

# Learning for Disparity Estimation through Feature Constancy

Zhengfa Liang\* Yiliu Feng\* Yulan Guo Hengzhu Liu Wei Chen  
 Linbo Qiao Li Zhou Jianfeng Zhang

National University of Defense Technology, Changsha, China

{liangzhengfa10, fengyiliu11, yulan.guo, hengzhuliu, chenwei,  
 qiao.linbo, zhouli06, jianfengzhang}@nudt.edu.cn

## Abstract

*Stereo matching algorithms usually consist of four steps, including matching cost calculation, matching cost aggregation, disparity calculation, and disparity refinement. Existing CNN-based methods only adopt CNN to solve parts of the four steps, or use different networks to deal with different steps, making them difficult to obtain the overall optimal solution. In this paper, we propose a network architecture to incorporate all steps of stereo matching. The network consists of three parts. The first part calculates the multi-scale shared features. The second part performs matching cost calculation, matching cost aggregation and disparity calculation to estimate the initial disparity using shared features. The initial disparity and the shared features are used to calculate the feature constancy that measures correctness of the correspondence between two input images. The initial disparity and the feature constancy are then fed into a sub-network to refine the initial disparity. The proposed method has been evaluated on the Scene Flow and KITTI datasets. It achieves the state-of-the-art performance on the KITTI 2012 and KITTI 2015 benchmarks while maintaining a very fast running time. Source code is available at <http://github.com/leonzfa/iResNet>.*

## 1. Introduction

Stereo matching aims to estimate correspondences of all pixels between two rectified images [1, 23, 6]. It is a core problem for many stereo vision tasks and has numerous applications in areas such as autonomous vehicles [29], robotics navigation [25], and augmented reality [33].

Stereo matching has been intensively investigated for several decades, with a popular four-step pipeline being developed. This pipeline includes matching cost calculation, matching cost aggregation, disparity calculation and dis-

parity refinement [23]. The four-step pipeline dominates existing stereo matching algorithms [24, 22, 9, 21], while each of its steps is important to the overall stereo matching performance. Due to the powerful representative capability of deep convolution neural network (CNN) for various vision tasks [14, 31, 15], CNN has been employed to improve stereo matching performance and outperforms traditional methods significantly [28, 32, 12, 17, 16].

Zbontar and LeCun [32] first introduced CNN to calculate the matching cost to measure the similarity of two pixels of two images. This method achieved the best performance on the KITTI 2012 [3], KITTI 2015 [18] and Middlebury [23, 24, 22, 9, 21] stereo datasets at that time. They argued that it is unreliable to consider only the difference of photometry in pixels or hand-crafted image features for matching cost. In contrast, CNN can learn more robust and discriminative features from images, and produces improved stereo matching cost. Following the work [32], several methods were proposed to improve the computational efficiency [16] or matching accuracy [28]. However, these methods still suffer from few limitations. **First**, to calculate the matching cost at all potential disparities, multiple forward passes have to be conducted by the network, resulting in high computational burden. **Second**, the pixels in occluded regions (i.e., only visible in one of the two images) cannot be used to perform training. It is therefore difficult to obtain a reliable disparity estimation in these regions. **Third**, several heuristic post-processing steps are required to refine the disparity. The performance and the generalization ability of these methods are therefore limited, as a number of parameters have to be chosen empirically.

Alternatively, the matching cost calculation, matching cost aggregation and disparity calculation steps can be seamlessly integrated into a CNN to directly estimate the disparity from stereo images [17, 12]. Traditionally, the matching cost aggregation and disparity calculation steps are solved by minimizing an energy function defined upon matching costs. For example, the Semi-Global Matching

\*Equal contribution.

(SGM) method [8] uses dynamic programming to optimize a path-wise form of the energy function in several directions. Both the energy function and its solving process are hand-engineered. Different from the traditional methods, Mayer et al. [17] and Kendall et al. [12] directly stacked several convolution layers upon the matching costs, and trained the whole neural network to minimize the distance between the network output and the groundtruth disparity. These methods can achieve higher accuracy and computational efficiency than the methods that use CNN for matching cost calculation only.

If all steps are integrated into a whole network for joint optimization, better disparity estimation performance can be expected. However, it is non-trivial to integrate the disparity refinement step with the other three steps. Existing methods [4, 19] used additional networks for disparity refinement. Specifically, once the disparity is calculated by CNN, one network or multiple networks are introduced to model the joint space of the inputs (including stereo images and initial disparity) and the output (i.e., refined disparity) to refine the disparity.

To bridge the gap between disparity calculation and disparity refinement, we propose to use feature constancy to identify the correctness of the initial disparity, and then perform disparity refinement using feature constancy. Here, “constancy” is borrowed from the area of optical estimation, where “grey value constancy” and “gradient constancy” are used [2]. “Feature constancy” means the correspondence of two pixels in feature space. Specifically, the feature constancy includes two terms, i.e., feature correlation and reconstruction error. The correlation between features extracted from left and right images is considered as the first feature constancy term, which measures the correspondence at all possible disparities. The reconstruction error in feature space is considered as the second feature constancy term estimated with the knowledge on initial disparity. Then, the disparity refinement task aims to improve the quality of the initial disparity given the feature constancy, this can be implemented by a small sub-network. These will be further explained in Sec. 3.3. Experiments on the Scene Flow and KITTI datasets have showed the effectiveness of our disparity refinement approach. Our method seamlessly integrates the disparity calculation and disparity refinement into one network for joint optimization, and improves the accuracy of the initial disparity by a notable margin. Our method achieves the state-of-the-art performance on the KITTI 2012 and KITTI 2015 benchmarks. It also has a very high running efficiency.

The contributions of this paper can be summarized as follows: 1) We integrate all steps of stereo matching into one network to improve accuracy and efficiency; 2) We perform disparity refinement with a sub-network using the feature constancy; 3) We achieve the state-of-the-art perfor-

mance on the KITTI benchmarks.

## 2. Related works

Over the last few years, CNN has been introduced to solve various problems in stereo matching. Existing CNN-based methods can broadly be divided into the following three categories.

### 2.1. CNN for Matching Cost Learning

In this category, CNN is used to learn the matching cost. Zbontar and LeCun [32] trained a CNN to compute the matching cost between two image patches (e.g.,  $9 \times 9$ ), which is followed by several post-processing steps, including cross-based cost aggregation, semi-global matching, left-right constancy check, sub-pixel enhancement, median filtering and bilateral filtering. This architecture needs multiple forward passes to calculate matching cost at all possible disparities. Therefore, this method is computationally expensive. Luo et al. [16] introduced a product layer to compute the inner product between the two representations of a siamese architecture, and trained the network as multi-class classification over all possible disparities to reduce computational time. Park and Lee [20] introduced a pixel-wise pyramid pooling scheme to enlarge the receptive field during the comparison of two input patches. This method produced more accurate matching cost than [32]. Shaked and Wolf [28] deepened the network for matching cost calculation using a highway network architecture with multi-level weighted residual shortcuts. It was demonstrated that this architecture outperformed several networks, such as the base network from MC-CNN [32], the conventional high-way network [30], ResNets [7], DenseNet [10], and the ResNets of ResNets [34].

### 2.2. CNN for Disparity Regression

In this category, CNN is carefully designed to directly estimate the disparity, which enables end-to-end training. Mayer et al. [17] proposed an encoder-decoder architecture for disparity regression. The matching cost calculation is seamlessly integrated into the encoder part. The disparity is directly regressed in a forward pass. Kendall et al. [12] used 3-D convolutions upon the matching costs to incorporate contextual information and introduced a differentiable “soft argmin” operation to regress the disparity. Both methods can run very fast, with 0.06s and 0.9s consumed on a single Nvidia GTX Titan X GPU, respectively. However, disparity refinement is not included in these networks, which limits their performance.

### 2.3. Multiple Sub-networks

In this category, different networks are used to handle the four steps for stereo matching. Shaked and Wolf

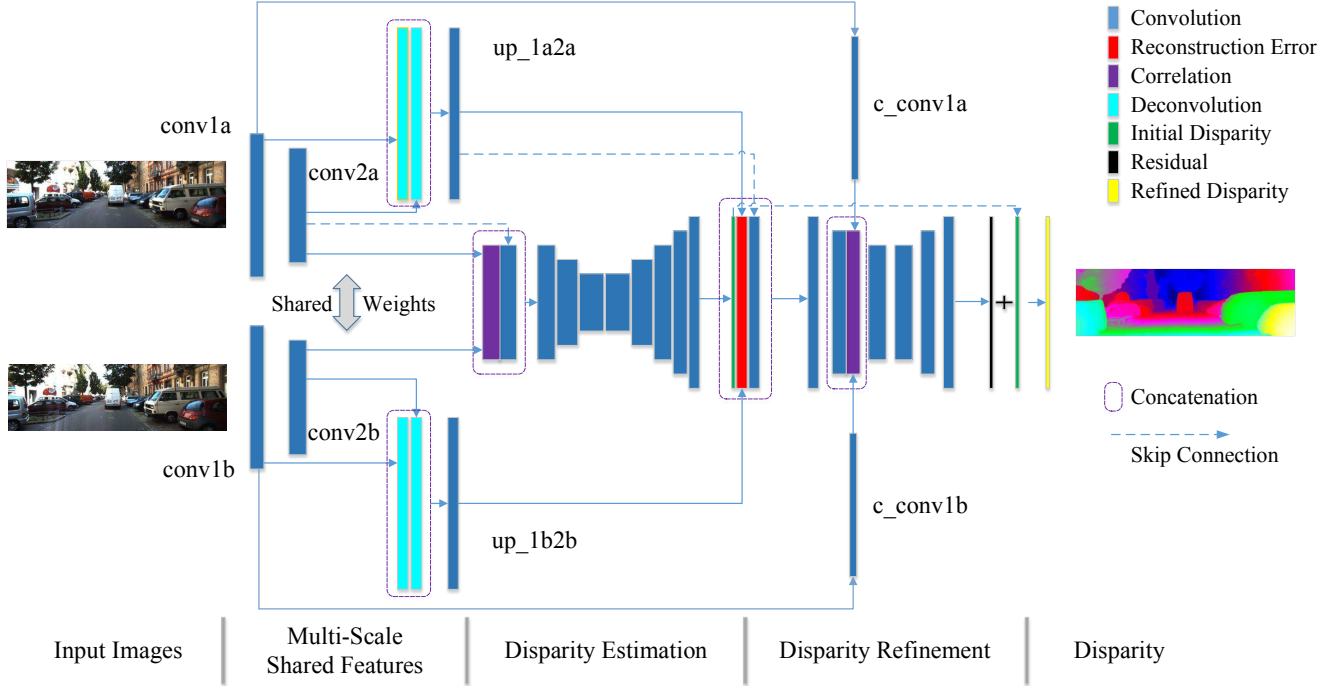


Figure 1. The architecture of our proposed network. It incorporates all of the four steps for stereo matching into a single network. Note that, the skip connections between encoder and decoder at different scales are omitted here for better visualization.

[28] used their highway network for matching cost calculation and an additional global disparity CNN to replace the “winner-takes-all” strategy used in conventional matching cost aggregation and disparity calculation steps. This method improves performance in several challenging situations, such as in occluded, distorted, highly reflective and sparse textured regions. Gidaris et al. [4] used the method in [16] to calculate the initial disparity, and then applied three additional neural networks for disparity refinement. Seki and Pollefeys [27] proposed SGM-Nets to learn the SGM parametrization. They obtained better penalties than the hand tuned method used in MC-CNN [32]. Peng et. al [19] built their work upon [17] by cascading an additional network for disparity refinement.

In our work, we incorporate all steps into one network. As a result, all steps can share the same features and can be optimized jointly. Besides, we introduce feature constancy into our network for improved disparity refinement using both feature correlation and reconstruction error. It is clearly demonstrated that better disparity estimation performance can be achieved by our method.

### 3. Approach

Different from existing methods that use multiple networks for different steps in stereo matching, we incorporate all steps into a single network to enable end-to-end train-

ing. The proposed network consists of three parts: multi-scale shared feature extraction, initial disparity estimation and disparity refinement. The framework of the proposed network is shown in Figure 1, and the network architecture is described in Table 1.

#### 3.1. Stem Block for Multi-scale Feature Extraction

The stem block extracts multi-scale shared features from the two input images for both initial disparity estimation and disparity refinement sub-networks. It contains two convolution layers with stride of 2 to reduce the resolution of inputs, and two deconvolution layers to up-sample the outputs of the two convolution layers to full-resolution. The up-sampled features are fused through an additional  $1 \times 1$  convolution layer. An illustration is shown in Figure 1. The outputs of this stem block can be divided into three types:

- 1) The outputs of the second convolution layer (i.e., *conv2a* for the left image and *conv2b* for the right image). Correlation with a large displacement (i.e., 40) is performed between *conv2a* and *conv2b* to capture the long-range but coarse-grained correspondence between two images. It is used by the first sub-network for initial disparity estimation.
- 2) The outputs of the first convolution layer (i.e., *conv1a* and *conv1b*). They are first compressed to fewer channels to obtain *c\_conv1a* and *c\_conv1b* through a con-

volution layer with a kernel size of  $3 \times 3$ , on which correlation with a small displacement (i.e., 20) is performed to capture short-range but fine-grained correspondence, which is complementary to the former correlation. Besides, these features also act as the **first feature constancy term** used by the second sub-network.

- 3) Multi-scale fusion features (i.e.,  $up\_1a2a$  and  $up\_1b2b$ ). They are first used as skip connection features to bring detailed information for the first sub-network. They are then used to calculate the **second feature constancy term** for the second sub-network.

### 3.2. Initial Disparity Estimation Sub-network

This sub-network generates a disparity map from “ $conv2a$ ” and “ $conv2b$ ” through an encoder-decoder architecture, which is inspired by DispNetCorr1D [17]. DispNetCorr1D can only produce disparity of half resolution. Using the full-resolution multi-scale fusion features as skip connection features, we are able to estimate initial disparity of full resolution. The multi-scale fusion features are also used to calculate the reconstruction error, as will be described in Sec. 3.3. In this sub-network, a correlation layer [17] is first introduced to calculate the matching costs in feature space (i.e., inner product between two feature maps at every possible disparity). There is a trade-off between accuracy and computational efficiency for matching cost calculation. That is, if matching cost is calculated using high-level features, more details are lost and several similar correspondences cannot be distinguished. In contrast, if matching cost is calculated using low-level features, the computational cost is high as feature maps are too large, and the receptive field is too small to capture robust features.

The matching cost is then concatenated with features from the left image. By concatenation, we expect the sub-network to consider low-level semantic information when performing disparity estimation over the matching costs. To some extent, it helps aggregate the matching cost and improves disparity estimation.

Disparity estimation is performed in the decoder part at different scales, where skip connection is introduced at each scale, as illustrated in Table 1. For the sake of computational efficiency, the multi-scale fusion features (described in Sec. 3.1) are only skip connected to the last layer of the decoder to perform full-resolution disparity estimation. This sub-network is called **DES-net**.

### 3.3. Disparity Refinement Sub-network

Although the disparity map estimated in Sec. 3.2 is already good, it still suffers from several challenges such as depth discontinuities and outliers. Consequently, disparity refinement is required to further improve the depth estimation performance.

Table 1. The detailed architecture of our network. Note that, in the “Input” column, “+” means concatenation into one bottom blob, while “,” means the inputs are fed into different bottom blobs, in which case the input channel number means the channel number of each bottom blob.

Type	Name	k	s	c I/O	Input
<i>Stem Block for Multi-scale Shared Features Extraction</i>					
Conv	conv1a	7	2	3/64	left image
	conv1b				right image
Deconv	up_1a	4	2	64/32	conv1a
	up_1b				conv1b
Conv	conv2a	5	2	64/128	conv1a
	conv2b				conv1b
Deconv	up_2a	8	4	128/32	conv2a
	up_2b				conv2b
Conv	up_1a2a	1	1	64/32	up_1a+up_2a
	up_1b2b				up_1b+up_2b
<i>Initial Disparity Estimation Sub-network</i>					
Corr	corr1d	1	1	128/81	conv2a, conv2b
Conv	conv_redir	1	1	128/64	conv2a
Conv	conv3	3	2	145/256	corr1d+conv_redir
Conv	conv3.1	3	1	256/256	conv3
Conv	conv4	3	2	256/512	conv3.1
Conv	conv4.1	3	1	512/512	conv4
Conv	conv5	3	2	512/512	conv4.1
Conv	conv5.1	3	1	512/512	conv5
Conv	conv6	3	2	512/1024	conv5.1
Conv	conv6.1	3	1	1024/1024	conv6
Conv	disp6	3	1	1024/1	conv6.1
Deconv	uconv5	4	2	1024/512	conv6.1
Conv	iconv5	3	1	1025/512	uconv5+disp6+conv5.1
Conv	disp5	3	1	512/1	iconv5
Deconv	uconv4	4	2	512/256	iconv5
Conv	iconv4	3	1	769/256	uconv4+disp5+conv4.1
Conv	disp4	3	1	256/1	iconv4
Deconv	uconv3	4	2	256/128	iconv4
Conv	iconv3	3	1	385/128	uconv3+disp4+conv3.1
Conv	disp3	3	1	128/1	iconv3
Deconv	uconv2	4	2	128/64	iconv3
Conv	iconv2	3	1	193/64	uconv2+disp3+conv2a
Conv	disp2	3	1	64/1	iconv2
Deconv	uconv1	4	2	64/32	iconv2
Conv	iconv1	3	1	97/32	uconv1+disp2+conv1a
Conv	disp1	3	1	32/1	iconv1
Deconv	uconv0	4	2	32/32	iconv1
Conv	iconv0	3	1	65/32	uconv0+disp1+up_1a2a
Conv	disp0	3	1	32/1	iconv0
<i>Disparity Refinement Sub-network</i>					
Warp	w_up_1b2b	-	-	32/32	up_1b2b, disp0
Conv	r_conv0	3	1	65/32	[up_1a2a-w_up_1b2b] +disp0+up_1a2a
Conv	r_conv1	3	2	32/64	r_conv0
Conv	c_conv1a	3	1	64/16	conv1a
	c_conv1b				conv1b
Corr	r_corr	1	1	16/41	c_conv1a, c_conv1b
Conv	r_conv1.1	3	1	105/64	r_conv1+r_corr
Conv	r_conv2	3	2	64/128	r_conv1.1
Conv	r_conv2.1	3	1	128/128	r_conv2
Conv	r_res2	3	1	128/1	r_conv2.1
Deconv	r_uconv1	4	2	128/64	r_conv2.1
Conv	r_iconv1	3	1	129/64	r_uconv1+r_res2+r_conv1.1
Conv	r_res1	3	1	64/1	r_iconv1
Deconv	r_uconv0	4	2	64/32	r_iconv1
Conv	r_iconv0	3	1	65/32	r_uconv1+r_res1+r_conv0
Conv	r_res0	3	1	32/1	r_iconv0

In this paper, we perform disparity refinement using feature constancy. Specifically, once the initial disparity  $disp_i$  is obtained using DES-net, we calculate the two feature constancy terms (i.e., feature correlation  $fc$  and reconstruction error  $re$ ). Then, the task of disparity refinement is to obtain the refined disparity  $disp_r$  considering these three types of



information, i.e.,

$$\{disp_i, fc, re\} \xrightarrow{CNN} disp_r \quad (1)$$

Specifically, the first feature constancy term  $fc$  is calculated as the correlation between feature maps of the left and right images (i.e.,  $c\_conv1a$  and  $c\_conv1b$ ).  $fc$  measures the correspondence of two feature maps at all displacements in disparity range that considered. It produces large values at correct disparities. The second feature constancy term  $re$  is calculated as the reconstruction error of the initial disparity, i.e., the absolute difference between the multi-scale fusion features (Sec. 3.1) of the left image and the back-warped features of the right image. Note that, to calculate  $re$ , only one displacement is conducted at each location in feature maps, which relies on the corresponding value of the initial disparity. If the reconstruction error is large, the estimated disparity is more likely to be incorrect or from occluded regions.

In practice, given the initial disparity produced by the disparity estimation sub-network (Sec. 3.2), the disparity refinement sub-network estimates the residual to the initial disparity. The summation of the residual and the initial disparity is considered as the refined disparity map. Since both the initial disparity and the two feature constancy terms are used to produce the disparity map  $disp_r$  (as shown in Eq. (1)), the disparity estimation performance is expected to be improved. This sub-network is called **DRS-net**. Note that, since the four steps for stereo matching are integrated into a single CNN network, end-to-end training is ensured.

### 3.4. Iterative Refinement

To extract more information from the multi-scale fusion features and to ultimately improve the disparity estimation accuracy, an iterative refinement approach is proposed. Specifically, the refined disparity map produced by the second sub-network (Sec. 3.3) is considered as a new initial disparity map, the feature constancy calculation and disparity refinement processes are then repeated to obtain a new refined disparity. This procedure is repeated until the improvement between two consecutive iterations is small. Note that, as the number of iterations is increased, the improvement decreases.

## 4. Experiments

In this section, we evaluate our method iResNet (iterative residual prediction network) on two datasets, i.e., Scene Flow [17] and KITTI [3, 18]. The Scene Flow dataset [17] is a synthesised dataset containing 35, 454 training image pairs and 4, 370 test image pairs. Dense groundtruth disparities are provided for both training and test sets. Besides, this dataset is sufficiently large to train a model without over-fitting. Therefore, the Scene Flow dataset [17]

is used to investigate different aspects of our method in Sec. 4.1. The KITTI dataset includes two subsets, i.e., KITTI 2012 and KITTI 2015. The KITTI 2012 dataset consists of 194 training image pairs and 195 test image pairs, while the KITTI 2015 dataset consists of 200 training image pairs and 200 test image pairs. These images were recorded in real scenes under different weather conditions. Our method is further compared to the state-of-the-art methods on the KITTI dataset (Sec. 4.2), with the best results being achieved.

Our method was implemented in Caffe [11]. All models were optimized using the Adam method [13] with  $\beta_1 = 0.9$ ,  $\beta_2 = 0.999$ , and a batch size of 2. “Multi-step” learning rate was used for the training. Specifically, for training on the Scene Flow dataset, the learning rate was initially set to  $10^{-4}$  and then reduced by a half at the 200k-th, 350k-th and 500k-th iterations, the training was stopped at the 650k-th iteration. This training procedure was repeated for an additional round to further optimize the model. For fine-tuning on the KITTI dataset, the learning rate was set to  $2 \times 10^{-5}$  for the first 20k iterations and then reduced to  $10^{-5}$  for the subsequent 120k iterations.

Data augmentation was also conducted for training, including spatial and chromatic transformations. The spatial transformations include translation, cropping and scaling, while the chromatic transformations includes color, contrast and brightness transformations. This data augmentation can help to learn a robust model against illumination changes and noise.

### 4.1. Ablation Experiments

In this section, we present several ablation experiments on the Scene Flow dataset to justify our design choices. For evaluation, we use the end-point-error (EPE), which is calculated as the average absolute distance between estimated and groundtruth disparity. We also use the percentage of disparities with their EPE larger than  $t$  pixels ( $> t$  px).

#### 4.1.1 Multi-scale Skip Connection

In Sec. 3, multi-scale skip connection is used to introduce features from different levels to improve disparity estimation and refinement performance. To demonstrate its effectiveness, the multi-scale skip connection scheme of our network was replaced by a single-scale skip connection scheme, the comparative results are shown in Table 3. It can be observed that the multi-scale skip connection scheme outperforms its single-scale counterpart, with the EPE being reduced from 2.55 to 2.50. That is because, the output of the first convolution layer contains high frequency information, it produces high reconstruction error for both regions along object boundaries and regions with large color changes. Note that, regions on an object surface

far from boundaries usually have a very accurate initial disparity estimation (i.e., the true reconstruction error is small), although large color changes occur due to texture variation. Therefore, the reconstruction errors given by the first convolution layer for these regions are inaccurate. In this case, multi-scale skip connection is able to improve the reliability of resulted reconstruction errors. Besides, introducing high-level features is also useful for feature constancy calculation, as higher-level features leverage more context information with a wide field of view.

#### 4.1.2 Feature Constancy for Disparity Refinement

To seamlessly integrate the initial disparity estimation sub-network (DES-net) and the disparity refinement sub-network (DRS-net) into a whole network, the feature constancy calculation and its subsequent sub-network play an important role. To demonstrate its effectiveness, we first removed all feature constancy used in our network (as shown in Table 1) and then retrained the model. The results are shown in Table 2. It can be observed that if no feature constancy is introduced for disparity refinement, the performance improvement is very small, with EPE being reduced from 2.81 to 2.72.

Then, we evaluate the importance of the three information in Eq. (1), i.e., initial disparity  $disp_i$ , feature correlation  $fc$ , and reconstruction error  $re$  produced by the initial disparity, as explained in Sec. 3. The results are shown in Table 2. It is observed that, the reconstruction error  $re$  plays the major role for the performance improvement. If the reconstruction error is removed, EPE is increased from 2.50 to 2.70. That is because this term provides the some knowledge about the initial disparity. That is, regions with poor initial disparity can be identified and then be contrapuntally refined. Besides, removing initial disparity  $disp_i$  or feature correlation  $fc$  from the disparity refinement sub-network slightly degrades the overall performance. Their EPE values are increased from 2.50 to 2.56 and 2.61, respectively. If all the three parts are incorporated, the disparity refinement network can achieve the best performance.

#### 4.1.3 Iterative Refinement

Iterative feature constancy calculation can further improve the overall performance. In practice, we do not train multiple DRS-nets. Instead, we directly stack another DRS-net, whose weights are exactly the same as the original DRS-net, at the top of the whole network. As the number of iterations is increased, the performance improvement is decreased rapidly. Specifically, the first iteration can reduce EPE from 2.50 to 2.46, while the second iteration can only reduce EPE from 2.46 to 2.45. Moreover, there is no obvious performance improvement after the third iteration of refinement. It can be concluded that: 1) Our framework can



Figure 2. Disparity refinement results on the Scene Flow test set under different iterations. The first row represents the input images, the second row shows the initial disparity without any refinement, the third and fourth rows show the refined disparity after 1 and 2 iterations, respectively. The last row gives the groundtruth disparity.

efficiently extract useful information for disparity estimation using feature constancy information; 2) The information contained in feature space is still limited. Therefore, it is possible to improve the disparity refinement performance by introducing more powerful features.

To further demonstrate the effectiveness of iterative refinement, the disparity estimation results for 2 iterations are shown in Figure 2. It can be observed that, lots of details cannot be accurately estimated in the initial disparity, e.g., the areas shown in red boxes. However, after two iterations of refinement, most inaccurate estimations can be corrected, and the refined disparity map looks more smooth.

Table 2. Comparative results on the Scene Flow dataset for networks with different settings on the disparity refinement sub-network. DES-net and DRS-net represent the initial disparity estimation sub-network and the disparity refinement sub-network, respectively.

Model	> 1px	> 3px	> 5px	EPE	Params.	Time (ms)
DES-net (without disparity refinement)	16.78	6.49	4.31	2.81	<b>42.76M</b>	<b>61</b>
DES-net + DRS-net without feature correlation $fc$ and $re$	15.65	6.12	4.11	2.72	43.30M	100
DES-net + DRS-net without reconstruction error $re$	15.54	6.10	4.10	2.70	43.34M	111
DES-net + DRS-net without feature correlation $fc$	11.13	5.32	3.79	2.56	43.31M	103
DES-net + DRS-net without initial disparity $disp_i$	11.64	5.37	3.81	2.61	43.34M	114
DES-net + DRS-net (iResNet)	10.24	4.93	3.54	2.50	43.34M	114
Refinement $\times 2$ (iResNet-i2)	9.42	4.64	3.37	2.46	43.34M	131
Refinement $\times 3$ (iResNet-i3)	<b>9.28</b>	<b>4.57</b>	<b>3.32</b>	<b>2.45</b>	43.34M	148

Table 3. Comparative results on the Scene Flow dataset for networks with single-scale and multi-scale skip connection.

Model	> 1px	> 3px	> 5px	EPE
Single-scale	10.90	5.23	3.74	2.55
Multi-scale	<b>10.24</b>	<b>4.93</b>	<b>3.54</b>	<b>2.50</b>

Table 4. EPE results on Scene Flow dataset achieved by the proposed iResNet method and the CRL method.

Method	EPE	Params.	Run time(ms)
CRL [19]	1.60	78.77M	162
iResNet	<b>1.40</b>	<b>43.11M</b>	<b>90</b>

#### 4.1.4 Feature Constancy vs Color Constancy

In this experiment, the superiority of feature constancy for disparity refinement over color constancy is demonstrated. Our method was compared to the cascade residual learning (CRL) method [19], which calculated the reconstruction error in the color space. Intuitively, calculating the reconstruction error in the feature space could obtain more robust results, since the learned features are more robust to noise and luminance changes. Besides, by sharing the features with the first network, the second network can be designed shadower, which subsequently leads to a more compact network. CRL [19] used one network (i.e., DispNetC [17]) for disparity calculation and an additional network (i.e., DispResNet [19]) for disparity refinement. For fair comparison, we also use DispNetC [17] without any fine-tuning as our disparity estimation sub-network. Note that, in their experiments on the Scene Flow dataset, disparity images (and their corresponding stereo pairs) with more than 25% of their disparity values greater than 300 were removed. We followed the same protocol in this comparison. Comparative results are shown in Table 4. The EPE result of CRL was taken from the paper [19], and its run time was tested on Nvidia Titan X (Pascal) using our implementation. It can be seen that our method (using feature constancy) significantly outperforms CRL (using color constancy). The EPE values achieved by our iResNet method and the CRL method are 1.40 and 1.60, respectively. Moreover, our method also requires fewer parameters and costs less computational time.

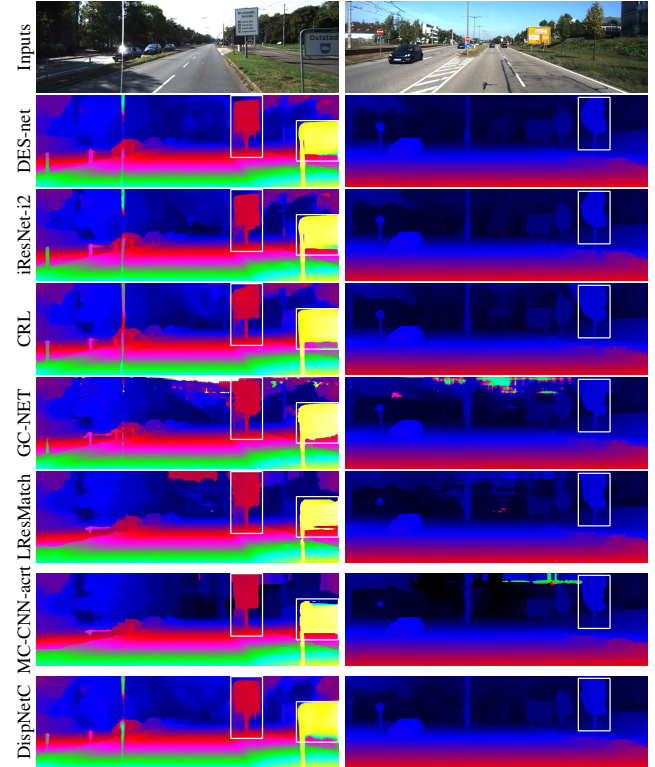


Figure 3. Comparison with other state-of-the-art methods on the KITTI 2015 dataset. The images in the first row are input images from KITTI 2015. Our iResNet-i2 refinement results (the third row) can greatly improve the initial disparity (the second row), and give better visualization effect than other methods, especially in the upper part of images, where there is no groundtruth in these regions.

#### 4.2. Benchmark Results

We further compared our method to several existing methods on the KITTI 2012 and KITTI 2015 benchmarks, including GC-NET [12], L-ResMatch [28], SGM-Net [27], SsSMnet [35], PBCP [26], Displets [5], and MC-CNN [32].

For the evaluation on KITTI 2012, we used the percentage of erroneous pixels in non-occluded (Out-Noc) and all (Out-All) areas. Here a pixel is considered to be erroneous

Table 5. Results on the KITTI 2012 dataset.

	> 2px		> 3px		> 4px		> 5px		Runtime (s)
	Out-Noc	Out-All	Out-Noc	Out-All	Out-Noc	Out-All	Out-Noc	Out-All	
GC-NET [12]	2.71	3.46	1.77	2.30	1.36	1.77	1.12	1.46	0.9
L-ResMatch [28]	3.64	5.06	2.27	3.40	1.76	2.67	1.50	2.26	48
SGM-Net [27]	3.60	5.15	2.29	3.50	1.83	2.80	1.60	2.36	67
SsSMnet [35]	3.34	4.24	2.30	3.00	1.82	2.39	1.53	2.01	0.8
PBCP [26]	3.62	5.01	2.36	3.45	1.88	2.28	1.62	2.32	68
Displets v2 [5]	3.43	4.46	2.37	3.09	1.97	2.52	1.72	2.17	265
MC-CNN-acrt [32]	3.90	5.45	2.43	3.63	1.90	2.85	1.64	2.39	67
DES-net (ours)	4.88	5.54	2.66	3.12	1.78	2.11	1.33	1.59	<b>0.05</b>
iResNet-i2 (ours)	<b>2.69</b>	<b>3.34</b>	<b>1.71</b>	<b>2.16</b>	<b>1.30</b>	<b>1.63</b>	<b>1.06</b>	<b>1.32</b>	0.12

if its disparity EPE is larger than  $t$  pixels ( $> t$  px). For the evaluation on KITTI 2015, we used the percentage of erroneous pixels in background ( $D1-bg$ ), foreground ( $D1-fg$ ) or all pixels ( $D1-all$ ) in the non-occluded or all regions. Here, a pixel is considered to be correct if its disparity EPE is less than 3 or 5% pixels.

The results are shown in Tables 5 and 6. For our method, the results for both of the DES-net (without disparity refinement) and the final iResNet-i2 model (with disparity refinement of 2 iterations) are presented. It is clear that the disparity refinement sub-network can consistently improve the performance by a notable margin. Moreover, our iResNet-i2 model achieves the best disparity estimation performance on both the KITTI 2012 and KITTI 2015 datasets in different scenarios. Note that, our method is also highly efficient, it achieves the shortest run time on the KITTI 2012 dataset. The overall run time tested on a single Nvidia Titan X (Pascal) GPU is only 0.12s.

Figure 3 illustrates few qualitative results achieved by our method and several state-of-the-art methods on the KITTI 2015 dataset. It can be observed that our method produces more smooth disparity estimation results. On one hand, our disparity refinement sub-network DRS-net can effectively improve the quality of the initial disparity estimated by DES-net, with many details being recovered. On the other hand, our method produces much better results than other methods in the upper parts of these images, as shown in white boxes in Figure 3. In fact, the upper parts of these images correspond to sky or regions beyond the working distance of a lidar. In that case, groundtruth disparity cannot be provided for these regions for training, making the learning based methods to deteriorate. From Figure 3, we can see that the performance of other methods in these regions is relatively poor. However, our method can still achieve acceptable results, with more accurate disparity estimation along boundaries. This also indicates that our method generalizes well on unseen data. To further demonstrate the generalization capability of our method, the iResNet-i2, CRL [19] and DispNetC [17] models are first trained on the KITTI 2015 training set and then tested on

Table 6. Results on the KITTI 2015 dataset.

	All Pixels			Non-Occluded Pixels		
	D1-bg	D1-fg	D1-all	D1-bg	D1-fg	D1-all
CRL [19]	2.48	3.59	2.67	2.32	3.12	2.45
GC-NET [12]	2.21	6.16	2.87	2.02	5.58	2.61
DRR [4]	2.58	6.04	3.16	2.34	4.87	2.76
SsSMnet [35]	2.70	6.92	3.40	2.46	6.13	3.06
L-ResMatch [28]	2.72	6.95	3.42	2.35	5.74	2.91
Displets v2 [5]	3.00	5.56	3.43	2.73	4.95	3.09
SGM-Net [27]	2.66	8.64	3.66	2.23	7.44	3.09
MC-CNN [32]	2.89	8.88	3.88	2.48	7.64	3.33
DispNetC [17]	4.32	4.41	4.34	4.11	3.72	4.05
DES-net (ours)	3.13	3.87	3.25	2.94	3.21	2.98
iResNet-i2 (ours)	<b>2.25</b>	<b>3.40</b>	<b>2.44</b>	<b>2.07</b>	<b>2.76</b>	<b>2.19</b>

Table 7. The generalization performance from KITTI 2015 benchmark to KITTI 2012 benchmark achieved by three methods.

Methods	KITTI 2015	KITTI 2012	Performance Drop
iResNet-i2 (ours)	<b>2.44</b>	<b>3.62</b>	<b>1.18</b>
CRL [19]	2.67	4.82	2.15
DispNetC [17]	4.34	9.64	5.3

the KITTI 2015 and 2012 test sets without fine-tuning. The achieved  $D1-all$  errors are shown in Table 7. It is clear that our method achieves the smallest performance drop.

## 5. Conclusion

In this work, we propose a network architecture to integrate the four steps of stereo matching for joint training. Our network first estimates an initial disparity, and then uses two feature constancy terms to refine the disparity. The refinement is performed using both feature correlation and reconstruction error, which effectively improves the quality of the initial estimated disparity. Moreover, by introducing the multi-scale shared features, our network can be designed more compact. Experimental results show that the proposed method achieves the state-of-the-art disparity estimation performance on the KITTI 2012 and KITTI 2015 benchmarks. Besides, our method is also highly efficient for calculation.



## References

- [1] S. T. Barnard and M. A. Fischler. Computational stereo. *Acm Computing Surveys*, 14(4):553–572, 1982. [1](#)
- [2] T. Brox, A. Bruhn, N. Papenberger, and J. Weickert. High accuracy optical flow estimation based on a theory for warping. In *European Conference on Computer Vision*, volume 3024, pages 25–36, 2004. [2](#)
- [3] A. Geiger. Are we ready for autonomous driving? the kitti vision benchmark suite. In *IEEE Conference on Computer Vision and Pattern Recognition*, pages 3354–3361, 2012. [1](#), [5](#)
- [4] S. Gidaris and N. Komodakis. Detect, replace, refine: Deep structured prediction for pixel wise labeling. In *IEEE Conference on Computer Vision and Pattern Recognition*, 2017. [2](#), [3](#), [8](#)
- [5] F. Guney and A. Geiger. Displets: Resolving stereo ambiguities using object knowledge. In *IEEE Conference on Computer Vision and Pattern Recognition*, pages 4165–4175, 2015. [8](#)
- [6] R. Hartley and A. Zisserman. *Multiple View Geometry in Computer Vision*. Cambridge University Press, 2003. [1](#)
- [7] K. He, X. Zhang, S. Ren, and J. Sun. Deep residual learning for image recognition. In *IEEE Conference on Computer Vision and Pattern Recognition*, pages 770–778, 2016. [2](#)
- [8] H. Hirschmüller. Stereo processing by semiglobal matching and mutual information. *IEEE Transactions on Pattern Analysis & Machine Intelligence*, 30(2):328, 2008. [2](#)
- [9] H. Hirschmüller and D. Scharstein. Evaluation of cost functions for stereo matching. In *IEEE Conference on Computer Vision and Pattern Recognition*, pages 1–8, 2007. [1](#)
- [10] G. Huang, Z. Liu, K. Q. Weinberger, and V. D. M. Laurens. Densely connected convolutional networks. In *IEEE Conference on Computer Vision and Pattern Recognition*, 2017. [2](#)
- [11] Y. Jia, E. Shelhamer, J. Donahue, S. Karayev, J. Long, R. Girshick, S. Guadarrama, and T. Darrell. Caffe: Convolutional architecture for fast feature embedding. *arXiv preprint arXiv:1408.5093*, 2014. [5](#)
- [12] A. Kendall, H. Martirosyan, S. Dasgupta, P. Henry, R. Kennedy, A. Bachrach, and A. Bry. End-to-end learning of geometry and context for deep stereo regression. In *IEEE Conference on Computer Vision and Pattern Recognition*, 2017. [1](#), [2](#), [8](#)
- [13] D. P. Kingma and J. Ba. Adam: A method for stochastic optimization. *arXiv preprint arXiv:1412.6980*, 2014. [5](#)
- [14] A. Krizhevsky, I. Sutskever, and G. E. Hinton. Imagenet classification with deep convolutional neural networks. In *International Conference on Neural Information Processing Systems*, pages 1097–1105, 2012. [1](#)
- [15] J. Long, E. Shelhamer, and T. Darrell. Fully convolutional networks for semantic segmentation. In *Computer Vision and Pattern Recognition*, pages 3431–3440, 2015. [1](#)
- [16] W. Luo, A. G. Schwing, and R. Urtasun. Efficient deep learning for stereo matching. In *IEEE Conference on Computer Vision and Pattern Recognition*, pages 5695–5703, 2016. [1](#), [2](#), [3](#)
- [17] N. Mayer, E. Ilg, P. Hausser, P. Fischer, D. Cremers, A. Dosovitskiy, and T. Brox. A large dataset to train convolutional networks for disparity, optical flow, and scene flow estimation. In *IEEE Conference on Computer Vision and Pattern Recognition*, pages 4040–4048, 2016. [1](#), [2](#), [3](#), [4](#), [5](#), [7](#), [8](#)
- [18] M. Menze and A. Geiger. Object scene flow for autonomous vehicles. In *IEEE Conference on Computer Vision and Pattern Recognition*, pages 3061–3070, 2015. [1](#), [5](#)
- [19] J. Pang, W. Sun, J. S. Ren, C. Yang, and Q. Yan. Cascade residual learning: A two-stage convolutional neural network for stereo matching. In *International Conference on Computer Vision Workshop*, 2017. [2](#), [3](#), [7](#), [8](#)
- [20] H. Park and K. M. Lee. Look wider to match image patches with convolutional neural networks. *IEEE Signal Processing Letters*, PP(99):1–1, 2017. [2](#)
- [21] D. Scharstein, H. Hirschmüller, Y. Kitajima, G. Krathwohl, N. Nei, X. Wang, and P. Westling. High-resolution stereo datasets with subpixel-accurate ground truth. In *German Conference on Pattern Recognition*, pages 31–42, 2014. [1](#)
- [22] D. Scharstein and C. Pal. Learning conditional random fields for stereo. In *IEEE Conference on Computer Vision and Pattern Recognition*, pages 1–8, 2007. [1](#)
- [23] D. Scharstein and R. Szeliski. A taxonomy and evaluation of dense two-frame stereo correspondence algorithms. *International Journal of Computer Vision*, 47:7–42, 2002. [1](#)
- [24] D. Scharstein and R. Szeliski. High-accuracy stereo depth maps using structured light. In *IEEE Computer Society Conference on Computer Vision and Pattern Recognition*, pages 195–202, 2003. [1](#)
- [25] K. Schmid, T. Tomic, F. Ruess, H. Hirschmüller, and M. Suppa. Stereo vision based indoor/outdoor navigation for flying robots. *IEEE/RSJ International Conference on Intelligent Robots and Systems*, pages 3955–3962, 2013. [1](#)
- [26] A. Seki and M. Pollefeys. Patch based confidence prediction for dense disparity map. In *British Machine Vision Conference (BMVC)*, volume 10, 2016. [8](#)
- [27] A. Seki and M. Pollefeys. SGM-Nets: Semi-global matching with neural networks. In *IEEE Conference on Computer Vision and Pattern Recognition Workshops*, 2017. [3](#), [8](#)
- [28] A. Shaked and L. Wolf. Improved stereo matching with constant highway networks and reflective confidence learning. In *IEEE Conference on Computer Vision and Pattern Recognition*, 2017. [1](#), [2](#), [3](#), [8](#)
- [29] S. Sivaraman and M. M. Trivedi. A review of recent developments in vision-based vehicle detection. *IEEE Intelligent Vehicles Symposium*, pages 310–315, 2013. [1](#)
- [30] R. K. Srivastava, K. Greff, and J. Schmidhuber. Highway networks. *CoRR*, abs/1505.00387, 2015. [2](#)
- [31] N. Wang and D. Y. Yeung. Learning a deep compact image representation for visual tracking. *Advances in Neural Information Processing Systems*, pages 809–817, 2013. [1](#)
- [32] J. Zbontar and Y. LeCun. Stereo matching by training a convolutional neural network to compare image patches. *Journal of Machine Learning Research*, 17(1-32):2, 2016. [1](#), [2](#), [3](#), [8](#)
- [33] N. Zenati and N. Zerhouni. Dense stereo matching with application to augmented reality. In *IEEE International*

- Conference on Signal Processing and Communications, pages 1503–1506, 2008. [1](#)
- [34] K. Zhang, M. Sun, X. Han, X. Yuan, L. Guo, and T. Liu. Residual networks of residual networks: Multilevel residual networks. IEEE Transactions on Circuits and Systems for Video Technology, PP(99):1–1, 2016. [2](#)
- [35] Y. Zhong, Y. Dai, and H. Li. Self-supervised learning for stereo matching with self-improving ability. CoRR, abs/1709.00930, 2017. [8](#)