $E_y(x_j)$ and $E_z(z_j)$, where x_j and z_j are unpaired samples randomly selected from \mathcal{X} and \mathcal{Z} .

Cycle Consistency (CC). Assumption A describes how to produce y and z from the input x , while there should be one more step to close the loop: that y and x should re-produce x . This is physically obvious and leads to the cycle consistency.

As shown in Fig. 2 (b), let y' and z' be the output of G_y and G_z , then we can simply fuse them again to produce $x' = y' \oplus z'$. Here we introduce the first cycle consistency constraint that enforces $x = x'$. Moreover, the new x' can be further separated into new y'' and z'' , resulting the second cycle consistency constraint of $y' = y''$ and $z' = z''$. In this manner, the physical correctness of the image layer separation & fusion processes is guaranteed.

GANs. Assumption A and B mainly focus on the sepa-

rability of different layers, while there is no guarantee that the separated images look realistic. Therefore, we adopt the widely successful Generative Adversarial Networks (GANs) in our framework. The aim is to ensure the output y and z show realistic visual appearances as real images in their respective domains \mathcal{Y} and \mathcal{Z} .

The GANs structure comprises generators and discriminators. For the former, we have $E_y \circ G_y$ and $E_z \circ G_z$ work as the generator to produce layer images y and z , as shown in Fig. 2 (b). As for the latter, we add two discriminators D_y and D_z . They are used to distinguish the outputs of the generators from the real images in domains \mathcal{Y} and \mathcal{Z} .

Loss Function Design. According to the techniques introduced above, we propose the detailed loss functions as following. The overall loss function is designed as:

$$\begin{aligned} & \min_{E, G} \max_D \mathcal{L}_{SS}(E_y, E_z) + \mathcal{L}_{GAN_{\mathcal{Y}}}(E_y, G_y, D_y) \\ & + \mathcal{L}_{GAN_{\mathcal{Z}}}(E_z, G_z, D_z) + \mathcal{L}_{CC_x}(E_y, G_y, E_z, G_z) \quad (4) \\ & + \mathcal{L}_{CC_y}(E_y, G_y) + \mathcal{L}_{CC_z}(E_z, G_z). \end{aligned}$$

The first part is related to the self-supervised training. Based on Equ. (3), the objective function can be designed as:

$$\begin{aligned} \mathcal{L}_{SS}(E_y, E_z) = & \lambda_1 d_\phi(E_y(y), E_y(x)) \\ & + \lambda_2 d_\phi(E_z(z), E_z(x)) \quad (5) \\ & + \lambda_3 (1 - d_\psi(E_y(x), E_z(x))), \end{aligned}$$

where the hyper-parameters λ_1 , λ_2 and λ_3 control the weights of different objective terms. We adopt $d_\phi(a, b) = |a - b|$ in practice, *i.e.*, we compute the $L1$ distance of two inputs a, b . Besides, d_ψ measures the distance between the latent codes

which are the outputs of E_Y and E_Z , and we use a modified sigmoid function to compute the distance:

$$d_\psi(a, b) = \frac{1}{1 + e^{g(a, b)}}, \text{ where } g(a, b) = -\frac{|a - b| - \alpha e^\alpha}{\alpha^2}. \quad (6)$$

Here, α controls the shape of the distance curve. Its effect is shown by the experimental results in the top of Fig. 3 (b).

The GANs-related objective functions are as below:

$$\begin{aligned} \mathcal{L}_{GAN_Y}(E_Y, G_Y, D_Y) &= \lambda_0 \mathbb{E}_{y \sim P_Y} [\log D_Y(y)] \\ &+ \lambda_0 \mathbb{E}_{f_y \sim E_Y(f_y|x)} [\log(1 - D_Y(G_Y(f_y)))], \\ \mathcal{L}_{GAN_Z}(E_Z, G_Z, D_Z) &= \lambda_0 \mathbb{E}_{z \sim P_Z} [\log D_Z(z)] \\ &+ \lambda_0 \mathbb{E}_{f_z \sim E_Z(f_z|x)} [\log(1 - D_Z(G_Z(f_z)))]. \end{aligned} \quad (7)$$

The above functions are in the form of conditional GANs objective functions. The only hyper-parameter λ_0 controls the impact of the entire GAN objective functions.

Finally, for the cycle-consistency constraint, the loss functions are based on the $L1$ differences:

$$\begin{aligned} \mathcal{L}_{CC_X} &= \lambda_4 |y' + z' - x|, \\ \mathcal{L}_{CC_Y} &= \lambda_5 |G_Y(E_Y(y' + z')) - y'|, \\ \mathcal{L}_{CC_Z} &= \lambda_6 |G_Z(E_Z(y' + z')) - z'|, \end{aligned} \quad (8)$$

where $y' = G_Y(E_Y(x))$ and $z' = G_Z(E_Z(x))$. The hyper-parameters λ_4 , λ_5 and λ_6 control the weights of these terms.

4 Implementation via CNNs

Our unsupervised SILS network consists of a two-stream generator and two discriminators, as shown in Fig. 2 (b). The generator is composed of two encoder/decoder pairs, *i.e.*, $E_Y \circ G_Y$ and $E_Z \circ G_Z$, to generate y and z from a single x . The discriminators determine whether the generated y and z look real enough. Below are the implementation details.

Due to the different tasks in our experiments, we design two types of generators in practice.

Generator for synthetic data. The synthetic data is rendered by using simple square and circle shapes, as shown in Fig. 3 (a). Therefore, the generator is designed to be relatively simple. The structures of $E_Y \circ G_Y$ and $E_Z \circ G_Z$ are the same, but they do not share parameters. In detail, each of the encoders E_Y and E_Z has five convolutional layers, which have $\{16, 32, 64, 128, 256\}$ filters of the size 4×4 and a stride of 2. Batch normalization (Ioffe and Szegedy 2015) and leaky ReLU activation are applied after every convolutional layer. Each of the decoders G_Y and G_Z has the same layers as the encoder but in a reverse order plus a final layer with 3 channels. The corresponding layers of the encoder and decoder have mirror-link connections as described in (Ronneberger, Fischer, and Brox 2015), yielding sharper results.

Generator for real image separation tasks. In order to handle real world tasks, we propose a more complex generator as an alternative. Overall, $E_Y \circ G_Y$ and $E_Z \circ G_Z$ still share the same structure, but not the parameters. In detail, the encoders E_Y and E_Z are implemented by using VGG-19 (Simonyan and Zisserman 2015). The decoders G_Y and G_Z have 4 convolution blocks in each, and each block comprises a convolution layer, a leaky ReLU activation layer and

an up-sample layer. The last block is followed by a fully convolutional network with 64 filters of size 3×3 , a stride of 1 and dilation ratios of $\{2, 4, 8, 16, 32, 1\}$. Instance normalization (Dmitry Ulyanov and Lempitsky 2016) and leaky ReLU activation are applied after each convolution layer. The output layer is a convolution layer with three 1×1 filters.

For the skip connection between the encoder and decoder, we select ‘conv1_2’, ‘conv2_2’ and ‘conv3_2’ layers in the encoder and make skip connections to the first block of the decoder. This strategy has been shown efficient for image synthesis and enhancement (Zhang, Ng, and Chen 2018).

Discriminator. We implement D_Y and D_Z as the discriminators. They share the same multi-branches structures as in (Liu, Breuel, and Kautz 2017), but do not share parameters. For the i -th branch in D_Y or D_Z , the input image is down-sampled by $i - 1$ times via average pooling and then fed to the network. Each branch has four convolution layers with $\{32, 64, 128, 32\}$ filters of size 4×4 and a stride of 2. In the end, the output features from different branches are fused together followed by a sigmoid activation. For the synthetic data case, the number of branches is simply 1, while for the real image separation tasks, the number of branches is set to 3.

Training. We use ADAM (Kingma and Ba 2015) optimizer with a learning rate of 0.0001 and momentums of 0.0 and 0.9. The default values of hyper-parameters in Equ. (5)-(8) are set to $\lambda_0 = 5.0$, $\lambda_1 = 0.5$, $\lambda_2 = 0.5$, $\lambda_3 = 1.0$, $\lambda_4 = 1.0$, $\lambda_5 = 1.0$, $\lambda_6 = 1.0$ and $\alpha = 1.4$.

Each mini-batch in training contains one input image from the domain \mathcal{X} , one layer image from the domain \mathcal{Y} and another one from domain \mathcal{Z} . Note that these three images do not form an image separation triplet. For the real world tasks, data augmentation operations are applied, including random scalings by factors within $0.8 \sim 1.2$, random croppings and random horizontal flippings.

Inheriting from GANs, the training solves a min-max problem to find a saddle point in the solution space. To make it stable, we use a gradient update strategy similar to (Liu, Breuel, and Kautz 2017) with gradient penalization to solve Equ. (4). Specifically, a gradient descent step first updates E_Y , E_Z , G_Y and G_Z with fixed D_Y and D_Z , then another gradient descent step updates D_Y and D_Z with fixed E_Y , E_Z , G_Y and G_Z .

5 Experiments

Experiments are conducted on the synthetic data and also real world single image separation tasks. The former aims at providing with intuitive and quantitative results, and the latter demonstrates the real world applicability of our method. Moreover, our method is extended to solve a more challenging task that involves multiple layers.

5.1 Datasets

We prepare three types of dataset: 1) synthetic dataset, 2) real world intrinsic image decomposition dataset and 3) real world reflection removal dataset. Note that in any case, we do not have the data triplet from the three domains, *i.e.*, input image with its two layers, in the training set.

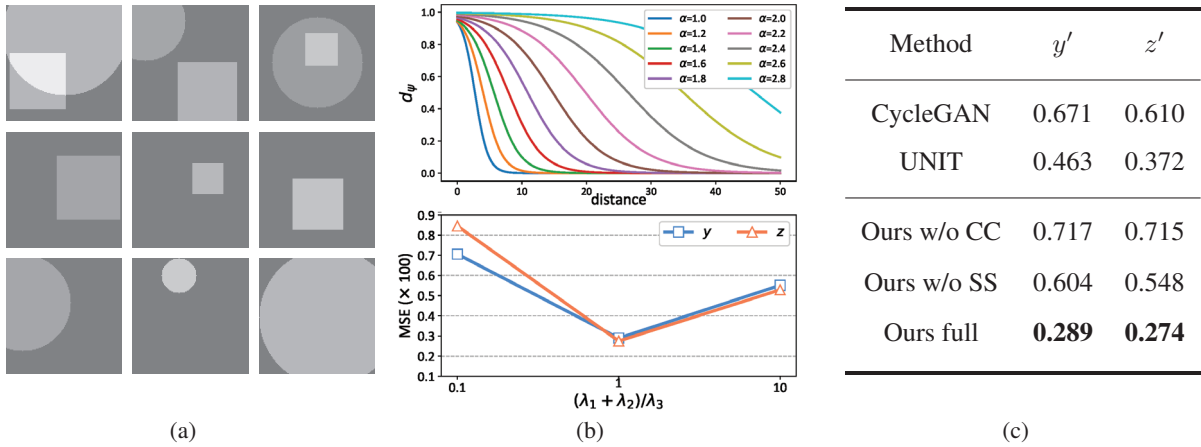


Figure 3: (a) Demonstration of the synthetic non-triplet data. Top row: input image x (a blend of circle and square), middle row: y (only square), and bottom row: z (only circle). Our purpose is to decompose x into $y \oplus z$ by learning from these data. (b) Visualization following the analysis in Sec. 5.2. Top: distance loss curve v.s. different α . Bottom: separation performance v.s. different hyper parameters in SS. (c) Illustration of how self-supervision (SS) and cycle consistency (CC) affect the performance, and a comparison with existing unsupervised methods. Note that the numbers indicate image MSE $\times 100$.

Examples of the synthetic data is shown in Fig. 3 (a). The image size is set to 128×128 . The squares and circles in the images are rendered with arbitrary sizes, positions and brightness levels. The training set includes 4000 images and the test set has 1000 images.

For intrinsic image decomposition, we use 220 images in the MIT intrinsic dataset (Grosse et al. 2011). To avoid the ground truth triplet in training, we divide the dataset into two subsets without overlapped scenes. Then we select the input images only from the first subset, and the layer images only from the second. The test samples are from the second subset. For reflection removal, we use the reflection removal benchmark dataset (Wan et al. 2017) with 454 images, and similarly select our data as above.

5.2 Evaluation on Synthetic Data

We conduct experiments with synthetic data to analyse our method in detail. Since the self-supervision (SS) is a major technical component of our method, we study its performance and optimal parameter settings.

First, we analyze the distance loss in Equ. (6). Results are shown in Fig. 3 (b), where $\text{loss} = 1 - d_\psi$ and distance indicates $|a - b|$. The curves show that the loss values, ranging from 0 to 1, indeed decline with increasing distance values but with different shapes. According to the curve shapes, we set $\alpha = 1.4$ empirically in the follow experiments.

We then study the balance between the intra-domain distance and inter-domain distance in SS. In Equ. (5), λ_1, λ_2 control the intra-domain distance loss while λ_3 controls the inter-domain distance loss. Therefore, we tune these hyper parameters and re-train the model without touching other settings. The resulting mean square errors (MSEs) are shown in the bottom of Fig. 3 (b), where $(\lambda_1 + \lambda_2)/\lambda_3 = 0.1$ produces the worst model while $(\lambda_1 + \lambda_2)/\lambda_3 = 1$ gives the best one. This proves that a good balance should consider both intra-domain and inter-domain distances of the feature, which con-

firms our Assumption B. In practice, we set $\lambda_1 = \lambda_2 = 0.5$ and $\lambda_3 = 1$ in the following experiment.

Next, we study the impacts of our SS and CC constraints by conducting an ablation study. Results are reported in Fig. 3 (c), showing that removing any of these constraints will increase the final error drastically, *i.e.*, twice as large or even more. This persuasively proves that both the SS and CC are essential to our method.

For more information, we also show the results of CycleGAN (Zhu et al. 2017) and UNIT (Liu, Breuel, and Kautz 2017) on our synthetic dataset. Their MSEs are 0.67/0.61 and 0.46/0.37, which are obviously larger than ours (0.289/0.274). This demonstrates the advantage of our method.

Finally, we consider the overall loss function in Equ. (4). We introduce two temporary weights w_1 and w_2 to balance each term, and the resulting loss function becomes

$$\min_{E, G} \max_D w_1 \mathcal{L}_{SS} + \mathcal{L}_{GAN} + w_2 \mathcal{L}_{CC}. \quad (9)$$

We conduct experiments with different w_1 and w_2 and the results are in Fig. 4, where we find that $w_1 = 1, w_2 = 10$

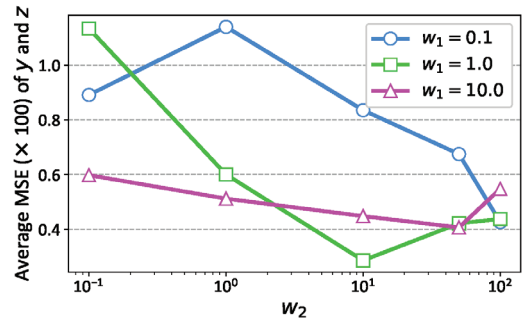


Figure 4: Results with different w_1 and w_2 in Equ. (9).

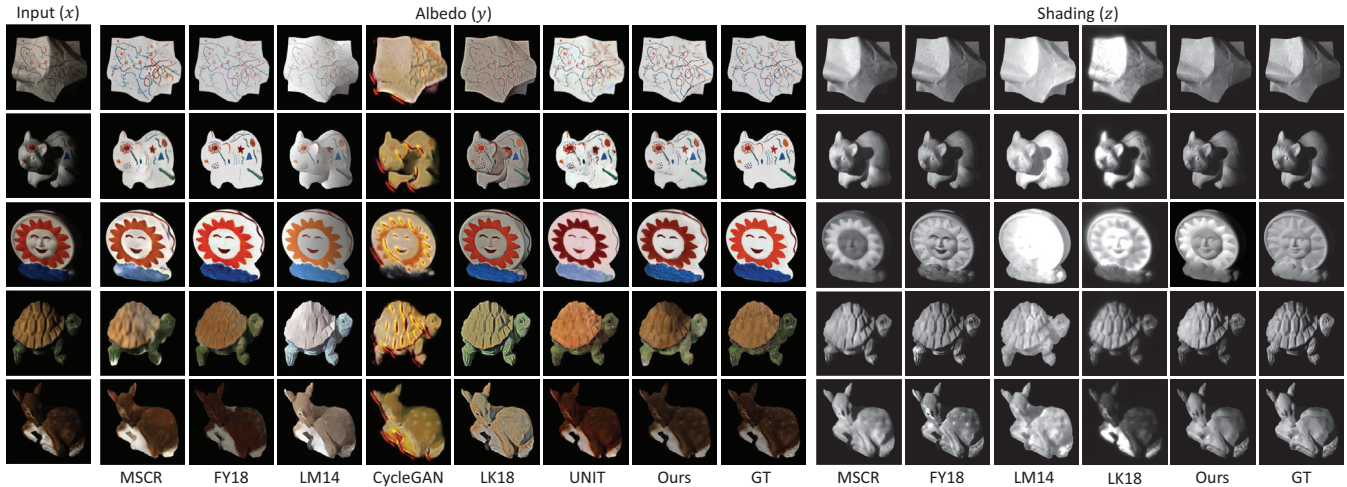


Figure 5: Qualitative comparisons on MIT dataset. From left to right: MSCR and FY18 are fully supervised method shown for reference. LM14, CycleGAN and UINT are unsupervised methods than can be directly compared with outs. Our results have a good visual quality and appear almost identical to the ground truth.

is an optimal choice. According to Equ. (8), this suggests $\lambda_4 = \lambda_5 = \lambda_6 = 10.0$ in the experiment.

5.3 Evaluation on Real Image Separation Tasks

Intrinsic image decomposition. We compare the proposed method with the most representative unsupervised methods, *i.e.*, LM14 (Li and Brown 2014), CycleGAN (Zhu et al. 2017), UNIT (Liu, Breuel, and Kautz 2017) and LK18 (Lettry, Vanhoey, and Van Gool 2018). They are trained on the same dataset as ours. Besides, we also show results of fully-supervised methods MSCR (Narihira, Maire, and Yu 2015) and FY18 (Fan et al. 2018) using their pre-trained models. Note that these fully-supervised methods are trained on ground truth data triplets and thus their results can only serve as a reference.

Visual comparison among all the methods is illustrated in Fig. 5. CycleGAN produces the worst images with bad textures and wrong colors. Its training may fail with a high probability. LK18 and LM14 perform better than CycleGAN while their visual qualities are still unsatisfactory. UNIT produces relatively good albedo layers. However, it cannot output the shading layer due to its image-to-image translation framework. On the other hand, our method generates both the albedo and shading layers with a good quality and they appear highly identical to the ground truth images. The visual quality of our results is even better than the two fully supervised methods, *i.e.*, MSCR and FY18, in some cases.

We also provide a qualitative comparison in Table 1, where LMSE is an error metric specifically designed for intrinsic image data introduced by (Grosse et al. 2009). Numerical results show that our method achieves state-of-the-art performance among these unsupervised methods. Its accuracy is even comparable to the fully supervised methods that rely on exact ground truth data triplets in training.

In this experiment, we also study our feature extraction results. Fig. 7 visualizes the distribution of the encoded features

		MSE			LMSE
Method		Albedo	Shading	Avg.	Total
Unsup.	LM14	0.0286	0.0227	0.0255	0.0366
	CycleGAN	0.0272	N/A	0.0272	0.407
	UNIT	0.0207	N/A	0.0207	0.0310
	LK18	0.0232	0.0166	0.0197	0.0379
	Ours	0.0167	0.0140	0.0154	0.025
Sup. (Ref.)	MSCR	0.0207	0.0124	0.0165	0.024
	FY18	0.0127	0.0085	0.0106	0.020

Table 1: Numerical comparison on MIT intrinsic dataset. Note that MSCR and FY18 are fully-supervised methods, which are trained on ground truth data triplets, and thus their results can only serve as a reference.

by using their 2D PCA vectors. It is clear that the encoder E_Y projects x and y into the same cluster (indicating the same latent space for albedo layer), and the encoder E_Z projects x and z into another cluster (latent space for shading layer). The two clusters stay apart from each other with a large distance. This verifies the correctness of our Assumption A and B and thus confirms the design of our method.

Reflection removal. Similar to the intrinsic image decomposition case, we compare our method with the unsupervised methods LM14, CycleGAN and UNIT. They are trained on the same dataset. For the fully-supervised methods, we use the pre-trained models of ZN18 (Zhang, Ng, and Chen 2018) and BDN (Yang et al. 2018) on the same dataset but with ground truth triplet data. Fig. 6 gives the visual comparison among unsupervised methods and ours. LM14 can separate both background and reflection layers from the input image, but the separation is clearly not ideal. Both the CycleGAN and UNIT only recover the background layer, where there still exist artifacts or remaining reflections in many image




Figure 6: Visual results by different unsupervised methods for reflection removal. From left to right, LM14, CycleGAN and UNIT. Note that CycleGAN and UINIT only recover background image without producing the reflection image.

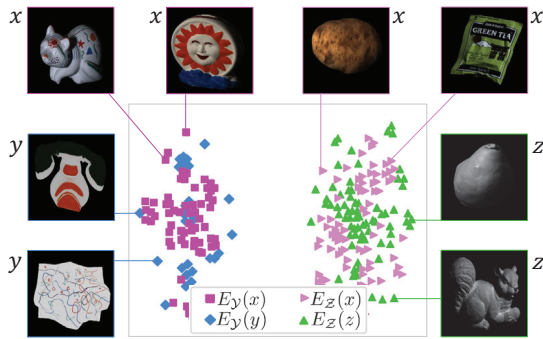


Figure 7: Visualization of our feature encoding distribution using the corresponding PCA vectors. We use t-SNE (Maaten and Hinton 2008) to aid the visualization of the latent space. Encoder E_y projects x and y into the same cluster (latent space for albedo layer), and encoder E_z projects x and z into another cluster (latent space for shading layer).

regions. Overall, our method achieves the best background cleanness and the visual quality of the reflection layer.

To make a qualitative comparison with state-of-the-art methods, we use peak-signal-to-noise-ratio (PSNR) (Huynh-Thu and Ghanbari 2008) and structural similarity index (SSIM) (Wang et al. 2004) as evaluation metrics. Table 2 shows that our method achieves the state-of-the-art performance against other unsupervised methods, and is also superior to the fully-supervised method BDN (Yang et al. 2018).

5.4 Evaluation on Multi-layer Separation

In this section, we extend our evaluation to the multi-layer case. We re-render the synthetic dataset with three different shapes, *i.e.*, square, circle and triangle, as shown in Fig. 8 (a). Totally 5000 images are synthesized without containing any triplet data, and other settings remain the same as described in Sec. 5.2. Accordingly, we modify our method by adding another network branch of a generator and a discriminator to handle the third layer. The example results are shown in Fig. 8 (b), which demonstrate the successfulness of our method to handle such a challenging problem.

		Background		Reflection	
Method		SSIM	PSNR	SSIM	PSNR
Unsup.	LM14	0.715	17.00	0.357	17.58
	CycleGAN	0.622	16.63	N/A	N/A
	UNIT	0.738	19.23	N/A	N/A
	Ours	0.821	20.47	0.425	19.31
Sup. (Ref.)	BDN	0.820	18.87	N/A	N/A
	ZN18	0.858	22.22	0.420	18.94

Table 2: Comparison on reflection removal dataset. Note that BDN and ZN18 are fully-supervised methods, which are trained on ground truth data triplets, and thus their results can only serve as a reference.




Figure 8: The extend experiment on three layers separation task. (a) Example of the blend image and each layer. (b) Two visual results obtained by the proposed method. From left to right: input image and separated layers.

6 Conclusion

In this paper, we propose an unsupervised single image layer separation method. It is unsupervised since it requires no ground truth data triplet, *i.e.*, input images and its layer images for training. To this end, we introduce two assumptions about the feature distributions of different layers, namely the shared information consistency and layer independence assumptions. Based on them, we design our method by using the GANs framework with self-supervision and cycle consistency constraints. Experimental results show that our method outperforms existing unsupervised methods in both synthetic and real world tasks. The method can also be extended to

solve a more challenging three-layer separation task.

References

- Bi, S.; Han, X.; and Yu, Y. 2015. An l1 image transform for edge-preserving smoothing and scene-level intrinsic decomposition. *TOG*.
- Cao, C.; Lu, F.; Li, C.; Lin, S.; and Shen, X. 2019. Makeup removal via bidirectional tunable de-makeup network. *IEEE TMM*.
- Dmitry Ulyanov, A. V., and Lempitsky, V. 2016. Instance normalization: The missing ingredient for fast stylization. In *arXiv preprint arXiv:1607.08022*.
- Fan, Q.; Yang, J.; Hua, G.; Chen, B.; and Wipf, D. 2017. A generic deep architecture for single image reflection removal and image smoothing. In *ICCV*.
- Fan, Q.; Yang, J.; Hua, G.; Chen, B.; and Wipf, D. 2018. Revisiting deep intrinsic image decompositions. In *CVPR*.
- Girshick, R. 2015. Fast r-cnn. In *ICCV*.
- Grosse, R.; Johnson, M. K.; Adelson, E. H.; and Freeman, W. T. 2009. Ground truth dataset and baseline evaluations for intrinsic image algorithms. In *ICCV*.
- Grosse, R.; Johnson, M. K.; Adelson, E. H.; and Freeman, W. T. 2011. Ground truth dataset and baseline evaluations for intrinsic image algorithms. In *ICCV*.
- Han, B.-J., and Sim, J.-Y. 2017. Reflection removal using low-rank matrix completion. In *CVPR*.
- Hao, Z.; You, S.; Li, Y.; Li, K.; and Lu, F. 2019. Learning from synthetic photorealistic raindrop for single image raindrop removal. In *ICCVW*.
- Hoshen, Y. Towards unsupervised single-channel blind source separation using adversarial pair unmix-and-remix. In *ICASSP*. IEEE.
- Huynh-Thu, Q., and Ghanbari, M. 2008. Scope of validity of psnr in image/video quality assessment. *Electronics letters*.
- Ioffe, S., and Szegedy, C. 2015. Batch normalization: Accelerating deep network training by reducing internal covariate shift. In *ICML*.
- Kingma, D. P., and Ba, J. 2015. Adam: A method for stochastic optimization. In *ICLR*.
- Lettry, L.; Vanhoey, K.; and Van Gool, L. 2018. Unsupervised deep single-image intrinsic decomposition using illumination-varying image sequences. In *WACV*.
- Li, Y., and Brown, M. S. 2013. Exploiting reflection change for automatic reflection removal. In *ICCV*.
- Li, Y., and Brown, M. S. 2014. Single image layer separation using relative smoothness. In *CVPR*.
- Li, Z., and Snavely, N. 2018. Learning intrinsic image decomposition from watching the world. In *CVPR*.
- Liu, M.-Y.; Breuel, T.; and Kautz, J. 2017. Unsupervised image-to-image translation networks. In *NIPS*.
- Lu, F.; Matsushita, Y.; Sato, I.; Okabe, T.; and Sato, Y. 2015. From intensity profile to surface normal: photometric stereo for unknown light sources and isotropic reflectances. *IEEE T-PAMI*.
- Lu, F.; Chen, X.; Sato, I.; and Sato, Y. 2017. Symps: Brdf symmetry guided photometric stereo for shape and light source estimation. *IEEE T-PAMI*.
- Maaten, L. v. d., and Hinton, G. 2008. Visualizing data using t-sne. *Journal of machine learning research*.
- Narihira, T.; Maire, M.; and Yu, S. X. 2015. Direct intrinsics: Learning albedo-shading decomposition by convolutional regression. In *ICCV*.
- Radford, A.; Metz, L.; and Chintala, S. 2016. Unsupervised representation learning with deep convolutional generative adversarial networks. In *ICLR*.
- Ronneberger, O.; Fischer, P.; and Brox, T. 2015. U-net: Convolutional networks for biomedical image segmentation. In *MICCAI*.
- Shih, Y.; Krishnan, D.; Durand, F.; and Freeman, W. T. 2015. Reflection removal using ghosting cues. In *CVPR*.
- Simonyan, K., and Zisserman, A. 2015. Very deep convolutional networks for large-scale image recognition.
- Tatarchenko, M.; Richter, S. R.; Ranftl, R.; Li, Z.; Koltun, V.; and Brox, T. 2019. What do single-view 3d reconstruction networks learn? In *CVPR*.
2017. Janner, michael and wu, jiajun and kulkarni, tejas d and yildirim, ilker and tenenbaum, josh. In *Advances in Neural Information Processing Systems*.
- Wan, R.; Shi, B.; Duan, L. Y.; Tan, A. H.; and Kot, A. C. 2017. Benchmarking single-image reflection removal algorithms. In *ICCV*.
- Wan, R.; Shi, B.; Li, H.; Duan, L.-Y.; and Kot, A. C. 2019. Face image reflection removal. In *CVPR*.
- Wang, Z., and Lu, F. 2019. Single image intrinsic decomposition with discriminative feature encoding. In *ICCVW*.
- Wang, Z.; Bovik, A. C.; Sheikh, H. R.; Simoncelli, E. P.; et al. 2004. Image quality assessment: from error visibility to structural similarity. *TIP*.
- Yang, J.; Gong, D.; Liu, L.; and Shi, Q. 2018. Seeing deeply and bidirectionally: A deep learning approach for single image reflection removal. In *ECCV*.
- Yi, Z.; Zhang, H.; Tan, P.; and Gong, M. 2017. Dualgan: Unsupervised dual learning for image-to-image translation. In *ICCV*.
- Zhang, X.; Ng, R.; and Chen, Q. 2018. Single image reflection separation with perceptual losses. In *CVPR*.
- Zhu, J.-Y.; Park, T.; Isola, P.; and Efros, A. A. 2017. Unpaired image-to-image translation using cycle-consistent adversarial networks. In *CVPR*.