

# ITSA: An Information-Theoretic Approach to Automatic Shortcut Avoidance and Domain Generalization in Stereo Matching Networks

WeiQin Chuah\* Ruwan Tennakoon\* Reza Hoseinnezhad\* Alireza Bab-Hadiashar\* David Suter†  
 RMIT University, Australia\* Edith Cowan University (ECU), Australia†  
 {wei.qin.chuah, ruwan.tennakoon, rezah, abh}@rmit.edu.au, d.suter@ecu.edu.au

## Abstract

*State-of-the-art stereo matching networks trained only on synthetic data often fail to generalize to more challenging real data domains. In this paper, we attempt to unfold an important factor that hinders the networks from generalizing across domains: through the lens of shortcut learning. We demonstrate that the learning of feature representations in stereo matching networks is heavily influenced by synthetic data artefacts (shortcut attributes). To mitigate this issue, we propose an Information-Theoretic Shortcut Avoidance (ITSA) approach to automatically restrict shortcut-related information from being encoded into the feature representations. As a result, our proposed method learns robust and shortcut-invariant features by minimizing the sensitivity of latent features to input variations. To avoid the prohibitive computational cost of direct input sensitivity optimization, we propose an effective yet feasible algorithm to achieve robustness. We show that using this method, state-of-the-art stereo matching networks that are trained purely on synthetic data can effectively generalize to challenging and previously unseen real data scenarios. Importantly, the proposed method enhances the robustness of the synthetic trained networks to the point that they outperform their fine-tuned counterparts (on real data) for challenging out-of-domain stereo datasets.*

## 1. Introduction

Stereo matching is a fundamental task in computer vision and is widely used for depth sensing in various applications such as augmented reality (AR), robotics and autonomous driving. In recent years, end-to-end trained Convolutional Neural Networks (CNNs) have achieved impressive results for this task as quantified by the performance on several publicly available stereo-matching benchmarks [6, 16, 18, 45, 52].

Generally, end-to-end stereo-matching networks require a large amount of labelled data for training. To overcome

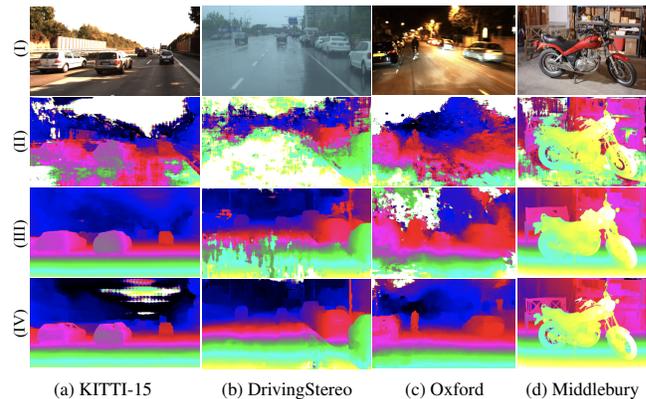


Figure 1. Comparison of disparity maps estimated by PSMNet [6] when it is trained under different settings and across multiple domains. Each column shows the results for a realistic domain namely: KITT1 2015 [26], DrivingStereo [47], Oxford Robotcar [24] and Middlebury [35]. Rows from top to bottom show a sample image (I), the prediction for the Scene Flow pre-trained model (II), KITT1-15 fine-tuned model (III), and the proposed ITSA optimized method (IV). Comparing these figures shows that PSMNet trained solely on synthetic data performs poorly on real data and fine tuning only improves the result for KITT1 dataset (still fails to generalize for other scenarios). The proposed method performs well across the board (best viewed in color).

this challenge, many state-of-the-art networks are initially trained on labelled synthetic data, commonly generated using game engines. However, models trained using synthetic data do not generalize well to unseen realistic domains. For example, the PSMNet [6] pre-trained on the Scene Flow dataset [25] performs poorly when tested on unseen realistic domains as illustrated in Fig. 1. Therefore, in practice, the networks trained with synthetic data are fine-tuned using labelled data from the relevant target domain. However, collecting even a relatively small amount of dense ground truth data in the real-world can be challenging for tasks like stereo-matching [22, 41]. Furthermore, to be practically useful in many applications, a stereo-matching model should be able to generalize effortlessly to different do-

mains like day and night times, varying weather conditions, etc. Collecting data for fine-tuning that cover all possible situations is both difficult and expensive. It is therefore highly desirable to remove the fine-tuning requirement.

It is known that neural networks, including stereo matching networks, can learn superficial shortcut features (or spurious correlations with the target labels), which prevent them from generalizing across different domains [2, 12]. We found that stereo matching networks trained on synthetic data are susceptible to exploiting shortcuts in synthetic data such as (1) consistent local statistics (RGB color features) between the left and right stereo images and (2) over-reliance on local chromaticity features (e.g. color, illumination, texture) of the reference stereo viewpoint. Detailed analysis and discussion are included in Sec. 4.2. Dependency on these shortcut cues, instead of the desirable semantic and structural representations, means that these networks would fail drastically when the spurious correlations between shortcuts and labels do not exist in a new (unseen) domain [33]. While several shortcut-removal approaches have been previously proposed [4, 17, 38], most of these methods are manually designed (e.g. carefully selected data augmentations [4, 17]) and rely on the assumption that the shortcuts could be identified in advance. However, shortcuts can be non-intuitive, task-specific, and difficult to identify [9, 27].

Our goal is to train a stereo matching network on synthetic data that can generalize to realistic scenes without the need for fine-tuning. To achieve this, we propose an information-theoretic approach to automatically restrict the shortcut-related information from being encoded from the input into the feature representations. The approach is based on the well known information bottleneck (IB) principle that proposes to optimize the following objective [1, 40]:

$$\operatorname{argmax}_{\theta} I(Y, Z; \theta) - \beta I(X, Z; \theta) \quad (1)$$

where  $Z$  is the encoding of input  $X$ ,  $Y$  is the target,  $I$  is mutual information and  $\beta \in [0, 1]$  is the hyperparameter that controls the size of the information bottleneck. While optimizing the IB objective leads to compressed feature representations, our empirical experiments showed that these compressed features are neither robust nor shortcut-invariant (details are provided in Sec. 3.3.1). Consequently, the IB optimized networks may still incorporate shortcuts and remain fragile when tested in unseen domains. The recently introduced robust IB criterion [31] encourages the learning of both robust and compressive features by replacing the mutual information in IB with statistical Fisher information. Robust IB is presented in the context of learning features that are robust to adversarial attacks and to the best of our knowledge it has not been used for domain generalization.

In our approach, we combine the task loss (e.g. smooth L1 loss) with Fisher information to learn a generalizable stereo matching model. Although such an objective can work in theory, straightforward optimization of the Fisher information by gradient descent requires computation of the second-order derivatives and is therefore computationally expensive for tasks with high dimensional inputs such as stereo matching and semantic segmentation. To overcome this shortcoming, we propose ITSA which consists of a novel loss term and perturbation technique to approximate the optimization of the Fisher information loss. The proposed ITSA is computationally efficient, and as we show by extensive experiments, it can promote the learning of shortcut-invariant features. Unlike the existing domain-invariant stereo matching networks [37, 53], the proposed ITSA does not involve significant network alteration and is model-agnostic. Therefore, as shown in the experiments section, it can be easily integrated with different stereo matching networks.

The empirical results show that stereo-matching networks trained on synthetic data, with the proposed ITSA, can generalise to realistic data without fine-tuning. Additional experiments on challenging out-of-domain stereo datasets (e.g. different adverse weathers and night scenes) show that our method also improves the overall robustness of the stereo matching networks and importantly even outperforms the networks fine-tuned on realistic domains when tested on these challenging datasets. The main contributions of this paper include:

- We show that learning feature representations that are less sensitive to input variations can significantly enhance the synthetic to realistic domain generalization, and robustness in stereo matching networks.
- We introduce a novel loss function that enables us to minimize the Fisher information, without computing the second-order derivatives.
- We also show that the application of the proposed framework is not limited to stereo matching task, and can be used in training models for non-geometry based vision problems such as semantic segmentation.

The rest of the paper is organized as follows. Sec. 2 describes the related work in the field of learning-based stereo matching networks, domain generalization and shortcut learning. Sec. 3 presents the proposed method for automatic shortcut avoidance and domain generalization. Experimental results and discussions are presented in Sec. 4, and Sec. 6 concludes the paper.

## 2. Related Work

### Learning-based Stereo Matching Networks

In recent years, end-to-end learned deep stereo match-

ing networks have excelled in most datasets and benchmarks [6, 18, 45, 52]. These networks generally have three sub-modules (1) feature extraction sub-network, (2) cost-volume generator, and (3) cost aggregation and refinement sub-network. There are two main types of stereo matching networks based on how the cost volume is generated.

*Correlation-based stereo matching networks* construct the cost volume by correlating the features extracted from the two views. Previously proposed correlation-based methods include DispNetC [25], iResNet [21], CRL [30], SegStereo [48], and AANet [45]. Although these methods are usually computationally efficient, semantics and structural information in the feature representations are lost due to the correlation operation [16]. As a result, the correlation-based stereo matching methods usually have inferior performance compared to the concatenated-based methods.

*Concatenation-based methods* use a cost volume that is a simple assembly of features extracted from the two views. Examples of the state-of-the-art concatenation-based stereo matching networks include PSMNet [6], GANet [52], GC-Net [18], StereoDrNet [5] and EMCUA [28]. While these networks can achieve superior performance in stereo matching, they require labelled samples from the target environments, for fine-tuning. Without fine-tuning, these networks cannot generalize to unseen test data.

To overcome this problem, Zhang *et al.* [53] proposed DSMNet, which employs Domain Normalization and non-local graph-based filtering layers to enforce the learning of structural features that are domain-invariant. Similarly, Shen *et al.* [37] introduced CFNet, an efficient network architecture with multi-scale cost volume fusion and refinement, to enforce the learning of robust structural representation for stereo matching. In contrast, we have identified shortcut learning [13] as a major factor that hinders stereo matching networks from generalizing across domains. In this work, we show that avoiding shortcut learning can effectively enhance the robustness of the stereo matching networks and enables a model to generalize across domains. This is evidenced by showing networks’ superior performance on challenging realistic data without fine-tuning.

### Single Domain Generalization

Domain generalization typically involves forcing DNNs to learn domain-invariant features, using data sampled from multiple source domains [20, 32]. On the other hand, single domain generalization is a more challenging problem because only one source domain is available for training. To solve this problem, Volpi *et al.* [42] proposed adversarial data augmentation (ADA), which aims to expand and diversify the distribution of training data. Specifically, ADA creates “fictitious” yet “challenging” new populations, simulating data sampled from novel domains, using adversarial

training. In a similar fashion, Qiao *et al.* [32] proposed a novel framework that employs ADA and meta-learning to enforce the learning of domain-invariant features. While these works focus on minimizing the domain differences, we are interested in learning robust and shortcut-invariant features that are transferable across different domains. To this end, we propose ITSA, an information-theoretic approach to prevent shortcut learning (see next section), particularly in the stereo matching networks.

### Shortcut Learning

Geirhos *et al.* [12] coined the term shortcut learning as a phenomenon where DNNs learn trivial solutions by relying on superficial features (shortcuts). These features are spuriously correlated with the target labels, without contributing to transferability across contexts. For example, image classification networks tend to rely on shortcuts such as backgrounds [2, 12] and textures [14, 43] to improve their performance. However, these networks fail to generalize to unseen domains, where the spurious correlations between shortcuts and labels are violated [33]. Similarly, we observed that stereo matching networks trained on synthetic data also have a tendency to exploit shortcuts to produce accurate depth results in synthetic domains. Consequently, these networks fail drastically when tested in unseen realistic environments.

Several attempts have been made to restrict the learning of identified shortcuts and generalize DNNs across domains [4, 7, 17, 38, 43]. These methods rely on having some shortcut-related prior knowledge and usually include data augmentations [4, 17], whitening transformation [7] or dropout-based regularization [38] as part of their solutions. However, shortcuts are non-trivial, task-specific, and are often difficult to be identified a priori [9, 27]. In contrast, our proposed method automatically avoids shortcut learning without requiring shortcut-related knowledge in advance.

## 3. Methodology

### 3.1. Problem Definition

In this work, we focus on the synthetic-to-realistic domain generalization for stereo matching. Given a synthetic stereo data set  $D_{syn}$  consisting of stereo image pairs  $\{x_{syn,l}^{(i)}, x_{syn,r}^{(i)}\}_{i=1}^n$  with corresponding ground-truth disparity  $\{y_{syn}^{(i)}\}_{i=1}^n$ , the goal is to design a robust and shortcut-invariant stereo matching network that can accurately predict disparity map  $\hat{y}^{(i)}$  for unseen realistic environments  $D_{real}$ .

Our approach to achieve synthetic-to-realistic domain generalization is to use an information-theoretic measure to automatically restrict the shortcut-related information from being included in feature representations.

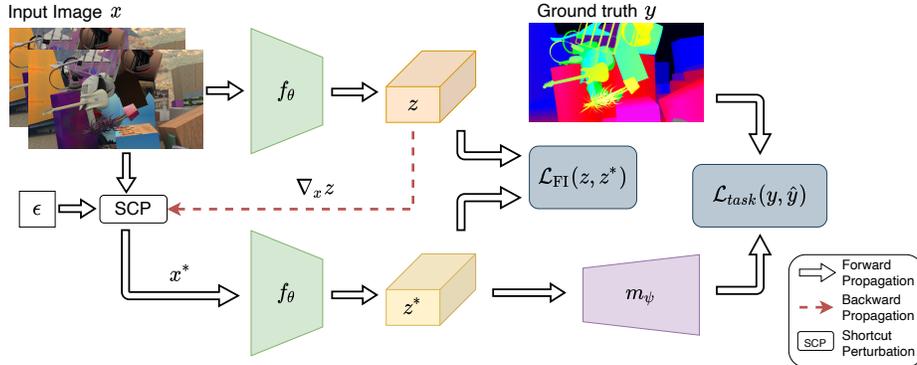


Figure 2. An overview of the proposed shortcut-avoidance strategy to achieve domain generalization in stereo matching networks. The parameters are shared across the two feature extractor networks  $f_\theta$  (best viewed in color).

### 3.2. Model

A typical stereo-matching network can be represented by the following equation:

$$\hat{y}^{(i)} = m_\psi \left( \mathbb{C} \left( f_\theta \left( x_l^{(i)} \right), f_\theta \left( x_r^{(i)} \right) \right) \right) \quad (2)$$

where  $f_\theta(\cdot)$  is the feature extraction sub-network,  $\mathbb{C}(\cdot)$  the cost volume and  $m_\psi(\cdot)$  the cost aggregation and refinement sub-network. The refined cost volumes are converted to disparity maps  $\hat{y}$  via the soft argmin [18] operation.

Our proposed method (ITSA) can be applied to any stereo-matching network that has the above structure. In the experiments section, we show the result of applying the proposed algorithm to different stereo-matching networks with concatenation cost (we observed similar results with correlation-based methods) volumes [6, 16, 37]. The high-level structure of the network including the proposed shortcut avoidance strategy is shown in Fig. 2.

### 3.3. Loss function

Our main contribution is the loss function devised to automatically restrict the shortcut-related information from being encoded in the learning process. As we explained earlier, the information bottleneck (IB) principle [1, 40] is typically used to compress features and would be a natural choice to achieve this objective.

The standard  $\mathcal{L}_{\text{IB}}$  loss defined in Eq. (1), which uses mutual information to quantify information content, was designed to extract features that are both concise and relevant for prediction. However, models trained by this loss are not robust to existence of artefacts that can generate shortcuts (similar to adversarial distortions mentioned in [31]).

To demonstrate the above point, we conducted a toy experiment. In this experiment, we investigated the efficacy of using IB loss for helping digit recognition networks (DRNs) to generalize from MNIST (source) [19] to MNIST-M [10]

Method	MNIST [19]	MNIST-M [10]
ERM	97.9 ± 0.14	40.9 ± 2.95
IB [1]	<b>99.0</b> ± 0.47	21.8 ± 0.21
RIB [31]	98.3 ± 0.13	52.8 ± 1.04
ITSA	98.1 ± 0.38	<b>56.9</b> ± 1.23

Table 1. Performance comparison of digit recognition networks optimized via empirical risk minimization (ERM), the information bottleneck (IB) [1], its robust variant (RIB) [31] and our proposed method (ITSA). While IB performs well in the in-domain tests, it performs poorly on out-of-domain tests.

(target) dataset. The former contains images of handwritten digits with black background, and the latter is created by combining the MNIST digits with randomly extracted color patches as their background. All networks were trained on the MNIST training set only and the top-1 accuracy (%) was employed for evaluation. The details of the experiment are included in the supplementary document. As shown in Tab. 1, the standard IB can effectively reduce over fitting, and achieves the best performance in the source domain. However, it fails to generalize its performance to the unseen domain. Importantly, it even performs worse than the baseline networks in the unseen target domain.

#### 3.3.1 Robust Information Bottleneck and Fisher Information

As our aim is to develop an IB based cost function that is not susceptible to existence of shortcuts in source data, we take inspiration from the robust IB principle [31]. Robust IB utilizes the statistical Fisher information  $\Phi(Z|X)$  of the extracted features  $Z$  parameterized by the inputs  $X$  as a more robust measure of information (in place of  $I(Z, X)$ ). The Fisher information  $\Phi(Z|X)$  is defined as:

$$\Phi(Z|X) = \int_{\mathcal{X}} \Phi(Z|X = x) p_X(x) dx, \quad (3)$$

where

$$\Phi(Z|X = x) = \int_{\mathcal{Z}} \|\nabla_x \log p_{Z|X}(z|x)\|_2^2 p_{Z|X}(z|x) dz. \quad (4)$$

The term  $\Phi(Z|X = x)$  in Eq. (3, 4) can be regarded as the sensitivity of the latent distribution  $p_{Z|X}(\cdot|x)$ , with respect to changes at the input  $x$ . Therefore, optimizing the Fisher information,  $\Phi(Z|X)$ , will minimize the average sensitivity of the latent distribution with respect to change of inputs  $X$ . As shortcuts are generated by data artefacts that are transient<sup>1</sup> by nature, they are sensitive to perturbations of input data [12]. As such, minimizing the Fisher information is a step towards promoting the learning of shortcut-invariant features. Our conjecture is supported by the results of the toy experiment included in Tab. 1. The DRNs constrained by the Fisher information (RIB) achieved better performance than the IB networks in the target domain.

In order to minimize the Fisher information expressed in Eq. (4), one has to compute second order derivatives such as  $\nabla_\theta \nabla_x \log p_{Z|X}(z|x)$ , which is computationally prohibitive for tasks with large dimensional inputs such as stereo matching, semantic segmentation, etc. [39]. To overcome this issue, we propose ITSA, a simple yet computationally feasible approach to promote the learning of shortcut-invariant features.

### 3.3.2 Approximating Fisher information

Optimizing the Fisher information  $\Phi(Z|X)$  measure defined in Eq. (3) is related to minimizing  $\Phi(Z|X = x)$ . By adding a regularization term such as  $\Phi(Z|X = x)$  to the loss function, we can penalize the transient features and discourages networks from learning shortcuts. To calculate this term, we employ a first order approximation as described below.

**Lemma 3.1.** *If  $\epsilon > 0$ ,  $u$  is a unit vector (i.e.  $\|u\| = 1$ , we refer to as the shortcut perturbation) and  $x^* = x + \epsilon u$ , then, subject to first order approximation:*

$$\Phi(Z|X = x) = \frac{\mathbb{E}_z \left[ \left| p_{Z|X=x^*}(z) - p_{Z|X=x}(z) \right|^2 \right]}{\epsilon^2 \cos^2 \psi} + \mathcal{V} \left[ \left\| \nabla_x \log p_{Z|X=x}(z) \right\|_2 \right] \quad (5)$$

where  $\mathbb{E}_z[v]$  and  $\mathcal{V}[v]$  are the expectation and variance of  $v$ , and  $\psi$  is the angle between  $u$  and  $\nabla_x p_{Z|X=x}$ .

Proof is given in the supplementary material.

The first term in the RHS of Eq. (5) will be minimized when the divergence (distance) between the two distributions,  $p_{Z|X=x}$  and  $p_{Z|X=x+\epsilon u}$ , is reduced. There are many

<sup>1</sup>We use transient to describe image attributes that are inconsistent across domains, and spuriously correlated with the true label. These features may include backgrounds, textures, image style, etc.

popular divergence measures between distributions, such as Kullback-Leibler divergence, Jensen-Shannon divergence, Total Variation, the Wasserstein distance, etc. In this work, we choose the Wasserstein distance: as the distributions  $p_{Z|X=x}$  and  $p_{Z|X=x+\epsilon u}$  may not have common supports and it leads to a simpler loss function.

In the case of a deterministic feature extractor, which is common in stereo matching networks, the distributions  $p_{Z|X=x}$  and  $p_{Z|X=x^*}$  can be seen as two degenerate distributions (i.e. Dirac delta distributions) located at points  $z = f_\theta(x)$  and  $z^* = f_\theta(x^*)$ . Furthermore, the  $\mathcal{V}[\cdot]$  in Eq. (5) will be zero. In this case, the Wasserstein- $p$  distance can be simplified as:

$$W_p(p_{Z|X=x^*}, p_{Z|X=x}) = \left( \|z^* - z\|_2^p \right)^{1/p}. \quad (6)$$

Using the above insights, we can see that minimizing  $\|z^* - z\|_2$  is a step towards minimizing  $\Phi(Z|X = x)$  (for  $p = 1$ ). Thus, we propose to promote the learning of robust and shortcut-invariant features in stereo matching networks, by optimizing the overall loss function defined below:

$$\mathcal{L} = \mathcal{L}_{smooth_{L1}}(\hat{y}, y) + \frac{\lambda}{2} \left( \mathcal{L}_{FI}(z_l, z_l^*) + \mathcal{L}_{FI}(z_r, z_r^*) \right) \quad (7)$$

where  $\hat{y}$  and  $y$  are the estimated and ground-truth disparity maps,  $\mathcal{L}_{FI}$  is our proposed Fisher information loss function defined as:

$$\mathcal{L}_{FI} = \sum_{i=1}^n \left\| z^{(i)} - z^{*(i)} \right\|_2 \quad (8)$$

and  $\mathcal{L}_{smooth_{L1}}$  is the smooth-L1 loss function commonly employed for optimizing stereo matching networks [6, 16, 52, 53].

### 3.3.3 Shortcut Perturbation (SCP)

In order to compute  $\mathcal{L}_{FI}$ , we need to define  $u$  (refer to as shortcut perturbation and is introduced in Lemma 3.1):  $u = \frac{\nabla_x z^{(i)}}{\|\nabla_x z^{(i)}\|_2}$  where  $\nabla_x z^{(i)}$  is the gradient of the extracted features  $z$  with respect to input. The shortcut-perturbed image can then be expressed as:

$$x^{*(i)} = x^{(i)} + \epsilon \frac{\nabla_x z^{(i)}}{\|\nabla_x z^{(i)}\|_2} \quad (9)$$

The above perturbation will put more weight on pixels that are sensitive to changes in the input. Intuitively, pixels with large absolute value of  $\nabla_x z$  will have significant impact in altering the statistics of encoded latent distributions and the extracted latent feature representations. Moreover, these pixels are also likely to include shortcuts as shortcuts are highly sensitive to perturbations of the input [12].

To examine the accuracy of the above approximations, we trained the digit recognition network of our toy experiment with the proposed SCP and  $\mathcal{L}_{FI}$  (ITSA). As the proposed method is specifically designed for domain generalization, our method can effectively generalize the network

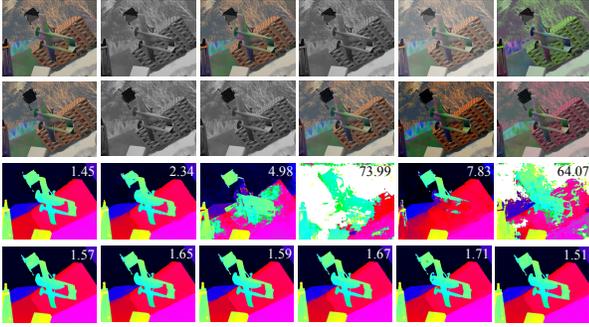


Figure 3. Examples of shortcuts in stereo matching networks. The left and right input images are included in the top two rows. The disparity maps estimated by the baseline PSMNet [6] are included in the third row and ITSA-PSMNet in the bottom row. The performance of the baseline PSMNet deteriorates substantially when the shortcut attributes are distorted or removed from the input stereo images. The corresponding EPE is displayed on the estimated disparity map. Best viewed in color and zoom in for details.

to unseen domains and achieve better performance (4%) than the robust information bottleneck as shown in Tab. 1.

## 4. Experiments

### 4.1. Experimental Settings

**Datasets and Metrics:** *Scene Flow* [25] is a large collection of synthetic stereo images with dense disparity ground truth. It contains FlyingThings3D, Driving and Monkka subsets, and provides 35,454 training and 4,370 testing images. In our experiments, all stereo matching networks are trained on the Scene Flow dataset *only*.

The realistic datasets used in our experiments include *KITTI2012* [11] and *KITTI2015* [26] containing 193 and 200 stereo images of outdoor driving scenes, *Middlebury* [35] containing 15 images of high resolution indoor scenes, and *ETH3D* [36] containing 27 low resolution, greyscale stereo images of both indoor and outdoor scenes. Furthermore, datasets covering different weather conditions provided by the *DrivingStereo* [47] dataset, and night-time provided by *Oxford Robotcar* [24]) were also included to evaluate the robustness of our proposed method. All the above datasets come with with sparse ground truth.

We evaluated the performance of disparity estimation using the D1 error rate (%), with different pixel threshold. The D1 metric computes the percentage of bad pixels (disparity end-point error larger than the threshold) in the left frame. Following the advice of data originators, a threshold of 3 pixels is selected for KITTI and DrivingStereo, 2 pixels for Middlebury, and 1 pixel for ETH3D.

**Baselines & Implementation Details:** We have selected three popular and top-performing stereo matching networks namely PSMNet [6], GwcNet [16] and CFNet [37] as the

Inputs	PSMNet [6]	GwcNet [16]	CFNet [37]
No Augment ( $x$ )	1.38	0.85	1.00
ACJ	13.98	3.13	1.34
GrayScale (L)	37.68	8.41	1.32
GrayScale (R)	9.82	2.25	1.09
SCP	5.84	2.90	2.55

Table 2. Analysis of the effect of data augmentation on the performance of stereo matching networks. All networks are only trained on the Scene Flow training set and the EPE metric is employed for evaluation. The results show that removing shortcut related artefacts (by data augmentation) negatively impact the performance of these networks. In particular, our proposed augmentation can even significantly impact robust methods (e.g. CFNet).

baseline networks for our experiments. We have selected these networks mainly due to the fact that PSMNet and GwcNet are well-studied, and commonly employed as a baseline in many prior works [44, 50, 54]; and CFNet is one of the recently proposed state-of-the-art stereo matching networks. The networks are implemented using PyTorch framework and are trained end-to-end with Adam ( $\beta_1 = 0.9, \beta_2 = 0.999$ ) optimizer. Similar to the original implementations of the selected networks, our data processing includes color normalization and random cropping the input images to size  $H = 256$  and  $W = 512$ . Following the original implementation of CFNet, asymmetric chromatic augmentation and asymmetric occlusion [46] are also employed for data augmentation in CFNet. The maximum disparity for PSMNet and GwcNet is set to 192, and for CFNet is set to 256. All models are trained from scratch for 20 epochs with learning rate set to 0.001 for the first 10 epochs and decreased by half for another 10 epochs. The batch size is set to 12 for training on 2 NVIDIA RTX 8000 Quadro GPUs. The models are trained using **synthetic data only** and directly tested using data from different realistic datasets. For all experiments included in the following sections, the hyper-parameters  $\lambda$  and  $\epsilon$  were set to 0.1 and 0.5 respectively. The hyper-parameter tuning experiments are detailed in the supplementary document. The code of our implementations is available at: <https://github.com/waychin-weiqin/ITSA>

### 4.2. Shortcuts in stereo matching networks

Our hypothesis is that the baseline stereo matching networks naively trained on synthetic data only, learn to exploit common artefacts of synthetic stereo images as shortcut features. These artefacts include (1) consistent local statistics (RGB color features) between the left and right stereo images and (2) over-reliance on local chromaticity features of the reference stereo viewpoint.

To empirically verify the above, we tested three baseline networks trained *only* with synthetic data (i.e. Scene flow), using augmented stereo inputs images. The augmented stereo images were derived from the Scene Flow

SCP	$\mathcal{L}_{FI}$	PSMNet KITTI-2012	GwcNet KITTI-2012	PSMNet KITTI-2015	GwcNet KITTI-2015
$\times$	$\times$	27.4	11.7	29.3	12.8
$\checkmark$	$\times$	8.1	5.3	8.6	5.9
$\checkmark$	$\checkmark$	<b>5.2</b>	<b>4.9</b>	<b>5.8</b>	<b>5.4</b>

Table 3. Ablation results on PSMNet [6] and GwcNet [16]. SCP is the proposed shortcut perturbations and  $\mathcal{L}_{FI}$  is the proposed loss function in Eq. (7). The D1 metric was used for evaluation.

test set using the following strategies: (1) Chromatic Augmentation (e.g. asymmetrical color jittering (ACJ) [46] and gray scaling) and (2) the shortcut-perturbation (SCP, explained in Sec. 3.3.3). If a network has learnt to utilize the transient attributes (related to shortcut), distorting those in the input space will negatively impact its performance. Experimental results, given in Tab. 2, showed that using these augmented images as inputs has substantially worsened the performance of the stereo matching networks.

Interestingly, the SCP images also deteriorate the performance of the best performing robust stereo matching networks such as CFNet [37]. In Sec. 4.5 and 4.5, we show that our method can enhance the robustness of CFNet and significantly improve its performance in unseen realistic environments and anomalous scenarios.

The qualitative results, shown in Fig. 3, demonstrate that the performance of baseline networks (third row) deteriorated significantly when the color features consistency between stereo viewpoints is violated. Moreover, as shown in the fourth column of Fig. 3, removing the chromaticity features from the reference image will causes substantial performance reduction in the baseline networks. In contrast, our proposed method reduces the exploitation of shortcut features and shows better robustness to adverse data augmentation scenarios, without using these shortcut-related knowledge (see last row of Fig. 3).

### 4.3. Ablation Study

This section presents the results of our study on the efficacy of each component of the proposed method. We first trained the baseline networks with the proposed shortcut-perturbation augmentation (SCP) *only*. Next, we trained the baseline networks with both the shortcut-perturbed stereo-images and the proposed loss function  $\mathcal{L}_{FI}$  in Eq. (7).

As shown in Tab. 3, the baseline networks, trained only with synthetic data, perform poorly when tested on KITTI data. The performance improved when shortcut-perturbations (SCP) were used in the training stage for input image augmentations. Further improvement in both baseline networks can be seen when using the proposed method i.e. SCP with the proposed loss function. We have not included CFNet in the ablation study as it is specifically designed for synthetic to real domain generalization.

Methods	KITTI		Middlebury			ETH3D
	2012	2015	Full	Half	Quarter	
HD <sup>3</sup> [49]	23.6	26.5	50.3	37.9	20.3	54.2
PSMNet [6]	27.4	29.3	60.4	29.1	19.6	16.1
GwcNet [16]	11.7	12.8	45.5	18.1	10.9	9.0
CasStereo [15]	11.8	11.9	40.6	-	-	7.8
GANet [52]	10.1	11.7	32.2	20.3	11.2	14.1
MS-PSMNet [3]	14.0	7.8	-	19.8	-	16.8
DSMNet [53]	6.2	6.5	21.8	13.8	<b>8.1</b>	6.2
MS-GCNet [3]	5.5	6.2	-	18.5	-	8.8
CFNet [37]	4.7	5.8	28.2	13.5	9.4	5.8
ITSA-PSMNet	5.2	5.8	28.4	12.7	9.6	9.8
ITSA-GwcNet	4.9	5.4	26.8	11.4	9.3	7.1
ITSA-CFNet	<b>4.2</b>	<b>4.7</b>	<b>20.7</b>	<b>10.4</b>	<b>8.5</b>	<b>5.1</b>
CFNet_RVC [37]	(1.6)	(2.0)	(16.1)	(10.1)	-	(3.7)

Table 4. Synthetic-to-realistic domain generalization evaluation using KITTI, Middlebury and ETH3D training sets. All methods are trained on the Scene Flow dataset and directly tested on the three real datasets. Pixel error rate with different threshold are employed: KITTI 3-pixel, Middlebury 2-pixel and ETH3D 1-pixel.

### 4.4. Synthetic-to-Realistic Domain Generalization Evaluation

In Tab. 4, we compare the synthetic-to-realistic domain generalization performance of our method with the state-of-the-art stereo matching networks [6, 15, 16, 37, 49, 52, 53] on the four realistic datasets. *All networks are trained on the synthetic Scene Flow training set only.* We found that the proposed ITSA substantially improved the domain generalization performance (6.8% – 23.5%) of the selected stereo networks (PSMNet [6] and GwcNet [16]), outperforming the state-of-the-art stereo matching networks in the realistic datasets. The improved networks also outperform DSMNet [53] on the KITTI 2012 [11] and KITTI 2015 [26] datasets, and achieve comparable performance as the CFNet on the Middlebury [35] datasets. In addition, we show that ITSA is even capable of further enhancing the robustness and cross-domain performance of CFNet [37], which was the best performing stereo matching networks in the Robust Vision Challenge 2020. The results of fine-tuned CFNet (CFNet\_RVC) are also included in Tab. 4 to show the upper bound performance of CFNet. Qualitative results comparison of the baseline and ITSA are included in Fig. 4.

### 4.5. Robustness to Anomalous Scenarios

Here, we analyze the robustness to anomalous conditions of a network trained on synthetic data with the proposed ITSA. The anomalous conditions include night-time, foggy and rainy weather conditions. In this comparison, we train the same network twice: (1) pre-train using synthetic data followed by fine-tuning on realistic KITTI 2015 dataset (common strategy), (2) train only using synthetic data with the proposed SCP and  $\mathcal{L}_{FI}$  (ITSA). We also included the pre-trained counterpart of CFNet [37] to illustrate the efficacy of our method in further enhancing the network robustness.

In Tab. 5, we show that the fine-tuned (FT) networks gen-

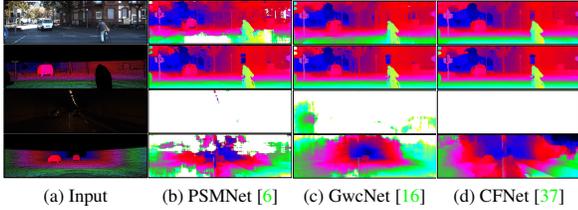


Figure 4. Qualitative results on KITTI 2015 stereo data. For each example, the results of the baseline networks are presented on the top row and the results from our method are included in the bottom row. The corresponding left image and ground truth are included in column (a). Our method can significantly improve the stereo matching performance even in scenario with poor lighting condition. Best viewed in color and zoom in for details.

Models	FT	ITSA	Sun	Cloud	Rain	Fog	Night	Avg
PSMNet [6]	✓	✗	<b>3.94</b>	<b>2.82</b>	11.51	6.50	16.66	8.28
	✗	✓	4.78	3.24	<b>9.43</b>	<b>6.31</b>	<b>8.56</b>	<b>6.46</b>
GwcNet [16]	✓	✗	<b>3.10</b>	<b>2.46</b>	12.34	5.98	25.33	9.84
	✗	✓	4.35	3.31	<b>9.78</b>	<b>5.88</b>	<b>9.41</b>	<b>6.55</b>
CFNet [37]	✗	✗	4.89	4.64	10.74	5.43	16.19	8.38
	✓	✗	<b>1.79</b>	<b>1.65</b>	<b>5.20</b>	<b>1.59</b>	11.56	<b>4.36</b>
	✗	✓	3.42	2.87	5.32	4.32	<b>8.95</b>	4.98

Table 5. Robustness evaluation on anomalous scenarios. Our method (ITSA) consistently enhances the robustness of selected stereo matching networks and outperform the fine-tuned (FT) models in the real-world anomalous scenarios including rainy and foggy weather and night-time. The performances were evaluated using the D1 metric.

erally has better performance when tested on data similar to the KITTI training data (sunny and cloudy). In contrast, our method (ITSA) can substantially improve the robustness and overall performance of the PSMNet [6] and GwcNet [16], without using the real-world data. The overall performance of fine-tuned CFNet is slightly better than its ITSA counterpart. However, as mentioned earlier, the proposed ITSA improves CFNet performance when only using synthetic data for training. The results demonstrate that our method could effectively improve the robustness and performance of existing stereo matching networks, and extends these networks to real-world applications, without using the real data for fine-tuning.

#### 4.6. Extension to Semantic Segmentation

Similar to stereo matching networks, semantic segmentation networks trained on synthetic data also fail to generalize to realistic data [32, 51]. Here, we show that the proposed ITSA can easily be extended to the semantic segmentation task to promote the learning of shortcut-invariant feature and enhance domain generalization. We have selected the commonly employed FCN [23] paired with ResNet-50 as the baseline network. The network was trained on synthetic GTAV [34] dataset only and evaluated on real Cityscapes [8] dataset. The mean intersection over

IBN-Net [29]	ISW [7]	DRPC [51]	ITSA
22.17	7.47 ↑	28.95	7.63 ↑
29.17		36.58	37.42
		32.45	4.97 ↑
		28.71	6.65 ↑
		35.36	

Table 6. Synthetic-to-realistic domain generalization performance comparison on semantic segmentation task. All networks were trained on the GTAV [34] synthetic dataset only and evaluated on the Cityscapes [8] validation set (G → C). The mean intersection over union (mIoU) metric was employed for evaluation.

union (mIoU) metric was employed for performance evaluation. As shown in Tab. 6, the proposed method (ITSA) can also improve the synthetic-to-realistic domain generalization performance of semantic segmentation networks and achieve comparable performance with the existing domain generalization methods (IBN-Net [29], ISW [7] and DRPC [51]). This further demonstrates the effectiveness of our proposed method in promoting shortcut-invariant features and enhancing the performance of domain generalization. The implementation details and qualitative results of our method are included in the supplementary document.

## 5. Limitations

Although the proposed method can significantly improve the performance of stereo matching networks without fine-tuning and even outperform their fine-tuned counterpart when tested in unseen challenging environments (e.g. rain and night-time), its performance remains fragile under extreme conditions (e.g. heavy rain and extreme low light), which may occur in real scenarios. This is reflected by the large errors reported in Tab. 5. By looking at samples with large errors, we noticed that those inaccuracies are largely due to having insufficient light source, lens glare/flares and reflection on specular surfaces (wet grounds).

## 6. Conclusion

In this work, we have presented ITSA: a novel information theory-based approach for domain generalization in stereo matching networks. To address the shortcut learning challenge, we propose to minimize the sensitivity of the extracted feature representations to the input perturbations, measured via the Fisher information. We further proposed an efficient algorithm to optimize the Fisher information objective. Experimental results show that the proposed method consistently promotes the learning of robust and shortcut-invariant features, and substantially enhances the performance of existing stereo matching networks in cross-domain generalization, even outperforming their fine-tuned counterparts in challenging scenarios. We also show that the proposed method can be easily extended for non-geometry based vision problems such as semantic segmentation.

**Acknowledgements** D. Suter acknowledges funding under Australian Research Council grant DP200103448.

## References

- [1] Alex Alemi, Ian Fischer, Josh Dillon, and Kevin Murphy. Deep variational information bottleneck. In *ICLR*, 2017. 2, 4
- [2] Sara Beery, Grant Van Horn, and Pietro Perona. Recognition in terra incognita. In *Proceedings of the European Conference on Computer Vision (ECCV)*, September 2018. 2, 3
- [3] Changjiang Cai, Matteo Poggi, Stefano Mattoccia, and Philippos Mordohai. Matching-space stereo networks for cross-domain generalization. In *2020 International Conference on 3D Vision (3DV)*, pages 364–373, 2020. 7
- [4] Fabio M Carlucci, Antonio D’Innocente, Silvia Bucci, Barbara Caputo, and Tatiana Tommasi. Domain generalization by solving jigsaw puzzles. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, pages 2229–2238, 2019. 2, 3
- [5] Rohan Chabra, Julian Straub, Christopher Sweeney, Richard Newcombe, and Henry Fuchs. Stereodnet: Dilated residual stereonet. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, pages 11786–11795, 2019. 3
- [6] Jia-Ren Chang and Yong-Sheng Chen. Pyramid stereo matching network. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, pages 5410–5418, 2018. 1, 3, 4, 5, 6, 7, 8
- [7] Sungha Choi, Sanghun Jung, Huiwon Yun, Joanne T Kim, Seungryong Kim, and Jaegul Choo. Robustnet: Improving domain generalization in urban-scene segmentation via instance selective whitening. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 11580–11590, 2021. 3, 8
- [8] Marius Cordts, Mohamed Omran, Sebastian Ramos, Timo Rehfeld, Markus Enzweiler, Rodrigo Benenson, Uwe Franke, Stefan Roth, and Bernt Schiele. The cityscapes dataset for semantic urban scene understanding. In *Proc. of the IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, 2016. 8
- [9] Nikolay Dagaev, Brett D Roads, Xiaoliang Luo, Daniel N Barry, Kaustubh R Patil, and Bradley C Love. A too-good-to-be-true prior to reduce shortcut reliance. *arXiv preprint arXiv:2102.06406*, 2021. 2, 3
- [10] Yaroslav Ganin, Evgeniya Ustinova, Hana Ajakan, Pascal Germain, Hugo Larochelle, François Laviolette, Mario Marchand, and Victor Lempitsky. Domain-adversarial training of neural networks. *The journal of machine learning research*, 17(1):2096–2030, 2016. 4
- [11] Andreas Geiger, Philip Lenz, and Raquel Urtasun. Are we ready for autonomous driving? the kitti vision benchmark suite. In *Conference on Computer Vision and Pattern Recognition (CVPR)*, 2012. 6, 7
- [12] Robert Geirhos, Jörn-Henrik Jacobsen, Claudio Michaelis, Richard Zemel, Wieland Brendel, Matthias Bethge, and Felix A Wichmann. Shortcut learning in deep neural networks. *arXiv preprint arXiv:2004.07780*, 2020. 2, 3, 5
- [13] Robert Geirhos, Jörn-Henrik Jacobsen, Claudio Michaelis, Richard Zemel, Wieland Brendel, Matthias Bethge, and Felix A Wichmann. Shortcut learning in deep neural networks. *Nature Machine Intelligence*, 2(11):665–673, 2020. 3
- [14] Robert Geirhos, Patricia Rubisch, Claudio Michaelis, Matthias Bethge, Felix A Wichmann, and Wieland Brendel. Imagenet-trained cnns are biased towards texture; increasing shape bias improves accuracy and robustness. *arXiv preprint arXiv:1811.12231*, 2018. 3
- [15] Xiaodong Gu, Zhiwen Fan, Siyu Zhu, Zuozhuo Dai, Feitong Tan, and Ping Tan. Cascade cost volume for high-resolution multi-view stereo and stereo matching. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 2495–2504, 2020. 7
- [16] Xiaoyang Guo, Kai Yang, Wukui Yang, Xiaogang Wang, and Hongsheng Li. Group-wise correlation stereo network. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, pages 3273–3282, 2019. 1, 3, 4, 5, 6, 7, 8
- [17] Dan Hendrycks and Thomas Dietterich. Benchmarking neural network robustness to common corruptions and perturbations. *Proceedings of the International Conference on Learning Representations*, 2019. 2, 3
- [18] Alex Kendall, Hayk Martirosyan, Saumitro Dasgupta, Peter Henry, Ryan Kennedy, Abraham Bachrach, and Adam Bry. End-to-end learning of geometry and context for deep stereo regression. In *Proceedings of the IEEE International Conference on Computer Vision*, pages 66–75, 2017. 1, 3, 4
- [19] Yann LeCun, Léon Bottou, Yoshua Bengio, and Patrick Haffner. Gradient-based learning applied to document recognition. *Proceedings of the IEEE*, 86(11):2278–2324, 1998. 4
- [20] Haoliang Li, Sinno Jialin Pan, Shiqi Wang, and Alex C Kot. Domain generalization with adversarial feature learning. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, pages 5400–5409, 2018. 3
- [21] Zhengfa Liang, Yiliu Feng, Yulan Guo, Hengzhu Liu, Wei Chen, Linbo Qiao, Li Zhou, and Jianfeng Zhang. Learning for disparity estimation through feature constancy. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, pages 2811–2820, 2018. 3
- [22] Rui Liu, Chengxi Yang, Wenxiu Sun, Xiaogang Wang, and Hongsheng Li. Stereogan: Bridging synthetic-to-real domain gap by joint optimization of domain translation and stereo matching. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 12757–12766, 2020. 1
- [23] Jonathan Long, Evan Shelhamer, and Trevor Darrell. Fully convolutional networks for semantic segmentation. In *Proceedings of the IEEE conference on computer vision and pattern recognition*, pages 3431–3440, 2015. 8
- [24] Will Maddern, Geoff Pascoe, Chris Linegar, and Paul Newman. 1 Year, 1000km: The Oxford RobotCar Dataset. *The International Journal of Robotics Research (IJRR)*, 36(1):3–15, 2017. 1, 6
- [25] Nikolaus Mayer, Eddy Ilg, Philip Hausser, 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*, pages 4040–4048, 2016. 1, 3, 6
- [26] Moritz Menze, Christian Heipke, and Andreas Geiger. Object scene flow. *ISPRS Journal of Photogrammetry and Remote Sensing (JPRS)*, 2018. 1, 6, 7
- [27] Matthias Minderer, Olivier Bachem, Neil Houlsby, and Michael Tschannen. Automatic shortcut removal for self-supervised representation learning. In *International Conference on Machine Learning*, pages 6927–6937. PMLR, 2020. 2, 3
- [28] Guang-Yu Nie, Ming-Ming Cheng, Yun Liu, Zhengfa Liang, Deng-Ping Fan, Yue Liu, and Yongtian Wang. Multi-level context ultra-aggregation for stereo matching. In *Proceedings of the IEEE conference on computer vision and pattern recognition*, pages 3283–3291, 2019. 3
- [29] Xingang Pan, Ping Luo, Jianping Shi, and Xiaoou Tang. Two at once: Enhancing learning and generalization capacities via ibn-net. In *Proceedings of the European Conference on Computer Vision (ECCV)*, pages 464–479, 2018. 8
- [30] Jiahao Pang, Wenxiu Sun, Jimmy SJ Ren, Chengxi Yang, and Qiong Yan. Cascade residual learning: A two-stage convolutional neural network for stereo matching. In *Proceedings of the IEEE International Conference on Computer Vision Workshops*, pages 887–895, 2017. 3
- [31] Ankit Pensia, Varun Jog, and Po-Ling Loh. Extracting robust and accurate features via a robust information bottleneck. *IEEE Journal on Selected Areas in Information Theory*, 1(1):131–144, 2020. 2, 4
- [32] Fengchun Qiao, Long Zhao, and Xi Peng. Learning to learn single domain generalization. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 12556–12565, 2020. 3, 8
- [33] Benjamin Recht, Rebecca Roelofs, Ludwig Schmidt, and Vaishaal Shankar. Do imagenet classifiers generalize to imagenet? In *International Conference on Machine Learning*, pages 5389–5400. PMLR, 2019. 2, 3
- [34] Stephan R Richter, Vibhav Vineet, Stefan Roth, and Vladlen Koltun. Playing for data: Ground truth from computer games. In *European conference on computer vision*, pages 102–118. Springer, 2016. 8
- [35] Daniel Scharstein, Heiko Hirschmüller, York Kitajima, Greg Krathwohl, Nera Nešić, Xi Wang, and Porter Westling. High-resolution stereo datasets with subpixel-accurate ground truth. In *German conference on pattern recognition*, pages 31–42. Springer, 2014. 1, 6, 7
- [36] Thomas Schöps, Johannes L. Schönberger, Silvano Galliani, Torsten Sattler, Konrad Schindler, Marc Pollefeys, and Andreas Geiger. A multi-view stereo benchmark with high-resolution images and multi-camera videos. In *Conference on Computer Vision and Pattern Recognition (CVPR)*, 2017. 6
- [37] Zhelun Shen, Yuchao Dai, and Zhibo Rao. Cfnet: Cascade and fused cost volume for robust stereo matching. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*, pages 13906–13915, June 2021. 2, 3, 4, 6, 7, 8
- [38] Baifeng Shi, Dinghui Zhang, Qi Dai, Zhanxing Zhu, Yadong Mu, and Jingdong Wang. Informative dropout for robust representation learning: A shape-bias perspective. In *International Conference on Machine Learning*, pages 8828–8839. PMLR, 2020. 2, 3
- [39] Yuge Shi, Jeffrey Seely, Philip HS Torr, N Siddharth, Awni Hannun, Nicolas Usunier, and Gabriel Synnaeve. Gradient matching for domain generalization. *arXiv preprint arXiv:2104.09937*, 2021. 5
- [40] Naftali Tishby and Noga Zaslavsky. Deep learning and the information bottleneck principle. In *2015 IEEE Information Theory Workshop (ITW)*, pages 1–5. IEEE, 2015. 2, 4
- [41] Alessio Tonioni, Oscar Rahnema, Thomas Joy, Luigi Di Stefano, Thalaiyasingam Ajanthan, and Philip HS Torr. Learning to adapt for stereo. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, pages 9661–9670, 2019. 1
- [42] Riccardo Volpi, Hongseok Namkoong, Ozan Sener, John C. Duchi, Vittorio Murino, and Silvio Savarese. Generalizing to unseen domains via adversarial data augmentation. In Samy Bengio, Hanna M. Wallach, Hugo Larochelle, Kristen Grauman, Nicolò Cesa-Bianchi, and Roman Garnett, editors, *Advances in Neural Information Processing Systems 31: Annual Conference on Neural Information Processing Systems 2018, NeurIPS 2018, December 3-8, 2018, Montréal, Canada*, pages 5339–5349, 2018. 3
- [43] Haohan Wang, Zexue He, Zachary L. Lipton, and Eric P. Xing. Learning robust representations by projecting superficial statistics out. In *International Conference on Learning Representations*, 2019. 3
- [44] Yan Wang, Wei-Lun Chao, Divyansh Garg, Bharath Hariharan, Mark Campbell, and Kilian Q Weinberger. Pseudo-lidar from visual depth estimation: Bridging the gap in 3d object detection for autonomous driving. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, pages 8445–8453, 2019. 6
- [45] Haofei Xu and Juyong Zhang. Aanet: Adaptive aggregation network for efficient stereo matching. *arXiv preprint arXiv:2004.09548*, 2020. 1, 3
- [46] Gengshan Yang, Joshua Manela, Michael Happold, and Deva Ramanan. Hierarchical deep stereo matching on high-resolution images. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, pages 5515–5524, 2019. 6, 7
- [47] Guorun Yang, Xiao Song, Chaoqin Huang, Zhidong Deng, Jianping Shi, and Bolei Zhou. Drivingstereo: A large-scale dataset for stereo matching in autonomous driving scenarios. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 899–908, 2019. 1, 6
- [48] Guorun Yang, Hengshuang Zhao, Jianping Shi, Zhidong Deng, and Jiaya Jia. Segstereo: Exploiting semantic information for disparity estimation. In *Proceedings of the European Conference on Computer Vision (ECCV)*, pages 636–651, 2018. 3
- [49] Zhichao Yin, Trevor Darrell, and Fisher Yu. Hierarchical discrete distribution decomposition for match density estimation. In *Proceedings of the IEEE/CVF Conference on Com-*

- puter Vision and Pattern Recognition*, pages 6044–6053, 2019. [7](#)
- [50] Yurong You, Yan Wang, Wei-Lun Chao, Divyansh Garg, Geoff Pleiss, Bharath Hariharan, Mark Campbell, and Kilian Q Weinberger. Pseudo-lidar++: Accurate depth for 3d object detection in autonomous driving. *arXiv preprint arXiv:1906.06310*, 2019. [6](#)
- [51] Xiangyu Yue, Yang Zhang, Sicheng Zhao, Alberto Sangiovanni-Vincentelli, Kurt Keutzer, and Boqing Gong. Domain randomization and pyramid consistency: Simulation-to-real generalization without accessing target domain data. In *Proceedings of the IEEE/CVF International Conference on Computer Vision*, pages 2100–2110, 2019. [8](#)
- [52] Feihu Zhang, Victor Prisacariu, Ruigang Yang, and Philip HS Torr. Ga-net: Guided aggregation net for end-to-end stereo matching. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, pages 185–194, 2019. [1](#), [3](#), [5](#), [7](#)
- [53] Feihu Zhang, Xiaojuan Qi, Ruigang Yang, Victor Prisacariu, Benjamin Wah, and Philip Torr. Domain-invariant stereo matching networks. In *European Conference on Computer Vision*, pages 420–439. Springer, 2020. [2](#), [3](#), [5](#), [7](#)
- [54] Youmin Zhang, Yimin Chen, Xiao Bai, Jun Zhou, Kun Yu, Zhiwei Li, and Kuiyuan Yang. Adaptive unimodal cost volume filtering for deep stereo matching. *arXiv preprint arXiv:1909.03751*, 2019. [6](#)