

SymmNet: A Symmetric Convolutional Neural Network for Occlusion Detection

Ang Li
bennie.522@stu.xjtu.edu.cn
Zejian Yuan
yuan.ze.jian@xjtu.edu.cn

Institute of Artificial Intelligence and
Robotics
Xi'an Jiaotong University
Xi'an, China

Abstract

Detecting the occlusion from stereo images or video frames is essential to many computer vision applications. Previous efforts focus on bundling it with the computation of disparity or optical flow, leading to a chicken-and-egg problem. In this paper, we leverage a convolutional neural network to liberate the occlusion detection task from the interleaved, traditional calculation framework. We propose a Symmetric Network (SymmNet) to directly exploit information from an image pair, without estimating disparity or motion in advance. The proposed network is structurally left-right symmetric to learn the binocular occlusion simultaneously, aimed at jointly improving both results. The extensive experiments show that our model achieves state-of-the-art results on detecting the stereo and motion occlusion.

1 Introduction

The problem of localizing the occluded and non-occluded areas over multi-view images or video sequences is of great interest for many computer vision tasks. The two most related tasks are stereo computation and optical flow estimation. The occluded pixels violate the inter-image correspondence constraint, resulting in ambiguous matching. State-of-the-art stereo and optical flow methods benefit from occlusion detection, either by explicitly excluding occluded pixels from disparity and motion computation [3, 8, 36] or by repairing these regions afterward [13, 44, 45]. Occlusion detection also has been applied to help improve the performance of other tasks, such as action recognition [40], object tracking [27] and 3D reconstruction [53].

Most of the existing methods take disparity or optical flow as an intermediary to estimate occlusion. The simplest but widely used left-right-cross-checking (LRC) [10, 37, 45] directly reasons occlusion from pre-computed disparity. This method assumes that the disparities of corresponding points in the left and right image agree with each other except for the pixels that arise from occlusion. For LRC, however, the lack of occlusion prior introduces difficulty into accurate disparity estimation. The imperfect disparity in turn easily leads to erroneous occlusion detection, and there is no chance to revise the result. Other approaches [9, 21, 36, 38, 43] iteratively refine their occlusion map by alternatively improving the disparity or motion accuracy. Kolmogorov and Zabih [21] explicitly model the occlusion based on the unique matching constraint and incorporate it into an energy-based disparity estimation

framework. Wang *et al.* [39] borrow the power from deep learning. They integrate a warp module for occlusion inference into an end-to-end trainable motion estimation network. This CNN-based method improves occlusion estimation coherently as learning accurate motion. Unlike methods above deterministically deciding occlusion from disparity or optical flow, learning based method [14] uses initial motion estimations as sources to produce features for a random forest occlusion classifier. Pérez-Rúa *et al.* [29] make plausible motions serve as a "soft" evidence for their occlusion model which is based on spatio-temporal reconstruction.

To some degree, previous occlusion detectors rely on an initial estimation of disparity or optical flow. Nevertheless estimating disparity or optical flow is definitely not an easy task due to the noise, low or repetitive textures and even occlusion itself. This motivates us to explore a solution to detect occlusion directly from stereo images or sequential frames. In this paper, we focus on the stereo situation. Inspired by the success of convolutional neural network (CNN) in the field of monocular depth [7, 22, 42] and camera localization [18, 47], we leverage CNN to free occlusion detection from disparity estimation.

We regard occlusion detection as a binary classification problem like [14] and propose a Symmetry Network (SymmNet) as the classifier. Compared with methods that infer occlusion after regressing the continuous disparity values or classifying disparity from hundreds of discrete labels, the high precision requirement is relaxed when directly determining the binary occlusion labels. The SymmNet is an hourglass architecture to exploit information from binocular images. We make the network left-right symmetrically infer the binocular occlusion cooperatively, so the left and right results can be jointly improved.

The contributions of this paper are mainly three-fold:

- This work is, to the best of our knowledge, the first to directly estimate occlusion regions from images without preliminary disparity or motion knowledge.
- We propose a SymmNet which takes an image pair as input to cooperatively reason binocular occlusion.
- We conduct an exhausted experimental analysis to verify our design, and our method achieves promising results for detecting stereo and motion occlusion.

2 Proposed Model

In binocular viewing of a scene, it is a common phenomenon that some portion of the scene can only be seen from one view. Fig. 1 shows an example. When projecting the points in the scene onto the two views, the point **b** appears only in the left image I_l and point **e** only in the right image I_r . The task of pixel-wise occlusion detection is to find these monocularly visible regions given a stereo image pair. The monocularly visible regions are so-called **occlusion**.

2.1 Occlusion detection with CNN

To infer the occlusion, what information is necessary? And is CNN capable of learning it? We argue that it is possible for CNN to learn occlusion from only one view's image in a stereo pair, while binocular images can provide more information.

Monocular clues.¹ It is theoretically workable to detect occlusion by digging out monocular information. First of all, a monocular image contains depth and camera configured information, which are two basic origins of occlusion. As shown in Fig. 1, points **a**, **b**, **c**, and **d**

¹The clues from one image in a stereo image pair, rather than an arbitrary monocular image.

are visible in image I_l , once we know the depth of these points as well as the pose of camera O_r , we can project them onto the right image I_r . **b** and **c** project to the same location, so the farther point **b** on I_l can be judged as occlusion. Secondly, the local structure of an image can assist in detection, since occlusion map tends to have a specific structure corresponding to the image. For example, occlusion most likely lies just adjacent to the edge of the closer object [22, 54] (except that the closer object is a thin stick), and the outer edge of occlusion always has the similar shape with the object edge. Furthermore, occlusion regions exhibit spatial coherence. Seldom does an isolated occluded pixel exist [34]. Fortunately, researchers have achieved prominent results for estimating monocular depth [2, 22, 22], camera pose [2, 18] and detecting edge [2, 41] by applying CNN on these problems. This suggests that we could deal with the occlusion detection task with deep learning from a single view image.

Binocular clues. Although learning from a monocular image is theoretically workable, a network bears too much uncertainty to effectively encode all the necessary information including the scene geometry, camera settings, and pictorial structure. Utilizing binocular images instead can better restrain this problem and potentially facilitate the detection in following aspects: (1) Occlusion in one image is the regions that have no correspondence in the other. Inspired by FlowNet [6] which learns optical flow from two stacked frames, we consider that feeding binocular images gives the neural network an opportunity to learn the correspondence. (2) Occlusion in one image and the depth of the other is symmetrically consistent, that is, one can trace the occlusion back to the other view’s depth. As indicated in Fig. 1, inversely project I_r to the left view according to the right depth, the being projected points (**a**_l, **c**_l, and **d**_l) are non-occluded, otherwise (the point **b**_l) is occluded. (3) Binocular images contain the information about the relative camera pose between two views, and CNN has the ability to learn it [25, 47]. Learning relative pose is favorable for enhancing the robustness to the changing of camera configurations.

Given a binocular image pair, how to design an occlusion detection network? Inspired by multi-task learning [6], we propose to simultaneously predict the binocular occlusion. Jointly inferring the occlusion for both views is helpful to improve the prediction accuracy, since it enables consistency cross-checking between two streams. This lies in the fact that occlusion can be inferred from the depth of either view. In another word, the depth of an image is sufficient for reasoning both views’ occlusion.

Based on the observations above, we propose a **Symmetric Network (SymmNet)** which makes stacked binocular images flow through a structurally left-right symmetric neural network to predict binocular occlusion. Fig. 2 illustrates the brief architecture of the proposed network, we will introduce details in the following sections.

2.2 Network architecture

We follow FlowNet [6] to build a fully convolutional network which consists of a contractive part and an expanding part with skip connections between them. The detailed layer-by-layer definition is listed in Table 1. Since determining occlusion probably relies on the information




Figure 1: **The occlusion reasoning diagram.** When binocular cameras O_l and O_r capture the scene, **b**_l and **c**_r are occlusions on the corresponding stereo images I_l and I_r .

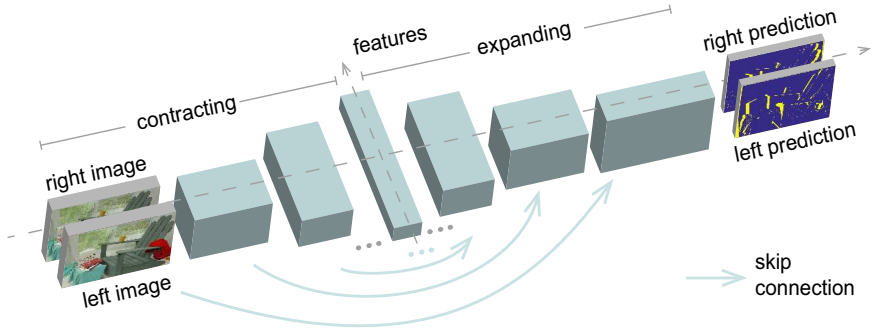


Figure 2: **SymmNet Architecture.** Intermediate layers and residual connections are omitted in the illustration.

from a wide field of view, the contractive part sub-samples the features to encode large structures. It contains 6 down-sampling layers with strides of 2 to progressively increase the receptive field and sub-samples the spatial size of feature maps by a factor of 64 in total. To obtain pixel-wise predictions with the original input resolution, in the expanding part we employ 6 deconvolutional layers to up-sample features. Each down- and up-sampling layer is followed with a convolutional layer for smoother results. For the sake of keeping fine local information, lower level features take part in higher level decoding through skip connections. ReLU comes after each layer to better cope with the gradient vanishing problem.

Being different from FlowNet, we take several strategies to prune the network for computational efficiency. Firstly, we compress the number of feature channels. The first layer has 16 filters. The length of the feature is doubled every time when the spatial size of feature map is down-sampled, and reaches a maximum of 512 filters at the last layer of the contractive part. Further, we replace the concatenation in the skip connection with addition. The feature length in the expanding part is correspondingly reduced to match that in the contractive part.

Another modification is that we include an extra up-sampling module at the end of the expanding part leading to full-resolution outputs rather than half-resolution. With up-sampled to the full resolution, original image features are concatenated with the features to the last convolutional layer. This is for the consideration that low-level features from images can aid occlusion localization.

A prediction layer follows the expanding part in series to generate a 4-channel output for two views' pixel-wise occlusion classification. Every 2 channels are normalized as probabilities by *softmax*. Then we can get the occlusion probability P_L and P_R for the left view L and the right view R . A pixel \mathbf{p} is inferred as occlusion if $P(\mathbf{p})$ is larger than a threshold τ .

2.3 Training details

To jointly train the binocular occlusion, we use the total binary-cross-entropy loss of both views as objective:

$$\mathbf{L} = -\frac{1}{2} \left(w_L^o \sum_{\mathbf{p}} \mathbf{1}(O_L(\mathbf{p}) = 1) \log(P_L(\mathbf{p})) + w_L^{\bar{o}} \sum_{\mathbf{p}} \mathbf{1}(O_L(\mathbf{p}) = 0) \log(1 - P_L(\mathbf{p})) \right. \\ \left. + w_R^o \sum_{\mathbf{p}} \mathbf{1}(O_R(\mathbf{p}) = 1) \log(P_R(\mathbf{p})) + w_R^{\bar{o}} \sum_{\mathbf{p}} \mathbf{1}(O_R(\mathbf{p}) = 0) \log(1 - P_R(\mathbf{p})) \right), \quad (1)$$

Name	Kernel	Str.	Ch I/O	OutRes	Input
Input					
input			6/6	$H \times W$	image pair
Contracting					
dwnsp1	8×8	2	6/16	$1/2H \times 1/2W$	input
conv1	3×3	1	16/16	$1/2H \times 1/2W$	dwnsp1
dwnsp2	6×6	2	16/32	$1/4H \times 1/4W$	conv1
conv2	3×3	1	32/32	$1/4H \times 1/4W$	dwnsp2
dwnsp3	6×6	2	32/64	$1/8H \times 1/8W$	conv2
conv3	3×3	1	64/64	$1/8H \times 1/8W$	dwnsp3
dwnsp4	4×4	2	64/128	$1/16H \times 1/16W$	conv3
conv4	3×3	1	128/128	$1/16H \times 1/16W$	dwnsp4
dwnsp5	4×4	2	128/256	$1/32H \times 1/32W$	conv4
conv5	3×3	1	256/256	$1/32H \times 1/32W$	dwnsp5
dwnsp6	4×4	2	256/512	$1/64H \times 1/64W$	conv5
conv6	3×3	1	512/512	$1/64H \times 1/64W$	dwnsp6

Name	Kernel	Str.	Ch I/O	OutRes	Input
Expanding					
upsp5	4×4	2	512/256	$1/32H \times 1/32W$	conv6
iconv5	3×3	1	256/256	$1/32H \times 1/32W$	upsp5+conv5
upsp4	4×4	2	256/128	$1/16H \times 1/16W$	iconv5
iconv4	3×3	1	128/128	$1/16H \times 1/16W$	upsp4+conv4
upsp3	4×4	2	128/64	$1/8H \times 1/8W$	iconv4
iconv3	3×3	1	64/64	$1/8H \times 1/8W$	upsp3+conv3
upsp2	4×4	2	64/32	$1/4H \times 1/4W$	iconv3
iconv2	3×3	1	32/32	$1/4H \times 1/4W$	upsp2+conv2
upsp1	4×4	2	32/16	$1/2H \times 1/2W$	iconv2
iconv1	3×3	1	16/16	$1/2H \times 1/2W$	upsp1+conv1
upsp0	4×4	2	16/8	$H \times W$	iconv1
iconv0	3×3	1	14/8	$H \times W$	upsp0 \oplus input
Prediction					
pr	3×3	1	8/4	$H \times W$	iconv0

Table 1: SymmNet architecture summary. Each layer except for the prediction layer pr is followed by ReLU. pr layer is followed by *softmax* to generate probability. This table is arranged from top to bottom, left to right. $+$ is the addition operation, \oplus is the concatenation operation in skip connection.

where O is ground-truth occlusion, $\mathbf{1}(\cdot)$ is indicating function, w^c is a class weight to make the loss adapt to the unbalanced number of occlusion and non-occlusion pixels. We adopt the bounded class weight [28] $w^c = 1/\ln(\varepsilon + q^c)$, where q^c is the proportion of class c (occlusion o or non-occlusion \bar{o}) in the training batch. ε is a hyper-parameter to limit the weight range.

We trained our model on the SceneFlow dataset [24], which consists of stereo image pairs rendered from synthetic sequences. The dataset is suitable for training the network for two reasons. One is that this dataset contains 35, 454 training and 4, 370 test pairs. It is large enough to train the model without over-fitting. The other reason is that it provides dense, perfect ground-truth disparity for both views, which can be used to generate binocular ground-truth occlusion. The ground-truth O_v for a view v is obtained by checking the left-right-consistency between its ground-truth disparity D_v and the other view’s, as

$$O_v(\mathbf{p}) = \mathbf{1}(|D_v(\mathbf{p}) - \hat{D}_v(\mathbf{p})| > \delta), \quad v \in \{L, R\}. \quad (2)$$

\hat{D}_v is the warped disparity from the other view v' . It is obtained by bilinear sampling mechanism [16] as $\hat{D}_v(\mathbf{p}) = \sum_{i \in \{t, b\}, j \in \{l, r\}} \omega_{ij} D_{v'}(\mathbf{t}^{ij})$. \mathbf{t}^{ij} is the 4-pixel neighbors of \mathbf{t} , which is the corresponding position of \mathbf{p} on view v' based on $D_v(\mathbf{p})$. ω_{ij} is the interpolation weight and $\sum_{i, j} \omega_{i, j} = 1$.

Training samples are randomly cropped patches with a spatial size of $H = 256$ and $W = 768$. The cropping process is for computational restriction. Besides, it is a data augmentation means, since the shape of the out-of-image occlusion at image boundary varies as cropping a patch at different locations. Accordingly, it should be noted that the ground-truth computation has to be done after cropping due to the varying out-of-image occlusion.

The network were optimized using the Adam [20] method ($\beta_1 = 0.9$ and $\beta_2 = 0.99$) and a constant learning rate of 1×10^{-2} for 10 epoches. The training batch contains 16 samples. ε in the class weight is empirically set to 1.5, δ in Eq. (2) is set to 1.

3 Experiment

In this section, we first test several variants of our method to verify the proposed pipeline. Then we compare the overall performance with several existing methods on SceneFlow [24]




Figure 3: **Example results of different architecture variants.** The first row in MonoNet(L/R) column is the result of MonoNetL, the second row is MonoNetR. The pink arrow in MonoNetL points to the fake occlusion occurring at the image edge. The arrow in MonoNetR points to the erroneous occlusion shape.

and Middlebury [10, 26, 30, 31, 32] dataset. Furthermore, we examine our model’s capacity to learn motion occlusion on MPI Sintel dataset [15]. We finally report the time and memory requirement of our architecture.

For evaluation, we report three metrics commonly used in occlusion detection task, which are precision (the percentage of true occluded pixels in detected occlusion), recall (the percentage of the detected occlusion pixels in the occluded regions) and Fscore (the harmonic average of precision and recall). When predicting occlusion, the threshold τ is set to 0.5 unless otherwise specified.

3.1 Architecture Analysis

To justify our design choices, we test several model variants of **SymmNet** on the SceneFlow test set. To be fair, we keep the parameter number of different architectures all the same except for the input and output layers. In Fig. 3, we provide example results and in Fig. 4 we visualize the precision-recall (PR) curves.

MonoNet(L/R). To investigate the role of monocular image input, we modify the **SymmNet** to take the single left image and right image separately as input and to predict the left occlusion (MonoNetL and MonoNetR). Fig. 3 shows that either image serves occlusion detection but in a different manner. The image of the homogeneous view, i.e., the left image, tends to provide more information about the object edges. This is a useful clue for determining the shape of occlusion, while gives rise to fake occlusion. The image of the cross view helps to tell the true occluded edges, while performs poor at estimating the shape.

SiameseNet. A Siamese architecture comprises two sub-networks with shared weights [4]. This architecture is widely adopted in the highly related task of stereo matching, in which each of the two branches concentrates on one view (the left or right) to extract unary features [23, 15] or to regularize the cost volume [19]. The Siamese structure, treating the binocular views indiscriminately, is efficient for stereo methods to encode the shared-knowledge as well as reduce the computational requirement. For occlusion detection, we construct a similar Siamese variant that each branch takes one view as input and outputs the corresponding occlusion map. The disordered results in Fig. 3 shows its incompetency for occlusion detection. This is because reasoning occlusion is subject to the viewpoint and




Figure 4: **Precision-recall (PR) curves.** The max Fscore on a PR curve is annotated. The curve that is closer to the upper-right-hand corner is better. **SymmNet** outperforms all of its variants.

we are aiming at finding the view-subjected discrimination. The shared-weights disable the network to learn the distinctive information. In addition, seeing one view for each shared-branch enforces the network only to involve monocular cues while omit the vital binocular information, such as the relative pose of the fellow camera.

AlterNet & HalfNet. Our model jointly learns binocular occlusion in order to make two streams help each other learn better. To verify this design, we construct two variants based on SymmNet for comparison. One is AlterNet that only outputs occlusion for a single view, while we iteratively interchange the stacked order of two input images to alternatively learn either the left or right occlusion as training. The other is HalfNet which consists of two separate networks, one for learning each view’s occlusion independently. Each sub-network in the HalfNet still takes binocular images as input, but the length of feature channel is half of that in SymmNet so as to keep the total model volume unchanged. AlterNet gets into trouble when learning the alternating views. As shown in Fig. 3, the result of the left view is approximately correct, while the right result gets a mess. HalfNet equally estimates both occlusion maps with good quality, while numerically performs slightly worse than SymmNet.

LRCNet. An alternative method to detect occlusion is accurately estimating disparity first and inferring occlusion from disparity instead. We replace the prediction layer in SymmNet with a regression layer to make the network learn binocular disparity and then apply LRC on the disparity to infer occlusion. We call this network LRCNet. This network lacks a module to directly regularize the shape of occlusion, thus there are evident holes in the occlusion regions as shown in Fig. 3. Moreover, the disparity results directly determine the occlusion detection quality, while learning disparity seems not easy. Among the variants, this network is the only one that is used for learning disparity rather than occlusion, whereas its performance is rather poor.

Discussion. Our SymmNet directly models occlusion from input images, rather than inferring occlusion at the following stage of disparity computation. This design, on the one hand, eases the problem in terms of engineering, as can be seen from the significant gap between the PR curves of LRCNet and ours in Fig. 4. On the other hand, it can be integrated into the disparity estimation framework at the very beginning, as suggested by Anderson and Nakayama that one senses occlusion at the earliest stages in the binocular visual system [14].

We learn complementary information from binocular images. Both images are indispensable for precise estimation, particularly for eliminating the fake occlusion and keeping the shape of occlusion. More importantly, two images work together, providing relative cues which are necessary for determining the discrimination between two views. This can be verified by comparing the results of SymmNet with MonoNets and SiameseNet in Fig. 3.

Furthermore, it should be noted that although we follow FlowNet to construct our structurally similar architecture, these two networks are different in nature. FlowNet only predicts for the reference view, while our SymmNet uses a unified architecture to reason the occlusion for both views. This symmetric design enables knowledge to transfer between the two views and boosts their performance coherently. The PR curve of our model covers the curves of all the variants from the upper-right, which shows the reasonability of our whole design.

3.2 Overall Performance

We compare our overall performance with that of two other occlusion detectors. We first run the method of Kolmogorov and Zabih [21] (KZ) which enforces the uniqueness constraint to detect the un-matched pixels as occlusion. We also compare with the LRC method. The initial disparity for LRC is obtained as MC-CNN [45], i.e., by extracting and matching the




Figure 5: **Qualitative comparison on SceneFlow dataset.** The true positive estimations are labelled in cyan, the false negative in magenta and the false positive in yellow. The green boxes in the first row label the errors occur at slanted planes, and the errors at textureless regions in the second row. The white boxes in the last column show the failure cases of our method.

deep features, followed cross-based cost aggregation [46] and semiglobal matching [9]. We use the code provided by the authors of these methods.

Validation on SceneFlow

We first evaluate the performance on SceneFlow test set. For a fair comparison, we fine-tune the MC-CNN model on the SceneFlow training set and report the best scores we can get. Fig. 5 exhibits the qualitative results. The performance of KZ and LRC relies on the quality of the initial disparity. KZ fails to recover the disparity of the slanted plane due to the first-order smoothness prior, MC-CNN encounters difficulty at matching the large textureless regions. Consequently, the occlusion detection error easily appears in the corresponding regions as shown in the first two rows in Fig. 5. Our method, directly predicting the occlusion regions, is free from the influence of the initial disparity estimation. Quantitative results also show the superiority of our method as summarized in Table 2.

Although our method obtains largely proper results, we fail to make the occlusion precisely coincide with the image edges, as shown in Fig. 5. We ascribe this failure to the smoothing effect brought by the convolution and contraction operations in the proposed architecture. Trying to rethink an architecture to preserve more details may help alleviate the problem. In addition, explicit matching evidence, rather than our implicit matching cues from stacked images, is potentially beneficial to accurate location.

Validation on Middlebury

Middlebury dataset provides stereo image pairs with dense ground-truth disparity of indoor scenes under controlled lighting conditions. Compared to the SceneFlow dataset, the scenes are more realistic, the lighting conditions and exposure settings are more complex. We collect 2845 image pairs with ground-truth disparity of both views, and split the collections into training set and validation set to conduct 10-fold-cross-validation. For our method, we test two configurations: (1) directly applying the model trained on SceneFlow and (2) fine-tuning the model for another 50 epoches on the Middlebury training set with learning rate set to 1×10^{-3} (SymmNet-MB). ϵ in the class weight is adjusted to 1.2 for the smaller occlusion.

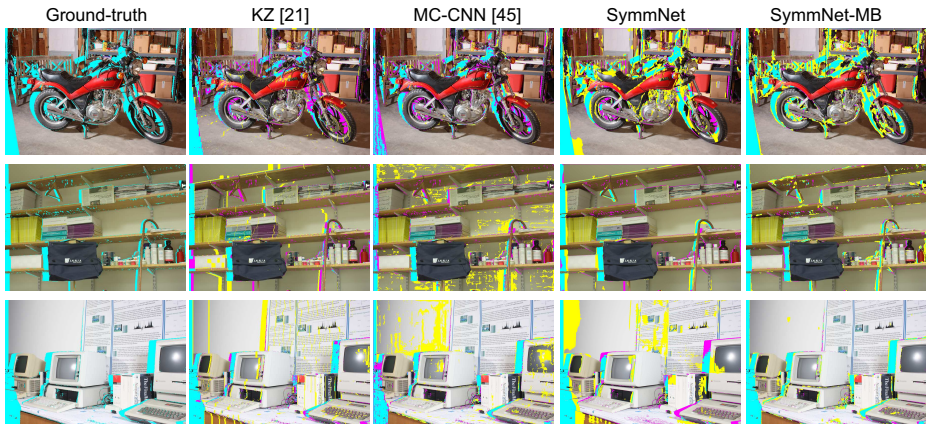


Figure 6: **Qualitative comparison on Middlebury dataset.** The true positive estimations are labelled in cyan, the false negative in magenta and the false positive in yellow.

Stereo Occlusion							Motion Occlusion		
	SceneFlow			Middlebury			MPI		
	Precision	Recall	F-score	Precision	Recall	F-score	Oracle 69	Global 69	
KZ [21]	0.554	0.609	0.580	0.585	0.628	0.605	Learning [24]	0.535	0.448
MC-CNN[45]+LRC	0.772	0.836	0.802	0.660	0.664	0.652	Depth Order [35]	0.465	0.449
SymmNet	0.799	0.919	0.873	0.584	0.737	0.666	Pérez-Rúa <i>et al.</i> [29]	0.550	0.540
SymmNet-MB	-	-	-	0.810	0.849	0.828	Ours-MPI	0.665	0.642

Table 2: **Quantitative evaluations.** All the evaluations are the higher, the better. We highlight the best scores in bold.

The qualitative results are shown in Fig. 6 and quantitative results are given in Table 2. Our fine-tuned model outperforms other methods on all the evaluation indexes. It is worth noting that our method reveals some robustness to the variation of camera configurations and environment. Since even without fine-tuning, our method can also generate comparable results.

3.3 Motion occlusion detection

Although our model is designed for detecting the occlusion in stereo settings, it can be effortlessly adapted to the task of motion occlusion detection by taking two consecutive frames as inputs. We demonstrate this ability on the MPI Sintel dataset. This dataset contains 69 sequences (3123 image pairs) equipped with ground-truth occlusion maps. We divide the dataset into training set and validation set for 10-fold-cross-validation. The hyper parameter c is set to 1.01 to fit the extremely unbalanced occlusion ratio in this experiment.

We compare with three motion occlusion detectors: the learning based method [24], a depth order based method [35] and the spatial-temporal reconstruction model of Pérez-Rúa *et al.* [29]. Following the evaluation methodology of Pérez-Rúa *et al.*, we test the average F-score over all 69 sequences when the threshold τ is set to maximize F-score (Oracle 69) and to 0.5 (Global 69). Our method excels all the other methods on both settings as shown in Table 2.

We provide several detection results in Fig. 7. Even though the true occlusion regions are much smaller and finer than those in stereo, we can still make a good prediction.




Figure 7: **Qualitative results of our method on MPI dataset.** From left to right: Average image of the two input frames; Occlusion ground-truth; Predicted occlusion map with proposed method.

3.4 Runtime and memory requirement

We test the runtime of our PyTorch implementation on a single NVIDIA Tesla M40 GPU. Training on SceneFlow dataset can be finished in two days. It takes 0.07s and requires 651M graphic memory to predict an image pair in size of 540×960 . The low requirement of time and memory makes our model an optional preprocess module for other tasks such as object tracking, human pose estimation and action recognition.

4 Conclusion

We have proposed a CNN model called SymmNet to detect occlusion from stereo images or video sequences. Unlike the traditional occlusion detectors which infer occluded pixels from pre-computed disparity or optical flow, our model directly learns from original images. The proposed SymmNet is left-right symmetric to jointly learn binocular occlusion by cooperatively extracting the binocular information. The experiment results have demonstrated the good ability of our method for stereo and motion occlusion detection.

We believe the proposed occlusion detector can be extended to facilitate other applications, such as stereo and optical flow. It would be an interesting future work to investigate the auxiliary role of occlusion based on our method.

Acknowledgement: This work was supported by the National Key R&D Program of China (No.2016YFB1001001) and the National Natural Science Foundation of China (No.61573280, No.91648121).

References

- [1] Barton L Anderson and Ken Nakayama. Toward a general theory of stereopsis: binocular matching, occluding contours, and fusion. *Psychological review*, 101(3):414, 1994.
- [2] Gedas Bertasius, Jianbo Shi, and Lorenzo Torresani. Deepedge: A multi-scale bifurcated deep network for top-down contour detection. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, 2015.
- [3] Michael Bleyer, Carsten Rother, and Pushmeet Kohli. Surface stereo with soft segmentation. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, 2010.
- [4] Jane Bromley, Isabelle Guyon, Yann LeCun, Eduard Säckinger, and Roopak Shah.

- Signature verification using a siamese time delay neural network. In *Proceedings of the Annual Conference on Neural Information Processing Systems (NIPS)*, 1993.
- [5] Rich Caruana. Multitask learning. In *Learning to learn*, pages 95–133. 1998.
- [6] Alexey Dosovitskiy, Philipp Fischer, Eddy Ilg, Philip Häusser, Caner Hazirbas, Vladimir Golkov, Patrick van der Smagt, Daniel Cremers, and Thomas Brox. FlowNet: Learning optical flow with convolutional networks. In *Proceedings of the IEEE International Conference on Computer Vision (ICCV)*, 2015.
- [7] David Eigen, Christian Puhrsch, and Rob Fergus. Depth map prediction from a single image using a multi-scale deep network. In *Proceedings of the Annual Conference on Neural Information Processing systems (NIPS)*, 2014.
- [8] Fabrice Heitz and Patrick Bouthemy. Multimodal estimation of discontinuous optical flow using markov random fields. *IEEE Transactions on Pattern Analysis and Machine Intelligence (TPAMI)*, 15(12):1217–1232, 1993.
- [9] Heiko Hirschmüller. Stereo processing by semiglobal matching and mutual information. *IEEE Transactions on Pattern Analysis and Machine Intelligence (TPAMI)*, 30(2): 328–341, 2008.
- [10] Heiko Hirschmüller and Daniel Scharstein. Evaluation of cost functions for stereo matching. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, 2007.
- [11] Heiko Hirschmüller, Peter R. Innocent, and Jonathan M. Garibaldi. Real-time correlation-based stereo vision with reduced border errors. *International Journal of Computer Vision (IJCV)*, 47(1-3):229–246, 2002.
- [12] Derek Hoiem, Alexei A. Efros, and Martial Hebert. Recovering occlusion boundaries from an image. *International Journal of Computer Vision (IJCV)*, 91(3):328–346, 2011.
- [13] Asmaa Hosni, Michael Bleyer, Margrit Gelautz, and Christoph Rhemann. Local stereo matching using geodesic support weights. In *Proceedings of the International Conference on Image Processing (ICIP)*, 2009.
- [14] Ahmad Humayun, Oisín Mac Aodha, and Gabriel J. Brostow. Learning to find occlusion regions. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, 2011.
- [15] Natan Jacobson, Yoav Freund, and Truong Q. Nguyen. An online learning approach to occlusion boundary detection. *IEEE Transactions Image Processing (TIP)*, 21(1): 252–261, 2012.
- [16] Max Jaderberg, Karen Simonyan, Andrew Zisserman, and Koray Kavukcuoglu. Spatial transformer networks. In *Proceedings of the Annual Conference on Neural Information Processing Systems (NIPS)*, 2015.
- [17] Alex Kendall and Roberto Cipolla. Geometric loss functions for camera pose regression with deep learning. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, 2017.

- [18] Alex Kendall, Matthew Grimes, and Roberto Cipolla. Posenet: A convolutional network for real-time 6-dof camera relocalization. In *Proceedings of the IEEE International Conference on Computer Vision (ICCV)*, 2015.
- [19] Alex Kendall, Hayk Martirosyan, Saumitro Dasgupta, and Peter Henry. End-to-end learning of geometry and context for deep stereo regression. In *Proceedings of the IEEE International Conference on Computer Vision (ICCV)*, 2017.
- [20] Diederik P Kingma and Jimmy Ba. Adam: A method for stochastic optimization. *arXiv preprint arXiv:1412.6980*, 2014.
- [21] Vladimir Kolmogorov and Ramin Zabih. Computing visual correspondence with occlusions using graph cuts. In *Proceedings of the IEEE International Conference on Computer Vision (ICCV)*, 2001.
- [22] Fayao Liu, Chunhua Shen, and Guosheng Lin. Deep convolutional neural fields for depth estimation from a single image. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, 2015.
- [23] Wenjie Luo, Alexander G. Schwing, and Raquel Urtasun. Efficient deep learning for stereo matching. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, 2016.
- [24] Nikolaus Mayer, Eddy Ilg, Philip Häusser, Philipp Fischer, Daniel Cremers, Alexey Dosovitskiy, and Thomas Brox. A large dataset to train convolutional networks for disparity, optical flow, and scene flow estimation. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, 2016.
- [25] Iaroslav Melekhov, Juha Ylioinas, Juho Kannala, and Esa Rahtu. Relative camera pose estimation using convolutional neural networks. In *Proceedings of the International Conference on Advanced Concepts for Intelligent Vision Systems (ACIVS)*, 2017.
- [26] Christopher J. Pal, Jerod J. Weinman, Lam C. Tran, and Daniel Scharstein. On learning conditional random fields for stereo - exploring model structures and approximate inference. *International Journal of Computer Vision (IJCV)*, 99(3):319–337, 2012.
- [27] Jiyang Pan and Bo Hu. Robust occlusion handling in object tracking. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, 2007.
- [28] Adam Paszke, Abhishek Chaurasia, Sangpil Kim, and Eugenio Culurciello. Enet: A deep neural network architecture for real-time semantic segmentation. *arXiv preprint arXiv:1606.02147*, 2016.
- [29] Juan-Manuel Pérez-Rúa, Tomas Crivelli, Patrick Bouthemy, and Patrick Pérez. Determining occlusions from space and time image reconstructions. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, 2016.
- [30] Daniel Scharstein and Richard Szeliski. A taxonomy and evaluation of dense two-frame stereo correspondence algorithms. *International Journal of Computer Vision (IJCV)*, 47(1-3):7–42, 2002.

- [31] Daniel Scharstein and Richard Szeliski. High-accuracy stereo depth maps using structured light. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, 2003.
- [32] Daniel Scharstein, Heiko Hirschmüller, York Kitajima, Greg Krathwohl, Nera Nesić, Xi Wang, and Porter Westling. High-resolution stereo datasets with subpixel-accurate ground truth. In *Proceedings of the German Conference on Pattern Recognition (GCPR)*, 2014.
- [33] Johannes L. Schönberger, Enliang Zheng, Jan-Michael Frahm, and Marc Pollefeys. Pixelwise view selection for unstructured multi-view stereo. In *Proceedings of the European Conference on Computer Vision (ECCV)*, 2016.
- [34] Andrew N. Stein and Martial Hebert. Occlusion boundaries from motion: Low-level detection and mid-level reasoning. *International Journal of Computer Vision (IJCV)*, 82(3):325–357, 2009.
- [35] Deqing Sun, Erik B. Sudderth, and Michael J. Black. Layered image motion with explicit occlusions, temporal consistency, and depth ordering. In *Proceedings of the Annual Conference on Neural Information Processing Systems (NIPS)*, 2010.
- [36] Jian Sun, Yin Li, and Sing Bing Kang. Symmetric stereo matching for occlusion handling. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, 2005.
- [37] Ralph Trapp, Siegbert Drüe, and Georg Hartmann. Stereo matching with implicit detection of occlusions. In *Proceedings of the European Conference on Computer Vision (ECCV)*, 1998.
- [38] Muninder Veldandi, Soumik Ukil, and Krishna Govinda Rao. Robust segment-based stereo using cost aggregation. In *Proceedings of the British Machine Vision Conference (BMVC)*, 2014.
- [39] Yang Wang, Yi Yang, Zhenheng Yang, Liang Zhao, and Wei Xu. Occlusion aware unsupervised learning of optical flow. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, 2018.
- [40] Daniel Weinland, Mustafa Özuysal, and Pascal Fua. Making action recognition robust to occlusions and viewpoint changes. In *Proceedings of the European Conference on Computer Vision (ECCV)*, 2010.
- [41] Saining Xie and Zhuowen Tu. Holistically-nested edge detection. *International Journal of Computer Vision (IJCV)*, 125(1-3):3–18, 2017.
- [42] Dan Xu, Elisa Ricci, Wanli Ouyang, Xiaogang Wang, and Nicu Sebe. Multi-scale continuous crfs as sequential deep networks for monocular depth estimation. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, 2017.
- [43] Qingxiong Yang, Liang Wang, Ruigang Yang, Henrik Stewénus, and David Nistér. Stereo matching with color-weighted correlation, hierarchical belief propagation, and occlusion handling. *IEEE Transactions on Pattern Analysis and Machine Intelligence (TPAMI)*, 31(3):492–504, 2009.

- [44] Xiaoqing Ye, Yuzhang Gu, Lili Chen, Jiamao Li, Han Wang, and Xiaolin Zhang. Order-based disparity refinement including occlusion handling for stereo matching. *IEEE Signal Processing Letters*, 24(10):1483–1487, 2017.
- [45] Jure Zbontar and Yann LeCun. Stereo matching by training a convolutional neural network to compare image patches. *Journal of Machine Learning Research (JMLR)*, 17:65:1–65:32, 2016.
- [46] Ke Zhang, Jiangbo Lu, and Gauthier Laffruit. Cross-based local stereo matching using orthogonal integral images. *IEEE Transactions Circuits and Systems for Video Technology (TCSVT)*, 19(7):1073–1079, 2009.
- [47] Tinghui Zhou, Matthew Brown, Noah Snavely, and David G. Lowe. Unsupervised learning of depth and ego-motion from video. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, 2017.