

PointFix: Learning to Fix Domain Bias for Robust Online Stereo Adaptation

Kwonyoung Kim¹ Jungin Park¹ Jiyoung Lee²
Dongbo Min³ Kwanghoon Sohn^{1*}

¹Yonsei University ²NAVER AI Lab ³Ewha Womans University
{kyk12, newrun, khsohn}@yonsei.ac.kr
lee.j@navercorp.com dbmin@ewha.ac.kr

Abstract. Online stereo adaptation tackles the domain shift problem, caused by different environments between synthetic (training) and real (test) datasets, to promptly adapt stereo models in dynamic real-world applications such as autonomous driving. However, previous methods often fail to counteract particular regions related to dynamic objects with more severe environmental changes. To mitigate this issue, we propose to incorporate an auxiliary point-selective network into a meta-learning framework, called *PointFix*, to provide a robust initialization of stereo models for online stereo adaptation. In a nutshell, our auxiliary network learns to fix local variants intensively by effectively back-propagating local information through the meta-gradient for the robust initialization of the baseline model. This network is model-agnostic, so can be used in any kind of architectures in a plug-and-play manner. We conduct extensive experiments to verify the effectiveness of our method under three adaptation settings such as short-, mid-, and long-term sequences. Experimental results show that the proper initialization of the base stereo model by the auxiliary network enables our learning paradigm to achieve state-of-the-art performance at inference.

Keywords: Online adaptation, stereo depth estimation, meta-learning

1 Introduction

Stereo depth estimation to predict 3D geometry for practical real-world applications such as autonomous driving [1] has been developed by handcrafted methods [11, 9, 43, 13] and deep stereo models based on supervised learning [23, 2, 34, 15] that leverage the excellent representation power of deep neural networks. In general, given that the high performance of deep networks is guaranteed when test and training data are derived from a similar underlying distribution [4, 22, 12, 20, 5], they demand a huge amount of annotated training data to reflect a real-world distribution. Acquiring groundtruth disparity maps, but unfortunately, is

* Corresponding author.

This work was supported by the National Research Foundation of Korea(NRF) grant funded by the Korea government(MSIP) (NRF-2021R1A2C2006703) and the Yonsei University Research Fund of 2021 (2021-22-0001).

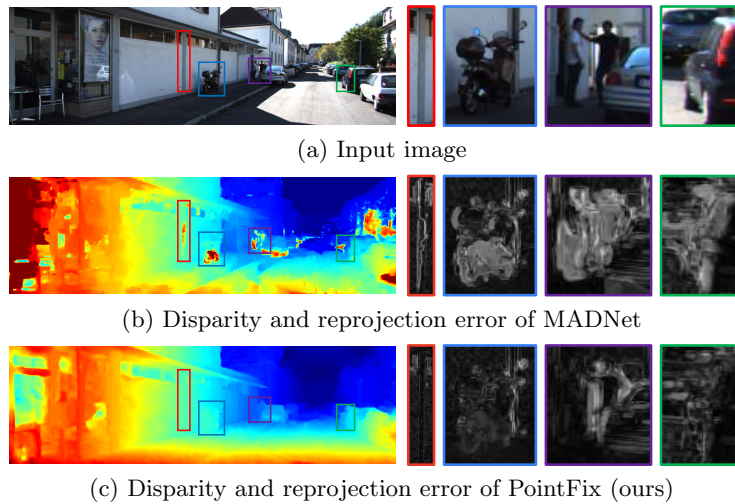


Fig. 1. Estimated disparity after online adaptation from MADNet and our PointFix. Our method has a much stronger adaptation ability, especially in local detail.

laborious and impractical [36]. Especially for autonomous driving, constructing datasets from all possible different conditions (*e.g.* weather and road conditions) is impossible while it is a very fatal problem [37]. To mitigate the aforementioned issues, an intuitive solution is to finetune the stereo model trained on a large-scale synthetic dataset that is easier to collect groundtruth. However, despite the help of large-scale synthetic datasets, most recent works [35, 26, 46, 44] have pointed out the limitation of fine-tuning that is incapable of collecting sufficient data in advance when running the stereo models in the *open world*. While domain generalization methods [42, 32] have shown promising results without real images, they require high computations to provide generalized stereo models and often fail to respond to continuously changing environments.

As an alternative solution, online stereo adaptation [37, 36, 39, 29, 45, 19] is proposed to incorporate unsupervised domain adaptation [6, 10] into a continual learning process [31]. Formally, a baseline network is trained offline using a large number of the labeled synthetic datasets (*e.g.* Synthia [30], FlyingThings3D [23]) and continually adapted to unlabeled unseen scenarios at test time in an unsupervised manner. To demonstrate a faster inference speed for real-world applications, MADNet [37, 29] proposed a lightweight network and modular adaptation framework to rely on self-supervision via reprojection loss [10]. Meanwhile, Learning-to-adapt (L2A) [36] introduced a new learning framework based on model agnostic meta-learning (MAML) [4] for the improved adaptation ability of the network by well-suited base parameters. It shows that the meta-learning framework has great potential in making the network parameters in learning process to make the parameters into a very adaptable state. Despite their great progress, most online adaptation methods [36, 37] have merely attempted to only impose a global average errors from the whole prediction as a

learning objective during an offline training without attention to a domain gap in local, showing poor initial performance.

In particular, we observe that given stereo images from a novel environment, incorrectly estimated disparities are concentrated on specific *local* regions, as depicted in Fig. 1(b). The domain shift [27] issue arises because the local context of test data (*e.g.* appearance deformations of objects, occlusion type, or the form of a shadow etc.) is significantly diverse from those deployed throughout the training process. This means that without taking such locally varying discrepancies between training and test data into account, the global adaptation strategy used in the existing methods [37, 36, 39, 29, 45, 19] has fundamental limitations in improving the adaptation performance. In addition, a plug-and-play algorithm to easily combine with evolving deep stereo networks is also needed.

In this paper, we propose a novel model-agnostic training method for robust online stereo adaptation, called *PointFix*, that can be flexibly built on the top of existing stereo models and learns the base stereo network on a meta-learning framework. Unlike the existing methods [29, 19, 3, 39] that focus on a new online adaptation strategy, we leverage the meta-learning strategy for *learning-to-fix* a base stereo network offline so that it can have generalized initial model parameters and respond to novel environments more robustly. Specifically, we incorporate an auxiliary point-selective network, termed PointFixNet into meta-learning to rectify the local detriment of the base network and alternately fix the base network by an additional update as in the online meta-learning methods [20, 5]. As a result, the parameters of the base network are updated to grasp and utilize the incoming local context and can be robust to the local variants at test time by preventing the network from being biased to global domain dependencies only.

In the experiment, we learn two base stereo models, DispNetC [23], MAD-Net [37], together with our framework on the synthetic data in the offline training, and then update the whole models (full adaptation) or the sub-module of the models (MAD adaptation) in the online adaptation using the unsupervised reprojection loss [6, 10] on the real-world dataset. Note that the proposed auxiliary network is not used in the online adaptation during inference, maintaining an original inference speed of the base stereo models. Given that our PointFix is a general and synergistic strategy that can be adopted with any kinds of stereo networks, it improves a generalization capability of the base stereo model to novel environments through the robust parameter initialization. Extensive experimental results show that PointFix outperforms recent state-of-the-art results by a significant margin on various adaptation scenarios including short-, mid-, and long-term adaptation. In addition, comparison with domain generalization methods [42, 32, 21] demonstrates the superiority of our PointFix in terms of both accuracy and speed.

2 Related Work

For depth estimation from stereo images, there is an extensive literature, but here we briefly introduce related work in the application of convolutional neural networks (CNNs). Modern approaches using CNNs are mostly categorized by

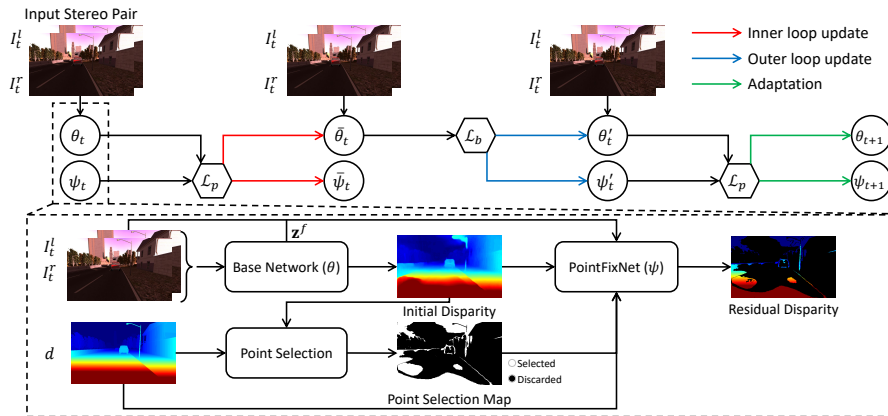


Fig. 2. The overall framework of the proposed *PointFix* and *PointFixNet*. We alternatively update the base network and *PointFixNet* underlying meta-learning. The detail of main flowchart is illustrated in the dotted box.

matching-based approaches [41, 15, 33] that learn how to match corresponding points, and regression-based methods [14, 2, 23] that learn to directly regress sub-pixel disparities. To further enhance the performance, some works [17, 16, 28] consider exploiting auxiliary network or module to assist base network by estimating confidence of prediction map from stereo inputs, prediction or cost volume of the base network. Although their results are promising, they have a limitation in retaining the superb accuracy in new domains [29].

Pointing out the domain shift issue, recent works [36, 37, 29, 45, 39, 3, 19] have proposed online stereo adaptation methods to consider more practical solution for real-world applications. They have argued that we should consider a new *open-world* scenario in which input frames are sequentially provided to the model with certain time intervals. By continually updating parameters at test time in an unsupervised manner, they observe the adaptability of the models in changing environments. As one of the pioneering works, [37, 29] has proposed a light weight model (*MADNet*) and its modular updating frameworks (*MAD*, *MAD++*) to improve the adaptation speed drastically. Following this idea, [39] has picked up the speed even more by implementing ‘Adapt or Hold’ mechanism based on deep Q-learning network [25].

Closely related work to ours is meta-learning based online adaptation approaches [36, 45] that train to learn proper model parameters to be better suitable for online adaptation. *L2A* [36] has directly incorporated online adaptation process into the inner loop updating process, and also learn a confidence measure to use adaptive weighted loss. However, *L2A* is inherently unstable during training due to multiple adaptation steps, especially when coupled with the lightweight model (*e.g.* *MADNet* [29]). Also, it is worth noting a difference in purpose between confidence-weighted adaptation of *L2A* and our *PointFixNet*. The role of the auxiliary network in *L2A* is to eliminate uncertain errors from inherently noisy reprojection loss, whereas ours is to stabilize the generality of the base model by generating a proper point-wise learning objective for particu-

lar bad pixels. We hereby deliver the gradient of the point-wise loss to the base stereo model to prevent the model from being learned by minimizing only the global errors and to remedy domain-invariant representations in local regions.

3 Problem Statement and Preliminaries

Online Stereo Adaptation Here we first formulate online stereo adaptation. Given stereo image pairs from source domain \mathcal{D}_s with available ground-truth maps, online stereo adaptation aims to learn a stereo model capable of adapting itself dynamically in a novel unseen domain \mathcal{D}_u . In the inference, given a set of parameters θ from the base stereo model trained on \mathcal{D}_s , the base network parameters are updated in a single iteration step t to adapt the stereo model with respect to continuous input sequence without ground-truth disparity maps:

$$\theta_{t+1} \leftarrow \theta_t - \alpha \nabla_{\theta_t} \mathcal{L}_u(\theta_t, I_t^l, I_t^r), \quad (1)$$

where \mathcal{L}_u is an unsupervised loss function, α is the learning rate, I_t^l , and I_t^r are left and right images of the t -th stereo pair from \mathcal{D}_u . We evaluate the performance of the online adaptation under the short-term (sequence-level), mid-term (environment-level), and long-term (full) settings. Note that following previous approaches [36, 37], we use a reprojection error [10] as the unsupervised learning objective \mathcal{L}_u .

Model Agnostic Meta-learning MAML [4] has proposed to learn initial base parameters that are suited to adapt to new domain with only few updates. This is attained by implementing a nested optimization which consists of an inner loop and an outer loop. Specifically, the inner loop updates the base parameters for each sample in batch separately in the standard gradient descent way. The outer loop performs the update of the base parameters using the sum of sample-specific gradients (meta-gradient) which are computed by the parameters updated in the inner loop. Formally, for a set of tasks \mathcal{T} , let the meta-training and meta-testing sets be $\mathcal{D}_\tau^{\text{train}}$ and $\mathcal{D}_\tau^{\text{test}}$ respectively. A set of parameters θ^* can be obtained for a specific task $\tau \in \mathcal{T}$ with a single gradient step in the inner loop:

$$\theta^* = \min_{\theta} \sum_{\tau \in \mathcal{T}} \mathcal{L}(\theta - \alpha \nabla_{\theta} \mathcal{L}(\theta, \mathcal{D}_\tau^{\text{train}}), \mathcal{D}_\tau^{\text{test}}), \quad (2)$$

where \mathcal{L} is a objective function for the task and α is the learning rate. They carries out meta-update, including inner and outer updates, during training and then adaptation for the target task after the the whole meta-learning processes are completely over. In this paper, we treat the disparity prediction for each stereo pair as a single task.

Recent works [20, 5] have introduced two stage meta-learning procedure (a.k.a online meta-learning) that implements meta-update and adaptation stages alternatively during training to achieve more efficient gradient path and better performance. Inspired by their alternative updating scheme, we propose a novel *learning-to-fix* strategy. Our objective is to get base network and PointFixNet parameters that enable the base network to quickly adapt regardless of local variants in unseen domains. The details are described in the following section.

Algorithm 1 Parameter update with PointFix loss

Input: $\mathcal{I} = \{(I_n^l, I_n^r, d_n)\}_{n=0}^{N-1}$; learning rate α ; base model parameters θ ; PointFix network parameters ψ **Output:** updated parameters $\bar{\theta}, \bar{\psi}$

```

1: function POINTUPDATE( $\mathcal{I}, \alpha, \theta, \psi$ )
2:    $\mathcal{L}_p^\tau \leftarrow 0$  ▷ Initialize loss
3:   for  $n = 0$  to  $N - 1$  do
4:      $\hat{d}, \mathbf{z}^b = \mathcal{F}(I_n^l, I_n^r | \theta)$ 
5:      $\mathbf{p}(\theta) = \{(i, j) \mid |(\hat{d}_{ij} - d_{n,ij})|_1 > 3\}$  ▷ Select points
6:      $\mathbf{z}^c = \mathcal{F}(I_n^l, \hat{d}, d_n | \psi^c)$ 
7:     for  $(i, j) \in \mathbf{p}(\theta)$  do
8:        $r_{ij} = \mathcal{F}(\mathbf{z}_{ij}^b, \mathbf{z}_{ij}^c | \psi^p)$  ▷ Residual disparity
9:        $\mathcal{L}_p^\tau \leftarrow \mathcal{L}_p^\tau + \mathcal{L}_p(r_{ij} + \hat{d}_{ij}, d_{n,ij})$  ▷ PointFix loss
10:    end for
11:  end for
12:   $\bar{\theta} \leftarrow \theta - \alpha \nabla_\theta \mathcal{L}_p^\tau$  ▷ Update base network
13:   $\bar{\psi} \leftarrow \psi - \alpha \nabla_\psi \mathcal{L}_p^\tau$  ▷ Update PointFixNet
14:  return  $\bar{\theta}, \bar{\psi}$ 
15: end function

```

4 PointFix: Learning to Fix

We design a novel meta-learning framework for online stereo adaptation to learn good base model and quickly adapt to novel environments (*i.e.*, unseen domain), especially concentrating on erroneous pixels. As illustrated in Fig. 2, we leverage off-the-shelf deep stereo model as a base network and incorporate an auxiliary network, PointFixNet, to make parameters of the base network robust to the local distortion.

4.1 Base Stereo Models

Our goal is to train base models offline to be more suitable for the online adaptation by correcting bias to the seen domain. We have employed two stereo networks as a base model: 1) DispNet-Corr1D [23] (shortened as DispNetC) and 2) MADNet [37]. Besides taking the initial disparity \hat{d} estimated from the base model, we extract intermediate features that is useful to exploit the fine-grained information. The intermediate features consist of the matching cost \mathbf{c} (same as correlation layer in DispnetC) and its corresponding left feature \mathbf{f}^l . They are concatenated and taken as a base feature $\mathbf{z}^b = \Pi(\mathbf{c}, \mathbf{f}^l)$, where $\Pi(\cdot, \cdot)$ is a concatenation operation. We note that in MADNet, the matching cost calculation is similar with the one in DispNetC, but before calculation their right features are warped with a disparity map on a coarse resolution to reduce search range and computation. However, to extract base features, we apply the same matching cost computation scheme used in DispNetC regardless of the base network to ensure the generality of our method.

In each inner loop of our method, we select a pixel i, j to fix local deformations by computing ℓ_1 loss between \hat{d} and the ground-truth d , such that the set of

Algorithm 2 Overall training procedure

Hyperparameters: batch size N ; max iteration K ; learning rate of inner and outer loop α, β
Input: pre-trained base model parameters θ ; source training dataset \mathcal{S}
Output: optimized base model parameters θ^*

```

1: function TRAINING( $\theta, \mathcal{S}$ )
2:   Initialize  $\theta$  and  $\psi$ .
3:   for  $k = 0$  to  $K - 1$  do
4:      $\mathcal{L}_k \leftarrow 0$  ▷ Initialize loss
5:      $\mathcal{I}_k \sim \mathcal{S}$  ▷ Sample a batch of size  $N$ 
6:     for  $n = 0$  to  $N - 1$  do
7:        $\theta_n \leftarrow \theta, \psi_n \leftarrow \psi$  ▷ Copy parameters
8:        $\bar{\theta}_n, \bar{\psi}_n \leftarrow \text{PointUpdate}(\mathcal{I}_{k,n}, \alpha, \theta_n, \psi_n)$  ▷ Inner loop update
9:        $\hat{d}_n, \mathbf{z}_{k,n}^b = \mathcal{F}(I_{k,n}^l, I_{k,n}^r | \bar{\theta}_n)$ 
10:       $\mathcal{L}_k \leftarrow \mathcal{L}_k + \mathcal{L}_b(\hat{d}_n, d_{k,n})$  ▷ Base loss
11:    end for
12:     $\theta' \leftarrow \theta - \beta \nabla_{\theta} \mathcal{L}_k$  ▷ Outer loop update
13:     $\psi' \leftarrow \psi - \beta \nabla_{\psi} \mathcal{L}_k$  ▷ Outer loop update
14:     $\theta, \psi \leftarrow \text{PointUpdate}(\mathcal{I}_k, \alpha, \theta', \psi')$  ▷ Adaptation
15:  end for
16:  return  $\theta^* \leftarrow \theta$ 
17: end function

```

points $\mathbf{p}(\theta)$ with the base parameter θ can be derived as:

$$\mathbf{p}(\theta) = \{(i, j) \mid |\hat{d}_{ij} - d_{ij}|_1 > 3\}. \quad (3)$$

Note that we represent \mathbf{p} as a function of θ to indicate that the selected point varies depending on θ updated in the learning procedure. In next section, we describe a way to leverage the set of points for correcting local distortions caused by the seen domain bias, depicted in Fig. 1.

4.2 PointFixNet

To mitigate the seen domain bias of the base network, we deploy an additional auxiliary network, called PointFixNet, which individually repairs a disparity by incurring a proper point-wise gradient. The PointFixNet consists of two modules: a feature extraction module (parameterized by ψ^c) that extracts feature \mathbf{z}^c from heterogeneous inputs; and a point-wise prediction module (parameterized by ψ^p) that generates residual disparity value of each point and back-propagates the point-wise errors. Specifically, the feature extraction module consists of three convolution layers and takes the left image, I^l , the initial disparity \hat{d} , and ground-truth d as inputs to integrate context around each erroneous pixel. Therefore, the feature \mathbf{z}^c can be obtained as follows:

$$\mathbf{z}^c = \mathcal{F}(I^l, \hat{d}, d | \psi^c), \quad (4)$$

where \mathcal{F} is a feed-forward process.

Then, the base feature \mathbf{z}^b from the base network and feature \mathbf{z}^c from the feature extraction module are concatenated and fed into the point-wise prediction module to generate the residual disparity value r_{ij} . Inspired by structure in [18], the module consists of four fully-connected (FC) layers to produce a single value for each pixel, such that:

$$r_{ij} = \mathcal{F}(\mathbf{z}_{ij}^b, \mathbf{z}_{ij}^c | \psi^p), \quad (5)$$

where $(i,j) \in \mathbf{p}(\theta)$. The final disparity for (i,j) -th pixel is obtained by adding r_{ij} to \hat{d}_{ij} . We note that FC layers share weights across all selected points.

4.3 Learning to Fix

The key idea underlying our framework is iterating *learning how to fix* first and then *fixing* alternatively. First, the parameters of the PointFixNet learn how to generate a proper gradient to the base model in a point-wise manner such that the base model can be improved with less domain bias and then secondly the base model is updated by the learned PointFixNet. This strategy is essential because if we keep network training only with point loss, the performance of prediction after PointFixNet can be guaranteed but the one after the base model may not. Thus, to enhance maximal performance of the base network, it is necessary to employ the alternative meta-learning structure. To this end, we deploy two loss functions: a base loss, \mathcal{L}_b , derived from the whole disparity map predicted by the base network and a point loss, \mathcal{L}_p , applied to the final disparity values. They are alternatively optimized to update θ and ψ by relying on the two-stage meta-learning scheme [20, 5], as described in Fig. 2 and Alg. 2.

In the inner loop, parameters are copied for each sample in batch, $\theta_n \leftarrow \theta$ and $\psi_n \leftarrow \psi$. Then we calculate the PointFix loss, to evaluate of the current parameters:

$$\bar{\theta}_n \leftarrow \theta_n - \alpha \nabla_{\theta_n} \sum_{(i,j) \in \mathbf{p}(\theta)} \mathcal{L}_p(\hat{d}_{ij} + r_{ij}, d_{ij}), \quad (6)$$

$$\bar{\psi}_n \leftarrow \psi_n - \alpha \nabla_{\psi_n} \sum_{(i,j) \in \mathbf{p}(\theta)} \mathcal{L}_p(\hat{d}_{ij} + r_{ij}, d_{ij}), \quad (7)$$

where α is a learning rate, $\mathbf{p}(\theta)$ is a set of selected points, and \mathcal{L}_p is a point-wise ℓ_1 loss between the final disparity and its corresponding ground-truth. Given in Alg. 1, since the PointFix loss is imposed on the local distortion of the erroneous pixels selected on the initial prediction, we can update base parameters that refer to the fine-grained details.

In the outer loop, we evaluate the performance of the updated base parameter after the inner loop. To measure the performance, we apply the conventional supervised loss between the initial disparity map and ground-truth. Following the procedure of [4], the parameters θ and ψ are updated based on sum of \mathcal{L}_k as follows:

$$\theta' \leftarrow \theta - \beta \nabla_{\theta} \mathcal{L}_k, \quad \psi' \leftarrow \psi - \beta \nabla_{\psi} \mathcal{L}_k, \quad (8)$$

where $\mathcal{L}_k = \sum_N \mathcal{L}_b(\hat{d}_n, d_{k,n})$ and k is the current iteration step. Note that the gradients are computed along with the parameters before being updated in the inner loop.

Unlike traditional MAML where the parameters are optimized via meta-update only, we deploy an additional update inspired by online meta-learning. θ' and ψ' are updated in the same way as the inner loop so that the parameters after final update, θ and ψ can be written as:

$$\theta \leftarrow \theta' - \alpha \nabla_{\theta'} \sum_{(i,j) \in \mathbf{p}(\theta')} \mathcal{L}_{\mathbf{p}}(\hat{d}_{ij} + r_{ij}, d_{ij}), \quad (9)$$

$$\psi \leftarrow \psi' - \alpha \nabla_{\psi'} \sum_{(i,j) \in \mathbf{p}(\theta')} \mathcal{L}_{\mathbf{p}}(\hat{d}_{ij} + r_{ij}, d_{ij}), \quad (10)$$

where $r_{ij} = \mathcal{F}(\mathbf{z}_{ij}^b, \mathbf{z}_{ij}^c, d_{ij} | \bar{\psi}^p)$ for $(i, j) \in \mathbf{p}(\theta')$. Finally, the networks are updated with PointFix loss which are, at the first stage, trained to generate proper back-propagation to enhance the performance of the base network in the next training step. At test time, we use the final parameters of the base network θ^* and perform adaptation according to Eq. (1).

5 Experiments

5.1 Experimental Settings

Datasets. In order to evaluate our method on realistic scenario, we use synthetic dataset for offline training and real dataset for the test. Therefore, the training and test data exist in completely different data distributions. Following the previous work [36], we train our networks using the Synthia [30] dataset and evaluate each model on the KITTI-raw [7] dataset and the subset of the DrivingStereo [40] dataset. All datasets are recorded in driving scene but the Synthia [30] is synthetic data, the KITTI and DrivingStereo are obtained from real world. The Synthia [30] dataset contains 50 sequences which have different combination of weathers, seasons and locations. To set similar disparity ranges in training and test, we resized Synthia [30] dataset images to half resolution as in [36]. We exploit stereo images from front direction only and there are 45,591 total number of stereo frames. The KITTI [7] dataset consists of 71 sequences and total 42,917 frames of stereo images and sparse depth maps. Different from the scenario on the KITTI [7] dataset, we present an additional adaptation scenario that the models adapt to various unseen weather conditions using the subset of the DrivingStereo [40] dataset. The DrivingStereo contains four different weather sequences (*i.e.*, cloudy, foggy, rainy, and sunny) that each sequence includes 500 stereo images with high quality labels obtained from multi-frame LiDAR points.

Metrics. We evaluate the performance using two popular evaluation metrics, the percentage of pixels with disparity outliers larger than 3 (D1-all) and average end point error (EPE). Following the scheme of [37, 36], we average each score from all the frames which belong to the same sequence and reset the model to the base parameters at the next sequence, based on the definition of a sequence for different evaluation protocols.

Evaluation protocols. We perform online stereo adaptation under three different settings according to the definition of the sequence, including short-, mid-,

Table 1. Mid-term adaptation: Performance comparison for several methods on the KITTI [7] dataset.

Method	Training	Adapt.	City		Residential		Campus		Road		Avg.	
			D1-all	EPE	D1-all	EPE	D1-all	EPE	D1-all	EPE	D1-all	EPE
DispNetC-GT	KITTI	No	1.94	0.68	2.43	0.77	5.43	1.10	1.67	0.69	2.87	0.81
MADNet-GT	KITTI	No	2.05	0.65	2.67	0.82	6.87	1.24	1.57	0.66	3.29	0.84
L2A-Disp.	Synthia	No	12.78	1.67	12.80	1.72	17.57	2.06	12.34	1.59	13.87	1.76
MADNet	Synthia	No	38.78	8.36	35.73	7.89	40.59	7.68	38.31	8.77	38.35	8.18
Ours-Disp.	Synthia	No	9.98	1.47	10.99	1.62	17.01	2.06	7.98	1.33	11.49	1.62
Ours-MAD.	Synthia	No	15.51	1.82	14.24	1.78	22.40	3.04	15.61	1.84	16.94	2.12
L2A-Disp.	Synthia	Full	2.05	0.78	2.57	0.86	4.43	1.07	1.63	0.77	2.67	0.87
MADNet	Synthia	Full	2.11	0.81	2.79	0.90	6.24	1.41	1.60	0.72	3.19	0.96
Ours-Disp.	Synthia	Full	2.03	0.99	2.46	0.83	4.21	1.02	1.58	0.74	2.57	0.90
Ours-MAD.	Synthia	Full	1.55	0.72	1.55	0.70	3.84	1.08	1.15	0.67	2.02	0.79
MADNet	Synthia	MAD	2.36	0.84	1.94	0.77	10.03	1.70	2.27	0.83	4.15	1.04
MADNet	Synthia	MAD++	1.95	0.80	1.86	0.76	8.57	1.65	1.94	0.80	3.56	0.99
Ours-MAD.	Synthia	MAD	1.63	0.74	1.62	0.73	4.16	1.12	1.23	0.69	2.16	0.82

and long-term adaptation. For **short-term adaptation**, each sequence is defined as a distinct sequence provided by the dataset (*e.g.* *2011_09_30_drive_0028_sync*). This setting is appeared in [36]. The sequences in **mid-term adaptation** are divided according to the environment (*i.e.*, *City, Residential, Campus, Road*). In **long-term adaptation**, we perform adaptation for all frames by concatenating all mid-term sequences. The mid- and long-term adaptation settings are shown in [37] as short- and long-term adaptation. The implementation details are provided in supplementary material.

5.2 Synthetic to Real Adaptation

We evaluate the performance corresponding to the different adaptation methods: *No adaptation (No)*, which measures the performance for all sequences without performing adaptation from the base parameters to estimate the capacity of the initial parameters; *Full adaptation (Full)*, which updates parameters of whole network; *MAD adaptation (MAD)*, which performs faster modular adaptation on a prediction of certain resolution selected at every iteration by their own handcrafted method as proposed in [37]; *MAD++ adaptation (MAD++)* is an extension from MAD and utilizes predictions obtained by handcrafted methods (*e.g.* SGM [11], WILD [38]) as proxy supervision. The cases of *-GT* are regarded as supervised learning that is fine-tuned on KITTI 2012 [8] and 2015 [24] datasets. In the experiment, (L2A, Ours)-Disp. and Ours-MAD. employ DispNetC and MADNet as base networks, respectively.

Mid-term adaptation. Table 1 shows the results according to the adaptation methods under mid-term adaptation setting. From the performance with *No adaptation* (row 3-6), we observe that our method helps to learn better base parameters compared to other methods using the same base network. Especially before the adaptation, with MADNet (row 4 and 6), our PointFix significantly improves the performance of the base network by 20.19% and 4.39 in terms of D1-all and EPE respectively, that demonstrates PointFix is effective in a generalization capability.

Table 2. Short-term and Long-term adaptation: Performance comparison for several methods on the KITTI [7] dataset.

Method	Training	Adapt.	Short-term		Long-term	
			D1-all	EPE	D1-all	EPE
DispNetC-GT	KITTI	No	2.38	0.77	2.32	0.75
MADNet-GT	KITTI	No	2.57	0.75	2.52	0.78
L2A-Disp.	Synthia	No	12.99	1.73	12.86	1.70
MADNet	Synthia	No	27.63	3.59	36.77	8.09
Ours-Disp.	Synthia	No	9.47	1.21	10.56	1.57
Ours-MAD.	Synthia	No	11.59	1.67	23.50	3.23
L2A-Disp.	Synthia	Full	2.64	0.84	2.37	0.84
MADNet	Synthia	Full	6.68	1.31	2.86	0.93
Ours-Disp.	Synthia	Full	2.62	0.81	2.30	0.82
Ours-MAD.	Synthia	Full	2.00	0.74	1.56	0.71
MADNet	Synthia	MAD	11.82	1.90	1.92	0.75
MADNet	Synthia	MAD++	9.56	1.61	1.70	0.75
Ours-MAD.	Synthia	MAD	2.64	0.87	1.47	0.70

The results of *Full adaptation* (row 7-10) and *MAD adaptation* (row 11-13) show the adaptation capability of the base network. Our PointFix outperforms previous methods with large margin in both metrics, achieving state-of-the-art performance on all domains. The PointFix with MAD adaptation (row 13) outperforms MADNet with MAD++ adaptation that leverages the additional supervision and even L2A and MADNet with full adaptation (row 7 and 8) while enabling fast inference by adapting only a few parameters. As pointed out in [37], all adapted models perform worse on *Campus* domain that has a small number of frames (1149) compared to the other domain (5674, 28067, 8027). The results on *Campus* domain show that PointFix adapt better than previous works with a small number of frames in all adaptation methods.

Short-term adaptation. In Table 2, to examine the adaptability in short sequences, we evaluate models on each sequence independently as represented in [36]. For each sequence, parameters are initialized at every beginning of sequences. Measured performance is first averaged in each sequence and then they are averaged out. Thanks to fast adaptation speed and inherent robustness of our framework, we surpass the performance than previous works with a large margin. Especially, due to its light weight structure, inherent weakness of MADNet is maximized in short-term environment (row 4 and 8) because they requires a number of frames to be adapted. Nevertheless, MADNet with our framework shows superior results. This suggests that our framework is worthy to be developed with light weight networks.

We conduct additional experiments on the DrivingStereo [40] dataset to hypothesize more difficult scenarios under various weather conditions. Specifically, we train all models on Synthia [30] dataset and evaluate the performance on each sequence including four types of weather conditions in the short-term adaptation setting. As shown in Table 3, our model outperforms MADNet [37] with a large margin for all novel weather conditions. In particular, our method represents

Table 3. Short-term adaptation: Performance comparison on the subset of DrivingStereo [40] under different weather conditions.

Method	Adapt.	cloudy		foggy		rainy		sunny		Avg.	
		D1-all	EPE	D1-all	EPE	D1-all	EPE	D1-all	EPE	D1-all	EPE
MADNet	No	56.83	19.16	70.14	23.85	54.20	19.20	51.30	16.08	58.12	19.57
Ours-MAD.	No	32.76	5.34	37.25	6.55	34.06	4.78	30.00	4.58	33.52	5.31
MADNet	Full	15.71	3.11	18.09	3.28	18.37	2.86	14.71	2.63	16.72	2.97
Ours-MAD.	Full	7.00	1.28	7.25	1.39	15.31	2.52	8.33	1.49	9.47	1.67
MADNet	MAD	28.46	6.65	33.56	6.13	31.34	5.79	27.10	6.15	30.11	6.18
Ours-MAD.	MAD	8.46	1.46	8.57	1.46	11.99	1.69	8.90	1.52	9.48	1.53

error rates of about half those of MADNet [37]. The implementation details and additional experimental results are provided in supplementary material.

Long-term adaptation. The adaptation on a long sequence followed by various environments without network resets can be regarded as the most practical scenario in the real world. To simulate this scenario, we report the results evaluated on the concatenation of four environments of the KITTI [7] dataset (~ 43000 frames) in Table 2. As analyzed in [37], the results show much smaller average errors than the mid- and short-term adaptation for all adaptation methods, as the length of the sequence increased. Among them, our PointFix shows drastically improved performance and significantly outperforms previous works. Therefore, PointFix framework can be further improved, continually adapting to the real world environment.

5.3 Analysis

Ablation study. To investigate the effectiveness of the components of within our model, we conduct ablation experiments on the KITTI [7] dataset according to PointFixNet and the meta-learning framework (ML), as shown in Table 4. Note that we use MADNet as the base network and evaluate the performance using the full adaptation under the short-term adaptation setting for all experiments in this section. As a baseline, we remove all components of the proposed method such that the first row in Table 4 corresponds to MADNet [37].

Table 4. Ablation studies for PointFixNet and the meta-learning framework (ML) evaluated on the KITTI [7] dataset under short-term adaptation setting.

PointFixNet	ML	Adapt.	D1-all	EPE
\times	\times	Full	6.68	1.31
\checkmark	\times	Full	8.06	1.47
\times	\checkmark	Full	3.12	0.96
\checkmark	\checkmark	Full	2.00	0.74

Effectiveness of PointFixNet. To validate the effectiveness of the point-wise backpropagation, we ablate PointFixNet and apply ℓ_1 loss between the initial disparity \hat{d} and groundtruth d instead of the point loss in Alg. 1. The comparison between the third row of Table 4 and the full use of components shows PointFixNet contributes 1.12% and 0.22 in terms of D1-all error and EPE and demonstrates fixing local detriments is simple yet effective to improve the robustness of the stereo model.




Fig. 3. Disparity maps predicted using MADNet as the base network on the KITTI [7] sequence. (a) Left images, (b) MADNet with MAD adaptation [37], (c) Ours-MAD with MAD adaptation. Red pixel values indicate closer objects.

Effectiveness of ML. As described in 4.3, the performance improvement of the base network using the point loss is not guaranteed without meta-learning. The results in the second row of Table 4 show poor performance even than the baseline. The comparison between the first and third rows of Table 4 further validates the effectiveness of the ML framework, showing significant performance improvements. Finally, the state-of-the-art performance is shown by demonstrating the advantage of the full use of all components of learning to fix the base network through meta-learning.

Convergence. In Fig. 3, we evaluate the qualitative results related to the convergence analysis. The results contain good initialized parameters (first column), fast adaptation (second column), and convergence to low errors (last column) as analyzed above. As the adaptation proceeds, the MADNet (row 2) estimates better prediction, yet still shows a high error while Ours-MAD (row 3) shows not only robust initial performance but also faster convergence to low errors.

To analyze and compare the adaptation cost corresponding to the methods, we visualize the adaptation performance over frames of the sequence from the KITTI [7] dataset in Fig. 4¹. In overall view, the results show that our PointFix adapts faster than [37] and converges with lower errors regardless of the adaptation method. Furthermore, the comparison between **MAD.-No** (green) and **Ours+Mad.-No** (brown) shows the effectiveness of the initial base parameters. The comparison between **MAD.-MAD** (yellow) and **Ours+Mad.-MAD** (purple) shows that the performance is improved by PointFix with MAD adap-

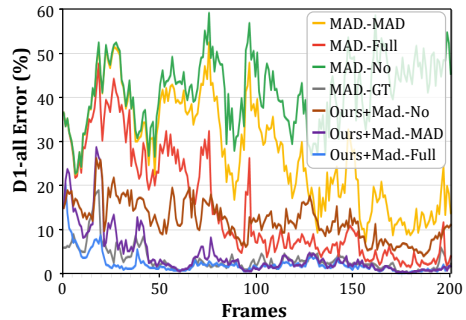


Fig. 4. D1-all error (%) across frames in sequence from the KITTI [7] dataset with respect to the adaptation methods.

¹ The results with DispNetC are shown in supplementary materials.

tation, while **MAD.-Full** (red) overtakes from the about 100-th frame. Finally, **Ours+Mad.-Full** (blue) adapts faster than all the other methods and converges to the low D1-all error, showing comparable performance with **MAD.-GT** (gray) fine-tuned with ground-truth. The additional experimental results and analysis are shown in supplementary material.

Comparison with domain generalization methods. To argue the practicality of online stereo adaptation, we compare our model with the state-of-the-art domain generalization (DG) methods [42, 32, 21] that the stereo models are trained on the synthetic dataset and evaluated on the unseen real dataset without the additional adaptation.

For a fair comparison, the models are pretrained on the SceneFlow [23] dataset and evaluated on the KITTI [7] dataset². Note that we report the performance of our model measured under the long-term adaptation setting. As shown in Table 5³, our model not only outperforms the generalization approaches in terms

of D1-all error but also shows about $\times 27$, $\times 8.3$, and $\times 1.6$ faster inference speed than [42], [32], and [21], respectively, despite of additional adaptation steps. While the domain generalization approaches [42, 32, 21] estimate the depth maps without the adaptation, they require a large number of parameters to obtain a generalized stereo model, making them impractical. This is worth noting that our method has high applicability to the practical application such as autonomous driving in terms of accuracy and inference speed.

Table 5. Comparisons of domain generalization methods and our method evaluated on the KITTI [7] dataset.

Method	Adapt.	D1-all	EPE	FPS
DSMNet [42]	N/A	1.59	0.68	1.30
CFNet [32]	N/A	1.93	0.97	4.27
Raft [21]	N/A	1.66	0.71	22.44
MADNet [37]	MAD	1.95	0.82	35.7
Ours-MAD.	MAD	1.47	0.70	35.7

6 Conclusion

In this paper, we proposed PointFix, a novel meta-learning framework to effectively adapt any deep stereo models in online setting. Compared with previous online stereo adaptation approaches facing global domain bias problem to synthetic data, our model can induce maximal performance of the base stereo networks by the proposed the auxiliary network PointFixNet and learning-to-fix strategy, that can adapt well to the fine-grained domain gap. Our extensive experiments show PointFix achieves state-of-the-art results, outperforming several online stereo adaptation methods in a wide variety of environments. In addition, the results demonstrate that PointFix is capable of improving the generalization ability of the stereo models.

² We conducted a custom experiment using a publicly available code for each paper.

³ For Raft-stereo [21], a real-time version was employed which shows a much faster inference speed.

References

1. Achtelik, M., Bachrach, A., He, R., Prentice, S., Roy, N.: Stereo vision and laser odometry for autonomous helicopters in gps-denied indoor environments. In: SPIE Defense, Security, and Sensing pp. 733219–733219 (2009)
2. Chang, J.R., Chen, Y.S.: Pyramid stereo matching network. In: CVPR (2018)
3. Chen, Y., Van Gool, L., Schmid, C., Sminchisescu, C.: Consistency guided scene flow estimation. In: ECCV (2020)
4. Finn, C., Abbeel, P., Levine, S.: Model-agnostic meta-learning for fast adaptation of deep networks. In: PMLR (2017)
5. Finn, C., Rajeswaran, A., Kakade, S., Levine, S.: Online meta-learning. In: ICML (2019)
6. Garg, R., Bg, V.K., Carneiro, G., Reid, I.: Unsupervised cnn for single view depth estimation: Geometry to the rescue. In: ECCV (2016)
7. Geiger, A., Lenz, P., Stiller, C., Urtasun, R.: Vision meets robotics: The kitti dataset. *Int. J. of Rob. Res.* **32**(11), 1231–1237 (2013)
8. Geiger, A., Lenz, P., Urtasun, R.: Are we ready for autonomous driving? the kitti vision benchmark suite. In: CVPR (2012)
9. Geiger, A., Roser, M., Urtasun, R.: Efficient large-scale stereo matching. In: ACCV (2010)
10. Godard, C., Mac Aodha, O., Brostow, G.J.: Unsupervised monocular depth estimation with left-right consistency. In: CVPR (2017)
11. Hirschmuller, H.: Stereo processing by semiglobal matching and mutual information. *IEEE Trans. Pattern Anal. Mach. Intell.* **30**(2), 328–341 (2008)
12. Hoffman, J., Wang, D., Yu, F., Darrell, T.: Fcns in the wild: Pixel-level adversarial and constraint-based adaptation. arXiv preprint arXiv:1612.02649 (2016)
13. Hu, X., Mordohai, P.: A quantitative evaluation of confidence measures for stereo vision. *IEEE Trans. Pattern Anal. Mach. Intell.* **34**(11), 2121–2133 (2012)
14. Kendall, A., Martirosyan, H., Dasgupta, S., Henry, P., Kennedy, R., Bachrach, A., Bry, A.: End-to-end learning of geometry and context for deep stereo regression. In: ICCV (2017)
15. Khamis, S., Fanello, S., Rhemann, C., Kowdle, A., Valentin, J., Izadi, S.: Stereonet: Guided hierarchical refinement for real-time edge-aware depth prediction. In: ECCV (2018)
16. Kim, S., Kim, S., Min, D., Sohn, K.: Laf-net: Locally adaptive fusion networks for stereo confidence estimation. In: CVPR (2019)
17. Kim, S., Min, D., Kim, S., Sohn, K.: Unified confidence estimation networks for robust stereo matching. *IEEE Trans. Image Process.* **28**(3), 1299–1313 (2018)
18. Kirillov, A., Wu, Y., He, K., Girshick, R.: Pointrend: Image segmentation as rendering. In: CVPR (2020)
19. Knowles, M., Peretroukhin, V., Greene, W.N., Roy, N.: Toward robust and efficient online adaptation for deep stereo depth estimation. In: ICRA (2021)
20. Li, D., Hospedales, T.: Online meta-learning for multi-source and semi-supervised domain adaptation. In: ECCV (2020)
21. Lipson, L., Teed, Z., Deng, J.: Raft-stereo: Multilevel recurrent field transforms for stereo matching. In: 3DV (2021)
22. Liu, M.Y., Tuzel, O.: Coupled generative adversarial networks. In: NIPS (2016)
23. Mayer, N., Ilg, E., Häusser, P., Fischer, P., Cremers, D., Dosovitskiy, A., Brox, T.: A large dataset to train convolutional networks for disparity, optical flow, and scene flow estimation. In: CVPR (2016)

24. Menze, M., Geiger, A.: Object scene flow for autonomous vehicles. In: CVPR (2015)
25. Mnih, V., Kavukcuoglu, K., Silver, D., Rusu, A.A., Veness, J., Bellemare, M.G., Graves, A., Riedmiller, M., Fidjeland, A.K., Ostrovski, G., et al.: Human-level control through deep reinforcement learning. *Nature* **518**(7540), 529–533 (2015)
26. Pang, J., Sun, W., Yang, C., Ren, J., Xiao, R., Zeng, J., Lin, L.: Zoom and learn: Generalizing deep stereo matching to novel domains. In: CVPR (2018)
27. Patricia, N., Carlucci, F.M., Caputo, B.: Deep depth domain adaptation: A case study. In: ICCV (2017)
28. Poggi, M., Kim, S., Tosi, F., Kim, S., Aleotti, F., Min, D., Sohn, K., Mattoccia, S.: On the confidence of stereo matching in a deep-learning era: a quantitative evaluation. *IEEE Trans. Pattern Anal. Mach. Intell.* (2021)
29. Poggi, M., Tonioni, A., Tosi, F., Mattoccia, S., Di Stefano, L.: Continual adaptation for deep stereo. *IEEE Trans. Pattern Anal. Mach. Intell.* pp. 1–1 (2021)
30. Ros, G., Sellart, L., Materzynska, J., Vazquez, D., Lopez, A.M.: The synthia dataset: A large collection of synthetic images for semantic segmentation of urban scenes. In: CVPR (2016)
31. Rusu, A.A., Rabinowitz, N.C., Desjardins, G., Soyer, H., Kirkpatrick, J., Kavukcuoglu, K., Pascanu, R., Hadsell, R.: Progressive neural networks. arXiv preprint arXiv:1606.04671 (2016)
32. Shen, Z., Dai, Y., Rao, Z.: Cfnet: Cascade and fused cost volume for robust stereo matching. In: CVPR (2021)
33. Sun, D., Yang, X., Liu, M.Y., Kautz, J.: Pwc-net: Cnns for optical flow using pyramid, warping, and cost volume. In: CVPR (2018)
34. Tankovich, V., Hane, C., Zhang, Y., Kowdle, A., Fanello, S., Bouaziz, S.: Hitnet: Hierarchical iterative tile refinement network for real-time stereo matching. In: CVPR (2021)
35. Tonioni, A., Poggi, M., Mattoccia, S., Di Stefano, L.: Unsupervised adaptation for deep stereo. In: ICCV (2017)
36. Tonioni, A., Rahnama, O., Joy, T., Stefano, L.D., Ajanthan, T., Torr, P.H.: Learning to adapt for stereo. In: CVPR (2019)
37. Tonioni, A., Tosi, F., Poggi, M., Mattoccia, S., Stefano, L.D.: Real-time self-adaptive deep stereo. In: CVPR (2019)
38. Tosi, F., Poggi, M., Mattoccia, S., Tonioni, A., di Stefano, L.: Learning confidence measures in the wild. *BMVC* (2017)
39. Wang, H., Wang, X., Song, J., Lei, J., Song, M.: Faster self-adaptive deep stereo. In: ACCV (2020)
40. Yang, G., Song, X., Huang, C., Deng, Z., Shi, J., Zhou, B.: Drivingstereo: A large-scale dataset for stereo matching in autonomous driving scenarios. In: CVPR (2019)
41. Zbontar, J., LeCun, Y., et al.: Stereo matching by training a convolutional neural network to compare image patches. *J. Mach. Learn. Res.* **17**(1), 2287–2318 (2016)
42. Zhang, Qi, X., Yang, R., Prisacariu, V., Wah, B., Torr, P.: Domain-invariant stereo matching networks. In: ECCV (2020)
43. Zhang, K., Lu, J., Lafruit, G.: Cross-based local stereo matching using orthogonal integral images. *IEEE Trans. Circuit Syst. Video Technol.* **19**(7), 1073–1079 (2009)
44. Zhang, Y., Khamis, S., Rhemann, C., Valentin, J., Kowdle, A., Tankovich, V., Schoenberg, M., Izadi, S., Funkhouser, T., Fanello, S.: Activestereonet: End-to-end self-supervised learning for active stereo systems. In: ECCV (2018)
45. Zhang, Z., Lathuilière, S., Pilzer, A., Sebe, N., Ricci, E., Yang, J.: Online adaptation through meta-learning for stereo depth estimation. arXiv preprint arXiv:1904.08462 (2019)

46. Zhou, C., Zhang, H., Shen, X., Jia, J.: Unsupervised learning of stereo matching. In: ICCV (2017)