

## Stereo Computation for a Single Mixture Image

Yiran Zhong<sup>1,3,4</sup>, Yuchao Dai<sup>2</sup>, and Hongdong Li<sup>1,4</sup>

<sup>1</sup> Australian National University

<sup>2</sup> Northwestern Polytechnical University

<sup>3</sup> Data61 CSIRO

<sup>4</sup> Australian Centre for Robotic Vision

{yiran.zhong,hongdong.li}@anu.edu.au daiyuchao@nwpu.edu.cn

**Abstract.** This paper proposes an original problem of *stereo computation from a single mixture image*— a challenging problem that had not been researched before. The goal is to separate (*i.e.*, unmix) a single mixture image into two constitute image layers, such that the two layers form a left-right stereo image pair, from which a valid disparity map can be recovered. This is a severely illposed problem, from one input image one effectively aims to recover three (*i.e.*, left image, right image and a disparity map). In this work we give a novel deep-learning based solution, by jointly solving the two subtasks of image layer separation as well as stereo matching. Training our deep net is a simple task, as it does not need to have disparity maps. Extensive experiments demonstrate the efficacy of our method.

**Keywords:** stereo computation, image separation, anaglyph, monocular depth estimation, double vision.

### 1 Introduction

Stereo computation (stereo matching) is a well-known and fundamental vision problem, in which a dense depth map  $D$  is estimated from two images of the scene from slightly different viewpoints. Typically, one of the cameras is in the left (denoted by  $I_L$ ) and the other in the right (denoted by  $I_R$ ), just like we have left and right two eyes. Given a single image, it is generally impossible to infer a disparity map, unless using strong semantic-dependent image priors such as those single-image depth-map regression works powered by deep-learning [1–3]. Even though these learning based monocular depth estimation methods could predict a reasonable disparity map from a single image, they all assume the input image to be *an original color image*.

In this paper, we propose a novel and original problem, assuming instead one is provided with one *single mixture image* (denoted by  $I$ ) which is a composition of an original stereo image pair  $I_L$  and  $I_R$ , *i.e.*  $I = f(I_L, I_R)$ , and the task is to simultaneously recover both the stereo image pair  $I_L$  and  $I_R$ , and an accurate dense depth-map  $D$ . Under our problem definition,  $f$  denotes different image composition operators that generate the mixture image, which is to be defined in details later. This is a very challenging problem, due to the obvious ill-pose

(under-constrained) nature of the task, namely, from one input mixture image  $I$  one effectively wants to recover three images ( $I_L$ ,  $I_R$ , and  $D$ ).

In theory it appears to be a blind signal separation (BSS) task, *i.e.*, separating an image into two different component images. However, conventional methods such as BSS using independent component analysis (ICA) [4] are unsuitable for this problem as they make strong assumptions on the statistical independence between the two components. Under our problem definition,  $I_L, I_R$  are highly correlated. In computer vision, image layer separation such as reflection and highlight removal [5, 6] are also based on the difference in image statistics, again, unsuitable. Another related topic is image matting [7], which refers to the process of accurate foreground estimation from an image. However it either needs human interaction or depends on the difference between foreground object and background, which cannot be applied to our task.

In this paper, we advocate a novel deep-learning based solution to the above task, by using a simple network architecture. We could successfully solve for a stereo pair  $L, R$  and a dense depth map  $D$  from a single mixture image  $I$ . Our network consists of an image separation module and a stereo matching module, where the two modules are optimized jointly. Under our framework, the solution of one module benefits the solution of the other module. It is worth-noting that the training of our network does not require ground truth depth maps.

At a first glance, this problem while intriguing, has pure intellectual interest only, not perhaps no practical use. In contrast, we show this is not the case: in this paper, we show how to use it to solve for three very different vision problems: double vision, de-anaglyphy and even monocular depth estimation.

The requirement for de-anaglyph is still significant. If search on Youtube, there are hundreds if not thousands of thousands of anaglyph videos, where the original stereo images are not necessarily available. Our methods and the previous work [8][9] enable the recovery of the stereo images and the corresponding disparity map, which will significantly improve the users' *real 3D experience*. As evidenced in the experiments, our proposed method clearly outperforms the existing work with a wide gap. Last but not least, our model could also handle the task of monocular depth estimation and it comes as a surprise to us: Even with one single mixture image, trained on the KITTI benchmark, our method produces the state of the art depth estimation, with results more better than those traditional two images based methods.

## 2 Setup the stage

In this paper we study two special cases of our novel problem of *joint image separation and stereo computation*, namely *anaglyph* (red-cyan stereo) and *diplopia* (double vision) (see Figure. 1), which have not been well-studied in the past.

- 1) **Double vision (aka. diplopia)**: *Double vision* is the simultaneous perception of two images (a stereo pair) of a single object in the form of a single mixture image. Specifically, under the double vision (diplopia) model (*c.f.* Fig. 1 (left column)), the perceived image  $I = f(I_L, I_R) = (I_L + I_R)/2$ , *i.e.*, the



**Fig. 1.** Examples of image separation for single image based stereo computation. **Left column:** A double-vision image is displayed here. **Right column:** A red-cyan stereo image contains channels from the left and right images.

image composition is  $f$  a direct average of the left and the right images. Note that the above equation shares similarity with the linear additive model in layer separation [5, 10, 11] for reflection removal and raindrop removal, we will discuss the differences in details later.

- 2) **Red-Cyan stereo (aka. anaglyph):** An anaglyph (*c.f.* Fig. 1 (right column)) is a single image created by selecting chromatically opposite colors (typically red and cyan) from a stereo pair. Thus given a stereo pair  $I_L, I_R$ , the image composition operator  $f$  is defined as  $I = f(I_L, I_R)$ , where the red channel of  $I$  is extracted from the red channel of  $I_L$  while its green and blue channels are extracted from  $I_R$ . De-anaglyph [8, 9] aims at estimating both the stereo pair  $I_L, I_R$  (color restoration) and computing its disparity maps.

At a first glance, the problem seems impossible as one has to generate two images plus a dense disparity map from *one single input*. However, since the two constitute images are not arbitrary but related by a valid disparity map. Therefore, they must be able to aligned well along the scanlines horizontally. For anaglyph stereo, existing methods [8, 9] exploit both image separation constraint and disparity map computation to achieve color restoration and stereo computation. Joulin and Kang [9] reconstructed the original stereo pairs given the input anaglyph by using a modified SIFT-flow method [12]. Williem *et.al.* [8] presented a method to solve the problem within iterations of color restoration and stereo computation. These works suggest that by properly exploiting the image separation and stereo constraints, it is possible to restore the stereo pair images and compute the disparity map from a single mixture image.

There is little work in computer vision dealing with double vision (diplopia), which is nonetheless an important topic in ophthalmology and visual cognition. The most related works seem to be layer separation [5, 10], where the task is to decompose an input image into two layers corresponding to the background image and the foreground image. However, there are significant differences between our problem and general layer separation. For layer separation, the two layers of the composited image are generally independent and statistically different. In contrast, the two component images are highly correlated for double vision.


Even though there have been remarkable progresses in monocular depth estimation, current state-of-the-art network architectures [1, 2] and [13] cannot be directly applied to our problem. This is because that they depend on a single left/right image input, which is unable to handle image mixture case investigated

in this work. Under our problem definition, the two tasks of image separation and stereo computation are tightly coupled: stereo computation is not possible without correct image separation; on the other hand, image separation will benefit from disparity computation.

In this paper, we present a unified framework to handle the problem of stereo computation for a single mixture image, which naturally unifies various geometric vision problems such as anaglyph, diplopia and even monocular depth estimation. Our network can be trained with the supervision of stereo pair images only without the need for ground truth disparity maps, which significantly reduces the requirements for training data. Extensive experiments demonstrate that our method achieves superior performances.

### 3 Our Method

In this paper, we propose an end-to-end deep neural network to simultaneously learn image separation and stereo computation from a single mixture image. It can handle a variety forms of problems such as anaglyph, de-diplopia and even monocular depth estimation. Note that existing work designed for either layer-separation or stereo-computation cannot be applied to our problem directly. This is because these two problems are deeply coupled, *i.e.*, the solution of one problem affects the solution of the other problem. By contrast, our formulation to be presented as below, jointly solves both problems.



**Fig. 2. Overview of our proposed stereo computation for a single mixture image framework.** Our network consists of an image separation module and a stereo computation module. Take a single mixture image as input, our network simultaneously separates the image into a stereo image pair and computes a dense disparity map.

#### 3.1 Mathematical Formulation

Under our mixture model, quality of depth map estimation and image separation are evaluated jointly and therefore, the solution of each task can benefit from each other. Our network model (*c.f.*, Fig. 2) consists of two modules, *i.e.* an

image separation module and a stereo computation module. During network training, only the ground-truth stereo pairs are needed to provide supervisions for both image separation and stereo computation.

By considering both the image separation constraint and the stereo computation constraint in network learning, we define the overall loss function as:

$$\mathcal{L}(\theta_L, \theta_R, \theta_D) = \mathcal{L}_C(\theta_L, \theta_R) + \mathcal{L}_D(\theta_D), \quad (1)$$

where  $\theta_L, \theta_R, \theta_D$  denote the network parameters corresponding to the image separation module (left image prediction and right image prediction) and the stereo computation module. A joint optimization of  $(\theta_L, \theta_R, \theta_D) = \arg \min \mathcal{L}(\theta_L, \theta_R, \theta_D)$  gives both the desired stereo image pair and the disparity map.

### 3.2 Image Separation

The input single mixture image  $I \in \mathbb{R}^{H \times W \times 3}$  encodes the stereo pair image as  $I = f(I_L, I_R)$ , where  $f$  is the image composition operator known a priori. To learn the stereo image pair from the input single mixture image, we present a unified end-to-end network pipeline. Specifically, denote  $\mathcal{F}$  as the learned mapping from the mixture image to the predicted left or right image parameterized by  $\theta_L$  or  $\theta_R$ . The objective function of our image separation module is defined as,

$$\alpha_c \mathcal{L}_c(\mathcal{F}(I; \theta_L), I_L) + \alpha_p \mathcal{L}_p(\mathcal{F}(I; \theta_L)), \quad (2)$$

where  $I$  is the input single mixture image,  $I_L, I_R$  are the ground truth stereo image pair. The loss function  $\mathcal{L}$  measures the discrepancy between the predicted stereo images and the ground truth stereo images. The object function for the right image is defined similarly.

In evaluating the discrepancy between images, various loss functions such as  $\ell_2$  loss [14], classification loss [15] and adversarial loss [16] can be applied. Here, we leverage the pixel-wise  $\ell_1$  regression loss as the content loss of our image separation network,

$$\mathcal{L}_c(\mathcal{F}(I; \theta_L), I_L) = |\mathcal{F}(I; \theta_L) - I_L|. \quad (3)$$

This loss allows us to perform end-to-end learning as compatible with the stereo matching loss and do not need to consider class imbalance problem or add an extra network structure as a discriminator.

Researches on natural image statistics show that a typical real image obeys sparse spatial gradient distributions [17]. According to Yang *et.al.* [5], such a prior can be represented as the Total Variation (TV) term in energy minimization. Therefore, we have our image prior loss:

$$\mathcal{L}_p(\mathcal{F}(I; \theta_L)) = |\mathcal{F}(I; \theta_L)|_{\text{TV}} = |\nabla \mathcal{F}(I; \theta_L)|, \quad (4)$$

where  $\nabla$  is the gradient operator.

We design a U-Net architecture [18] for image separation, which has been used in various conditional generation tasks. Our image separation module consists of

22 convolutional layers. Each convolutional layer contains one convolution-relu pair except for the last layer and we use element-wise add for each skip connection to accelerate the convergence. For the output layer, we utilize a “tanh” activation function to map the intensity value between  $-1$  and  $1$ . A detailed description of our network structure is provided in the supplemental material.

The output of our image separation module is a 6 channels image, where the first 3 channels represent the estimated left image  $\mathcal{F}(I; \theta_L)$  and the rest 3 channels for the estimated right image  $\mathcal{F}(I; \theta_R)$ . When the network converges, we could directly use these images as the image separation results. However, for the de-anaglyph task, as there is extra constraint (the mixture happens at channel level), we could leverage the color prior of an anaglyph that the desired image separation (colorization) can be further improved by warping corresponding channels based on the estimated disparity maps.

For the monocular depth estimation task, only the right image will be needed as the left image has been provided as input.

### 3.3 Stereo Computation

The input to the stereo computation module is the separated stereo image pair from the image separation module. The supervision of this module is the ground truth stereo pairs rather than the inputs. The benefit of using ground truth stereo pairs for supervision is that it makes the network not only learn how to find the matching points, but also makes the network to extract features that are robust to the noise from the generated stereo images.

Fig. 2 shows an overview of our stereo computation architecture, we adopt a similar stereo matching architecture from Zhong *et al.* [19] without its consistency check module. The benefit for choosing such a structure is that their model can converge within 2000 iterations which makes it possible to train the entire network in an end-to-end fashion. Additionally, removing the need of ground truth disparity maps enables us to access much more accessible stereo images.

Our loss function for stereo computation is defined as:

$$\mathcal{L}_D = \omega_w(\mathcal{L}_w^l + \mathcal{L}_w^r) + \omega_s(\mathcal{L}_s^l + \mathcal{L}_s^r), \quad (5)$$

where  $\mathcal{L}_w^l, \mathcal{L}_w^r$  denote the image warping appearance loss,  $\mathcal{L}_s^l, \mathcal{L}_s^r$  express the smoothness constraint on the disparity map.

Similar to  $\mathcal{L}_c$ , we form a loss in evaluating the image similarity by computing the pixel-wise  $\ell_1$  distance between images. We also add a structural similarity term SSIM [20] to improve the robustness against illumination changes across images. The appearance loss  $\mathcal{L}_w^l$  is derived as:

$$\mathcal{L}_w^l(I_L, I_L'') = \frac{1}{N} \sum \lambda_1 \frac{1 - \mathcal{S}(I_L, I_L'')}{2} + \lambda_2 |I_L - I_L''|, \quad (6)$$

where  $N$  is the total number of pixels and  $I_L''$  is the reconstructed left image.  $\lambda_1, \lambda_2$  balance between structural similarity and image appearance difference.

According to [2],  $I_L''$  can be fully differentially reconstructed from the right image  $I_R$  and the right disparity map  $d_R$  through bilinear sampling [21].

For the smoothness term, similar to [2], we leverage the Total Variation (TV) and weight it with image’s gradients. Our smoothness loss for disparity field is:

$$\mathcal{L}_s^l = \frac{1}{N} \sum |\nabla_u d_L| e^{-|\nabla_u I_L|} + |\nabla_v d_L| e^{-|\nabla_v I_L|}. \quad (7)$$

### 3.4 Implementation Details

We implement our network in TensorFlow [22] with 17.1M trainable parameters. Our network can be trained from scratch in an end-to-end fashion with a supervision of stereo pairs and optimized using RMSProp [23] with an initial learning rate of  $1 \times 10^{-4}$ . Input images are normalized with pixel intensities level ranging from -1 to 1. For the KITTI dataset, the input images are randomly cropped to  $256 \times 512$ , while for the Middlebury dataset, we use  $384 \times 384$ . We set disparity level to 96 for the stereo computation module. For weighting loss components, we use  $\alpha_c = 1, \alpha_p = 0.2, \omega_w = 1, \omega_s = 0.05$ . We set  $\lambda_1 = 0.85, \lambda_2 = 0.15$  throughout our experiments. Due to the hardware limitation (Nvidia Titan Xp), we only use batch size 1 during network training.

## 4 Experiments and Results

In this section, we validate our proposed method and present experimental evaluation for both de-anaglyph and de-diplopia (double vision). For experiments on anaglyph images, given a pair of stereo images, the corresponding anaglyph image can be generated by combining the red channel of the left image and the green/blue channels of the right image. Any stereo pairs can be used to quantitatively evaluate the performance of de-anaglyph. However, since we also need to quantitatively evaluate the performance of anaglyph stereo matching, we use two stereo matching benchmarking datasets for evaluation: Middlebury dataset [24] and KITTI stereo 2015 [25]. Our network is initially trained on the KITTI Raw dataset with 29000 stereo pairs that listed by [2] and further fine-tuned on Middlebury dataset. To highlight the generalization ability of our network, we also perform qualitative experiments on random images from Internet. For de-diplopia (double vision), we synthesize our inputs by averaging stereo pairs. Qualitative and quantitative results are reported on KITTI stereo 2015 benchmark [25] as well. Similar to the de-anaglyph experiment, we train our initial model on the KITTI raw dataset.

### 4.1 Advantages of joint optimization

Our framework consists of image separation and stereo computation, where the solution of one subtask benefits the solution of the other subtask. Direct stereo computation is impossible for a single mixture image. To analyze the advantage of joint optimization, we perform ablation study in image separation without

stereo computation and the results are reported in Table 1. Through joint optimization, the average PSNR increases from 19.5009 to 20.0914, which demonstrates the benefit of introducing the stereo matching loss in image separation.


Metric	Image separation only	Joint optimization
PSNR	19.5009	20.0914

**Table 1.** Ablation study of image separation on KITTI.

## 4.2 Evaluation of Anaglyph Stereo

We compare the performance of our method with two state-of-the-art de-anaglyph methods: Joulin *et.al.* [9] and Williem *et.al.* [8]. Evaluations are performed on two subtasks: stereo computation and image separation (color restoration).


**Stereo Computation.** We present qualitative comparison of estimated disparity maps in Fig. 3 for Middlebury [24] and in Fig. 4 for KITTI 2015 [25]. Stereo pairs in Middlebury are indoor scenes with multiple handcrafted layouts and the ground truth disparities are captured by highly accurate structural light sensors. On the other hand, the KITTI stereo 2015 consists of 200 outdoor frames in their training set, which is more challenging than the Middlebury dataset. The ground truth disparity maps are generated by sparse LIDAR points and CAD models.



**Fig. 3.** Qualitative stereo computation results on the Middlebury dataset by our method. From left to right: input anaglyph image, ground truth disparity map, disparity map generated by Williem *et.al.* [8] and our method.

On both datasets, our method can generate more accurate disparity maps than previous ones from visual inspection. It can be further evidenced by the quantitative results of bad pixel percentage that shown in Table. 2 and Fig. 5. For the Middlebury dataset, our method achieves 32.55% performance leap than Williem *et.al.* [8] and 352.28% performance leap than Joulin *et.al.* [9]. This is reasonable as Joulin *et.al.* [9] did not add disparity into its optimization. For the KITTI dataset, we achieve an average bad pixel ratio (denoted as D1\_all) of 5.96% with 3 pixel thresholding across 200 images in the training set as opposed to 13.66% by Joulin *et.al.* [9] and 14.40% by Williem *et.al.* [8].

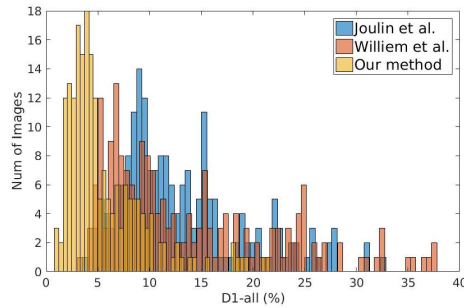




**Fig. 4.** Qualitative disparity map recovery results on KITTI-2015 of our method. Top row: Input anaglyph image and ground truth disparity map. Bottom row: Result of Williem *et.al.* [8] and our result.

Method	Tsukuba	Venus	Cones	Teddy	Mean
Joulin [9]	14.02	25.90	23.49	43.85	26.82
Williem [8]	12.53	2.24	8.00	8.68	7.86
Ours	<b>10.44</b>	<b>1.25</b>	<b>4.51</b>	<b>7.51</b>	<b>5.93</b>

**Table 2.** Performance comparison in disparity map estimation for de-anaglyph on the Middlebury dataset. We report the bad pixel ratio with a threshold of 1 pixel. Disparities are scaled according to the provided scaling factor on the Middlebury dataset.



**Fig. 5.** Disparity map estimation results comparison on the KITTI stereo 2015 dataset.

**Image Separation.** As an anaglyph image is generated by concatenating the red channel from the left image and the green and blue channels from the right image, the original color can be found by warping the corresponding channels based on the estimated disparity maps. We leverage such a prior for de-anaglyph and adopt the post-processing step from Joulin *et.al.* [9] to handle occluded regions. Qualitative and quantitative comparison of image separation performance

are conducted on the Middlebury and KITTI datasets. We employ the Peak Signal-to-Noise Ratio (PSNR) to measure the image restoration quality.

Qualitative results for both datasets are provided in Fig. 6 and Fig. 7. Our method is able to recover colors in the regions where ambiguous colorization options exist as those areas rely more on the correspondence estimation, while other methods tend to fail in this case.



**Fig. 6.** Qualitative image separation results on the KITTI-2015 dataset. Top to bottom: Input, ground truth, result from Williém *et.al.* [8], our result. Our method successfully recovers the correct color of the large textureless region on the right of the image while the other method fails.

Table 3 and Table 4 report the performance comparison between our method and state-of-the-art de-anaglyph colorization methods: Joulin *et.al.* [9] and Williém *et.al.* [8] on the Middlebury dataset and on the KITTI dataset correspondingly. For the KITTI dataset, we calculated the mean PSNR throughout the total 200 images of the training set. Our method outperforms others with a notable margin. Joulin *et.al.* [9] is able to recover relatively good restoration results when the disparity level is small, such as Tsukuba, Venus, and KITTI. When the disparity level doubled, its performance drops quickly as for Cone and Teddy images. Different with Williém *et.al.* [8], which can only generate disparity maps at pixel level, our method is able to further optimize the disparity map to sub-pixel level, therefore achieves superior performance in both stereo computation and image restoration (separation).

Method	Tsukuba		Venus		Cones		Teddy	
	Left	Right	Left	Right	Left	Right	Left	Right
Joulin [9]	30.83	32.88	29.66	31.97	21.52	24.54	21.16	24.59
Williém [8]	30.32	31.12	31.41	34.93	23.68	26.17	23.14	31.05
Ours	<b>32.99</b>	<b>35.31</b>	<b>35.05</b>	<b>37.74</b>	<b>26.31</b>	<b>30.17</b>	<b>28.44</b>	<b>35.53</b>

**Table 3.** Performance comparisons (PSNR) in image separation (restoration) for the task of de-anaglyph on the Middlebury dataset.



**Fig. 7.** Qualitative comparison in image separation (restoration) on the Middlebury dataset. The first column shows the input anaglyph image and the ground truth image. The results of Williem *et.al.* [8] (top) and our method (bottom) with their corresponding error maps are shown in the second and the third column.

Dataset	View	Joulin [9]	Williem [8]	Ours
KITTI	Left	25.13	24.57	<b>26.30</b>
	Right	27.19	26.94	<b>28.76</b>

**Table 4.** Performance comparisons (PSNR) in image separation (restoration) for the task of de-anaglyph on the KITTI dataset.

**Anaglyph in the wild.** One of the advantages of conventional methods is their generalization capability. They can be easily adapt to different scenarios with or without parameter changes. Deep learning based methods, on the other hand, are more likely to have a bias on specific dataset. In this section, we provide qualitative evaluation of our method on anaglyph images downloaded from the Internet to illustrate the generalization capability of our method. Our method, even though trained on the KITTI dataset which is quite different from all these images, achieves reliable image separation results as demonstrated in Fig. 8. This further confirms the generalization ability of our network model.

### 4.3 Evaluation for double-vision unmixing

Here, we evaluate our proposed method for unmixing of double-vision image, where the input image is the average of a stereo pair. Similar to anaglyph, we evaluate our performance based on the estimated disparities and reconstructed stereo pair on the KITTI stereo 2015 dataset. For disparity evaluation, we use the



**Fig. 8.** Qualitative stereo computation and image separation results for real world anaglyph images downloaded from the Internet. Left to right: input anaglyph images, disparity maps of our method, image separation results of our method.


oracle disparity maps (that are computed with clean stereo pairs) as a reference in Fig. 9. The mean bad pixel ratio of our method is 6.67%, which is comparable with the oracle’s performance as 5.28%. For image separation, we take a layer separation method [26] as a reference. A quantitative comparison is shown in Table 5. Conventional layer separation methods tend to fail in this scenario as the statistic difference between the two mixed images is minor which violates the assumption of these methods. Qualitative results of our method are shown in Fig. 10.

Dataset	View	Li <i>et.al.</i> [26]	Ours
KITTI	Left	17.03	<b>24.38</b>
	Right	7.67	<b>24.55</b>


**Table 5.** Performance comparison in term of PSNR for the task of image restoration from double-vision images on the KITTI dataset.

## 5 Beyond Anaglyph and Double-Vision

Our problem definition also covers the problem of monocular depth estimation, which aims at estimating a depth map from a single image [27, 2, 28, 3]. Under this setup, the image composition operator  $f$  is defined as  $I = f(I_L, I_R) = I_L$  or



**Fig. 9.** Stereo computation results on the KITTI 2015 dataset. The oracle disparity map is computed with clean stereo images.



**Fig. 10.** Qualitative diplopia unmixing results by our proposed method. Top to bottom, left to right: input diplopia image, ground truth left image, restored left image by our method, and our estimated disparity map.

$I = f(I_L, I_R) = I_R$ , *i.e.*, the mixture image is the left image or the right image. Thus, monocular depth estimation is a special case of our problem definition.


We evaluated our framework for monocular depth estimation on the KITTI 2015 dataset. Quantitative results and qualitative results are provided in Table 6 and Fig. 11, where we compare our method with state-of-the-art methods [1], [29] and [2]. Our method, even designed for a much more general problem, outperforms both [1] and [29] and achieves quite comparable results with [2].

## 6 Conclusion

This paper has defined a novel problem of stereo computation from a single mixture image, where the goal is to separate a single mixture image into a pair of stereo images—from which a legitimate disparity map can be estimated. This problem definition naturally unifies a family of challenging and practical problems such as anaglyph, diplopia and monocular depth estimation. The problem goes beyond the scope of conventional image separation and stereo computation. We have presented a deep convolutional neural network based framework

Methods	Abs Rel	Sq Rel	RMSE	RMSE log	$\delta < 1.25$	$\delta < 1.25^2$	$\delta < 1.25^3$
Zhou <i>et.al.</i> [29]	0.183	1.595	6.709	0.270	0.734	0.902	0.959
Garg <i>et.al.</i> [1]	0.169	1.080	5.104	0.273	0.740	0.904	0.962
Godard <i>et.al.</i> [2]	0.140	0.976	4.471	0.232	0.818	0.931	0.969
Ours	0.126	0.835	3.971	0.207	0.845	0.942	0.975

**Table 6.** Monocular depth estimation results on the KITTI 2015 dataset using the split of Eigen *et.al.* [27]. Our model is trained on 22,600 stereo pairs from the KITTI raw dataset listed by [2] by 10 epochs. Depth metrics are from Eigen *et.al.* [27]. Our performance is better than the state-of-the-art method [2].



**Fig. 11.** Qualitative monocular estimation evaluations on the KITTI-2015 dataset: Top to bottom: left image, ground truth, results from Zhou *et.al.*[29], results from Garg *et.al.*[1], results from Godard *et.al.*[2] and our results. Since the ground truth depth points are very sparse, we interpolated it with a color guided depth painting method [30] for better visualization.

that jointly optimizes the image separation module and the stereo computation module. It is worth noting that we do not need ground truth disparity maps in network learning. In the future, we will explore additional problem setups such as “alpha-matting”. Other issues such as occlusion handling and extension to handle video should also be considered.

**Acknowledgements** Y. Zhong’s PhD scholarship is funded by Data61. Y. Dai is supported in part by National 1000 Young Talents Plan of China, Natural Science Foundation of China (61420106007, 61671387), and ARC grant (DE140100180). H. Li’s work is funded in part by ACRV (CE140100016). The authors are very grateful to NVIDIA’s generous gift of GPUs to ANU used in this research.

## References

1. Garg, R., Kumar, B.V., Carneiro, G., Reid, I.: Unsupervised cnn for single view depth estimation: Geometry to the rescue. In: Proc. Eur. Conf. Comp. Vis. (2016) 740–756
2. Godard, C., Mac Aodha, O., Brostow, G.J.: Unsupervised monocular depth estimation with left-right consistency. In: Proc. IEEE Conf. Comp. Vis. Patt. Recogn. (2017)
3. Li, B., Shen, C., Dai, Y., van den Hengel, A., He, M.: Depth and surface normal estimation from monocular images using regression on deep features and hierarchical crfs. In: Proc. IEEE Conf. Comp. Vis. Patt. Recogn. (June 2015)
4. Bronstein, A.M., Bronstein, M.M., Zibulevsky, M., Zeevi, Y.Y.: Sparse ica for blind separation of transmitted and reflected images. *Intl J. Imaging Science and Technology* **15**(1) (2005) 84–91
5. Yang, J., Li, H., Dai, Y., Tan, R.T.: Robust optical flow estimation of double-layer images under transparency or reflection. In: Proc. IEEE Conf. Comp. Vis. Patt. Recogn. (June 2016)
6. Li, Z., Tan, P., Tan, R.T., Zou, D., Zhou, S.Z., Cheong, L.F.: Simultaneous video defogging and stereo reconstruction. In: Proc. IEEE Conf. Comp. Vis. Patt. Recogn. (June 2015) 4988–4997
7. Levin, A., Lischinski, D., Weiss, Y.: A closed-form solution to natural image matting. *IEEE Trans. Pattern Anal. Mach. Intell.* **30**(2) (Feb 2008) 228–242
8. Williem, Raskar, R., Park, I.K.: Depth map estimation and colorization of anaglyph images using local color prior and reverse intensity distribution. In: Proc. IEEE Int. Conf. Comp. Vis. (Dec 2015) 3460–3468
9. Joulin, A., Kang, S.B.: Recovering stereo pairs from anaglyphs. In: Proc. IEEE Conf. Comp. Vis. Patt. Recogn. (2013) 289–296
10. Li, Y., Brown, M.S.: Single image layer separation using relative smoothness. In: Proc. IEEE Conf. Comp. Vis. Patt. Recogn. (June 2014)
11. Fan, Q., Yang, J., Hua, G., Chen, B., Wipf, D.: A generic deep architecture for single image reflection removal and image smoothing. In: Proc. IEEE Int. Conf. Comp. Vis. (2017)
12. Liu, C., Yuen, J., Torralba, A.: Sift flow: Dense correspondence across scenes and its applications. *IEEE Trans. Pattern Anal. Mach. Intell.* **33**(5) (May 2011) 978–994
13. Xie, J., Girshick, R., Farhadi, A.: Deep3d: Fully automatic 2d-to-3d video conversion with deep convolutional neural networks. In: Proc. Eur. Conf. Comp. Vis. (2016) 842–857
14. Iizuka, S., Simo-Serra, E., Ishikawa, H.: Let there be color!: Joint end-to-end learning of global and local image priors for automatic image colorization with simultaneous classification. *ACM Trans. Graph.* **35**(4) (July 2016) 110:1–110:11
15. Zhang, R., Isola, P., Efros, A.A.: Colorful image colorization. In: Proc. Eur. Conf. Comp. Vis. (2016)
16. Isola, P., Zhu, J.Y., Zhou, T., Efros, A.A.: Image-to-image translation with conditional adversarial networks. In: Proc. IEEE Conf. Comp. Vis. Patt. Recogn. (2017)
17. Levin, A., Weiss, Y.: User assisted separation of reflections from a single image using a sparsity prior. *IEEE Trans. Pattern Anal. Mach. Intell.* **29**(9) (September 2007) 1647–1654

18. Ronneberger, O., Fischer, P., Brox, T.: U-net: Convolutional networks for biomedical image segmentation. *International Conference Medical Image Computing and Computer-Assisted Intervention* (2015) 234–241
19. Zhong, Y., Dai, Y., Li, H.: Self-supervised learning for stereo matching with self-improving ability. In: *arXiv:1709.00930*. (2017)
20. Wang, Z., Bovik, A.C., Sheikh, H.R., Simoncelli, E.P.: Image quality assessment: from error visibility to structural similarity. *IEEE Trans. Image Proc.* **13**(4) (April 2004) 600–612
21. Jaderberg, M., Simonyan, K., Zisserman, A., kavukcuoglu, k.: Spatial transformer networks. In: *Proc. Adv. Neural Inf. Process. Syst.* (2015) 2017–2025
22. Abadi, M., Agarwal, A., *et.al.*, P.B.: Tensorflow: Large-scale machine learning on heterogeneous distributed systems. *CoRR* **abs/1603.04467** (2016)
23. Tieleman, T., Hinton, G.: Lecture 6.5—RmsProp: Divide the gradient by a running average of its recent magnitude. *COURSERA: Neural Networks for Machine Learning* (2012)
24. Scharstein, D., Szeliski, R.: A taxonomy and evaluation of dense two-frame stereo correspondence algorithms. *Int. J. Comp. Vis.* **47**(1-3) (April 2002) 7–42
25. Menze, M., Geiger, A.: Object scene flow for autonomous vehicles. In: *Proc. IEEE Conf. Comp. Vis. Patt. Recogn.* (2015)
26. Li, Y., Brown, M.S.: Single image layer separation using relative smoothness. In: *Proc. IEEE Conf. Comp. Vis. Patt. Recogn.* (2014) 2752–2759
27. Eigen, D., Puhrsch, C., Fergus, R.: Depth map prediction from a single image using a multi-scale deep network. In: *Proc. Adv. Neural Inf. Process. Syst.* (2014)
28. Li, B., Dai, Y., He, M.: Monocular depth estimation with hierarchical fusion of dilated cnns and soft-weighted-sum inference. *Pattern Recognition* **83** (2018) 328 – 339
29. Zhou, T., Brown, M., Snavely, N., Lowe, D.G.: Unsupervised learning of depth and ego-motion from video. In: *Proc. IEEE Conf. Comp. Vis. Patt. Recogn.* (2017)
30. Yang, J., Ye, X., Li, K., Hou, C., Wang, Y.: Color-guided depth recovery from rgb-d data using an adaptive autoregressive model. In: *IEEE Trans. Image Proc.* Volume 23. (Aug 2014) 3443–3458