

Unsupervised Hierarchical Domain Adaptation for Adverse Weather Optical Flow

Hanyu Zhou¹, Yi Chang^{1*}, Gang Chen², Luxin Yan¹,

¹ National Key Laboratory of Science and Technology on Multispectral Information Processing, School of Artificial Intelligence and Automation, Huazhong University of Science and Technology

² School of Computer Science and Engineering, Sun Yat-sen University
hyzhou@hust.edu.cn, yichang@hust.edu.cn, cheng83@mail.sysu.edu.cn, yanluxin@hust.edu.cn

Abstract

Optical flow estimation has made great progress, but usually suffers from degradation under adverse weather. Although semi/full-supervised methods have made good attempts, the domain shift between the synthetic and real adverse weather images would deteriorate their performance. To alleviate this issue, our start point is to unsupervisedly transfer the knowledge from source clean domain to target degraded domain. Our key insight is that adverse weather does not change the intrinsic optical flow of the scene, but causes a significant difference for the warp error between clean and degraded images. In this work, we propose the first unsupervised framework for adverse weather optical flow via hierarchical motion-boundary adaptation. Specifically, we first employ image translation to construct the transformation relationship between clean and degraded domains. In motion adaptation, we utilize the flow consistency knowledge to align the cross-domain optical flows into a motion-invariance common space, where the optical flow from clean weather is used as the guidance-knowledge to obtain a preliminary optical flow for adverse weather. Furthermore, we leverage the warp error inconsistency which measures the motion misalignment of the boundary between the clean and degraded domains, and propose a joint intra- and inter-scene boundary contrastive adaptation to refine the motion boundary. The hierarchical motion and boundary adaptation jointly promotes optical flow in a unified framework. Extensive quantitative and qualitative experiments have been performed to verify the superiority of the proposed method.

Introduction

Optical flow is an important task in computer vision, which aims to estimate per-pixel motion in a video sequence and has been widely applied for autonomous driving (Sun et al. 2020), object tracking (Fan et al. 2019). Existing optical flow methods have achieved great progress under good weather conditions in Fig. 1 (a), but may suffer from challenges under adverse weather in Fig. 1 (b), such as rain and fog. The main reason is that degradation factors have unexpectedly violated the brightness and gradient constancy assumptions. The goal of this work is to estimate optical flow under adverse weather, such as rain, fog, and snow as shown in Fig. 1 (d).

*Corresponding Author

Copyright © 2023, Association for the Advancement of Artificial Intelligence (www.aaai.org). All rights reserved.

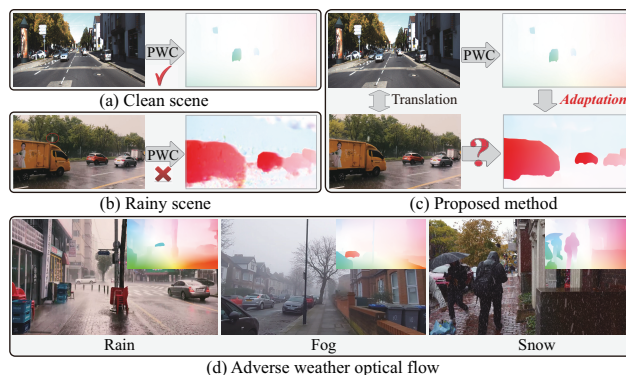


Figure 1: (a) and (b) Classical PWC-Net obtains good flow on the clean scene yet fails to produce satisfactory optical flow under the real degraded scene. (c) The proposed method HMBA-FlowNet bypasses the difficulty of directly estimating adverse weather optical flow, and proposes the first unsupervised domain adaptation method to transfer the knowledge from clean domain. (d) Optical flow could be estimated by the proposed method under various real adverse scenes.

The existing methods have made attempts to get rid of the negative influence of adverse weather for optical flow estimation. The optimization-based optical estimation method (Li, Tan, and Cheong 2018) focuses on specific weather conditions rainy or foggy with sophisticated hand-crafted prior, which makes them less suitable for various adverse weather. The learning-based methods require the ground truth labels of clean image, optical flow (Li et al. 2019) or auxiliary gyroscope information (Li, Luo, and Liu 2021), which are not easy to be acquired in real-world scenes. Although previous learning-based methods have achieved great progress, the domain shift between the synthetic and real adverse weather images restricts their performance in real adverse scenes.

To overcome the domain shift issue, our start point is to formulate this task as unsupervised domain adaptation where we can transfer motion knowledge from source clean domain to target degraded domain in Fig. 1 (c). To our knowledge, we are the first to handle adverse weather optical flow in an unsupervised framework. *Our key insight is that the inherent motion is invariant but the warp error is variant between clean and degraded images.* This intuitive observation is rea-

sonable. On one hand, the optical flow of the same scene should be essentially consistent no matter what weather conditions are. On the other hand, the warp error measuring the alignment of the image structure of two adjacent frames is very small for the clean image while the warp error for the degraded image would be enormous due to degradation. The motion consistency and boundary inconsistency between both domains provide us the powerful knowledge for adaptation.

In this work, we first formulate the degraded optical flow estimation into an ill-posed problem via *maximum-a-posterior*. The mathematical formulation offers us a heuristic guidance on how to principally design the network and unsupervised loss. Consequently, we propose an unsupervised hierarchical motion-boundary adaptation (HMBA-FlowNet) method for optical flow under adverse weather, which consists of three main parts: clean and degraded image translation, motion adaptation, and boundary adaptation.

Specifically, we first learn the unpaired image translation between the clean and degraded domains. In the motion adaptation step, we encode the cross-domain clean-degraded images into a motion-invariance common feature space, and then decoder them into the same optical flow via the consistency loss. The motion adaptation could capture the cross-domain flow consistency relationship and transfer the motion information from the clean domain to the degraded domain, thus encouraging preliminary optical flow results for the real degraded images. In the boundary adaptation step, we propose a joint intra- and inter-scene boundary contrastive adaptation in the warp error space. The intra-scene boundary adaptation is to minimize the discrepancy between different patches in the same scene and maximize the distance between the clean and degraded domains, meanwhile, the inter-scene boundary contrastive adaptation encourages the feature manifolds of different scenes in the same domain to be closer to each other. Both intra- and inter-scene boundary contrastive adaptation boost the model to effectively distinguish the boundary from the warp error caused by degradation, thus refining the inherent motion boundary in the real degraded domain. We summarize the main contributions as follows:

- We design a joint motion-consistency and boundary-inconsistency alignment mechanism for the adaptation, and propose an unsupervised hierarchical motion-boundary adaptation framework for adverse weather optical flow. This is a pioneer exploration to transfer the reliable knowledge from clean domain to adverse weather domain from the unsupervised domain adaptation paradigm.
- We propose a motion consistency adaptation (MCA) to learn the cross-domain motion-invariance feature which acts as a fidelity to ensure the preliminary motion, and a boundary contrastive adaptation (BCA) to effectively distinguish the motion boundary from the warp error caused by degradation which works as a regularization to refine the inherent motion boundary. The joint adaptation mechanism benefits obtaining better optical flow.
- The proposed method consistently outperforms the competing methods on real datasets by a large margin. Moreover, the proposed method generalizes well for different adverse weather conditions, such as rain, fog, and snow.

Related Work

Optical Flow under Adverse Weather. In recent years, the learning-based optical flow approaches (Dosovitskiy et al. 2015; Ren et al. 2017; Sun et al. 2018; Luo et al. 2021; Goodfellow et al. 2014; Liu et al. 2019a) have made great progress. Although these methods have achieved satisfactory results in clean scenes, they would suffer from degradation under adverse weather. An intuitive solution to this challenging task is to perform the image deraining (Fu et al. 2017a; Yang et al. 2017; Zhang and Patel 2018) or dehazing (Ren et al. 2018; Liu et al. 2019c) with subsequent optical flow estimation. However, existing restoration methods are not designed for optical flow and possible over-smoothness or residual artifacts may result in a negative influence.

RobustFlow (Li, Tan, and Cheong 2018) was the first optical flow estimation method in rainy scenes by constructing the handcrafted residue channel and its colored version prior that is invariant to rain streaks. Further, to improve feature representation, Li *et al.* (Li et al. 2019) proposed to automatically learn the rain- and veiling-invariant features in a supervised manner. Given that it is difficult to collect ground truth for real scenes, Yan *et al.* (Yan, Sharma, and Tan 2020) presented a semi-supervised method DenseFogFlow, and its key idea is to model the flow consistency between clean images and corresponding rendered foggy images. GyroFlow (Li, Luo, and Liu 2021) resorted to a hardware scheme to handle adverse weather, and innovatively utilized the gyroscope to obtain ego-motion labels of the camera for weakly-supervised optical flow estimation. Compared with previous methods, the proposed method is the first to handle the optical flow under adverse weather from an unsupervised perspective which generalizes well for various real adverse weather images.

Domain Adaptation. Domain adaptation (Chen et al. 2021a) aims to effectively tackle the distribution discrepancy between the source and target domains, which has been widely applied in classification (Chen et al. 2021b), detection (Xie et al. 2021) and semantic (Gao et al. 2022; Ma et al. 2022). In fact, the previous methods have utilized different invariance properties (Li et al. 2019; Yan, Sharma, and Tan 2020) between the clean and specific degraded images, which aim to transfer motion knowledge from the source clean domain to the target adverse weather domain. Compared with previous methods, the proposed method not only explores the motion-consistency for preliminary motion estimation, but also excavates the boundary-inconsistency prior to better refine the motion boundary. Our hierarchical motion-boundary adaptation empowers us with better knowledge transfer ability in a unified coarse-to-fine optical flow framework.

Hierarchical Adaptation for Optical Flow

Problem and Motivation

The goal of this work is to estimate optical flow under real adverse weather scenes. Since adverse weather degradation breaks the brightness constancy and gradient constancy assumptions of optical flow, most of the existing optical flow methods usually suffer from degradation under real adverse weather scenes. To alleviate this issue, we start with the mathematical formulation for the degraded optical flow to reveal

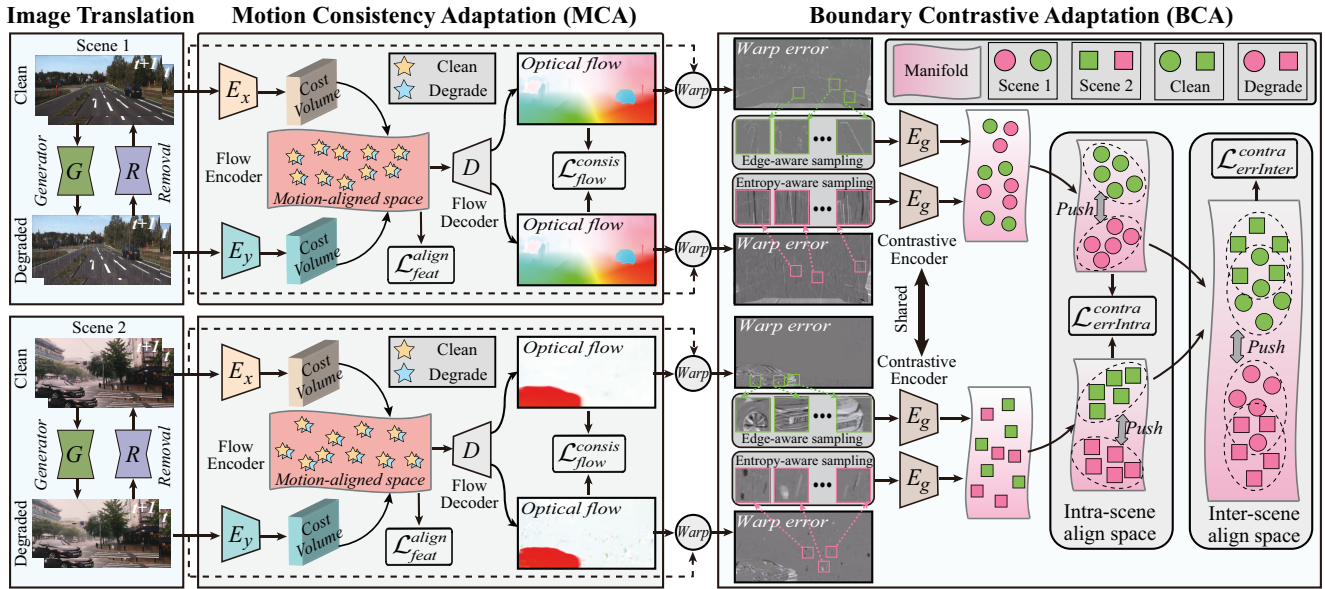


Figure 2: The overall architecture mainly contains image translation, motion-boundary adaptation. We first employ the image translation between clean and degraded domains. In motion adaptation, we align the cross-domain motion features into a common space, and then enforce the optical flow consistency loss to transfer motion knowledge. In boundary adaptation, we propose intra- and inter-scene boundary contrastive adaptation to model the warp error inconsistency to further improve motion boundary.

the optical flow consistency and warp error inconsistency priors. And then, we propose an unsupervised hierarchical motion-boundary domain adaptation for optical flow under adverse weather scenes in Fig. 2. We first employ the unpaired image translation between the clean domain and the degraded domain. Consequently, we transfer motion knowledge from the clean domain to the degraded domain via motion adaptation, and align the distribution of motion boundary of intra- and inter-scene via boundary adaptation.

Mathematical Formulation

Given an adverse weather image Y , we simply assume the degradation as the linear additive decomposition model:

$$Y = X + D, \quad (1)$$

where X , Y , and D denote clean image, degraded image, and adverse component, respectively. The motion is usually defined as the linear additive relationship of both the foreground object and background scene (Lv et al. 2018). We assume that clean image X and adverse component D are independent from each other. For example, the motion of the rain/snow is independent from the foreground object and background. Thus, we enforce optical flow operator F on Eq. (1):

$$F(Y) = F(X) + F(D), \quad (2)$$

That is to say, the optical flow of the degraded image can be inferred from that of the clean image. Thus, we can formulate the optical flow under adverse weather as an ill-posed inverse problem via *maximum-a-posterior* (Murphy 2012):

$$\min_{X, R, F} \underbrace{\gamma \|F(Y) - F(X)\|_1}_{\text{flow consistent}} + \underbrace{\delta P_x(F(X))}_{\text{clean flow}} + \underbrace{\rho P_d(F(D))}_{\text{degraded flow}} \quad s.t. \quad X = R(Y), \quad (3)$$

where R is the image translation, the first term is flow consistency between clean and degraded images, the second and

third terms P_x , P_d denote the optical flow prior knowledge of clean image and degradation, respectively. We further relax the constrained problem Eq. (3) into an unconstrained one:

$$\min_{X, R, F} \underbrace{\alpha \|X - R(Y)\|_1}_{\text{image translate}} + \underbrace{\beta P(X)}_{\text{image}} + \underbrace{\gamma \|F(Y) - F(X)\|_1}_{\text{flow consistent}} + \underbrace{\delta P_x(F(X))}_{\text{clean flow}} + \underbrace{\rho P_d(F(D))}_{\text{degraded flow}}, \quad (4)$$

where the first two terms are to translate the image between clean and degraded domains, and the intention of the third and fourth terms is to estimate the optical flow from clean domain and construct motion adaptation between both the domains, and the fifth term is the flow prior about the degradation. $\alpha, \beta, \gamma, \delta, \rho$ are balance hyper-parameters. Next, we will describe how we design each term and overall network to mimic the optimization function Eq. (4).

Unpaired Image Translation

Our key idea is to transfer motion knowledge from clean to degraded domain. Thus, the image translation between clean and degraded domains is essential to build the relationship between them. Intuitively, we can employ a classical generative network (Isola et al. 2017) to mimic the first two terms in Eq. (4). In this work, instead of using the conventional GAN, we employ the unpaired CycleGAN (Zhu et al. 2017), which takes two unpaired scenes as input for robust image translation and abundant knowledge transfer. Thus, given two adjacent clean images $X = [X_1, X_2]$ and degraded images $Y = [Y_1, Y_2]$, the image translation loss can be expressed as:

$$\mathcal{L}_{img}^{tran} = \|R(G(X)) - X\|_1 + \|G(R(Y)) - Y\|_1, \quad (5)$$

where R, G is the degraded-to-clean and clean-to-degraded image translation networks, respectively. We follow (Zhu

et al. 2017) to express the image prior with adversarial loss:

$$\begin{aligned} \mathcal{L}_{img}^{prior} = & \mathbb{E}_x[\log D_x(\mathbf{X})] + \mathbb{E}_y[\log(1 - D_x(R(\mathbf{Y})))] \\ & + \mathbb{E}_y[\log D_y(\mathbf{Y})] + \mathbb{E}_x[\log(1 - D_y(G(\mathbf{X})))] \end{aligned} \quad (6)$$

where D_x and D_y are discriminators to make the images lying on the natural clean and degraded manifolds, respectively.

Motion Consistency Adaptation

According to Eq. (4), the third and fourth terms enlighten us to bypass the difficulty to solve this challenging problem in the degraded domain, but instead, we can resort to the motion knowledge in the clean domain. The physical meaning of the two terms is to transfer optical flow from clean domain to degraded domain via the motion consistency. We employ the simple yet effective PWC-Net (Sun et al. 2018) as our backbone to learn optical flow. We first take two encoders E_x and E_y after the cost volume layer to align the cross-domain inputs into a motion-invariance common feature space. It is worth noting that, although we share the same architecture for optical flow of clean and degraded domains, their weights of encoders have better not be shared (please refer to the discussion). We enforce the cost volume feature consistency loss to align the correlation features between both domains:

$$\begin{aligned} \mathcal{L}_{feat}^{align} = & \|C(E_y(G(\mathbf{X}))) - C(E_x(\mathbf{X}))\|_1 \\ & + \|C(E_y(\mathbf{Y})) - C(E_x(R(\mathbf{Y})))\|_1, \end{aligned} \quad (7)$$

where $C(\cdot)$ is the cost volume layer that extracts the most discriminative motion features. The motion-aligned embedding features could provide a homogeneous adaptation preparation for the image-level flow adaptation. Consequently, we enforce one shared flow decoder D to estimate optical flows of both domains from the motion-aligned feature space with the optical flow consistency loss as follows:

$$\begin{aligned} \mathcal{L}_{flow}^{consis} = & \|D(C(E_y(G(\mathbf{X})))) - D(C(E_x(\mathbf{X})))\|_1 \\ & + \|D(C(E_x(R(\mathbf{Y})))) - D(C(E_y(\mathbf{Y})))\|_1. \end{aligned} \quad (8)$$

Eq. (7) and Eq. (8) jointly achieve the flow consistency adaptation of the third term in Eq. (4). As for the clean optical flow prior, namely the fourth term in Eq. (4), the brightness and gradient constancy assumptions still hold for clean images. Thus, we employ the photometric loss (Yu, Harley, and Derpanis 2016) to compute the brightness difference between two adjacent images $\mathbf{X}_1, \mathbf{X}_2$ for non-occluded pixels:

$$\begin{aligned} \mathcal{L}_{flow}^{prior} = & \sum \psi(\mathbf{X}_1 - \mathbf{w}(\mathbf{X}_2)) \odot (1 - O_f) / \sum (1 - O_f) \\ & + \sum \psi(\mathbf{X}_2 - \mathbf{w}(\mathbf{X}_1)) \odot (1 - O_b) / \sum (1 - O_b), \end{aligned} \quad (9)$$

where \mathbf{w} is warping operator, ψ is a sparse L_p norm ($p = 0.4$), O_f and O_b are the forward and backward occlusion mask by checking forward-backward consistency, and \odot is a matrix element-wise multiplication. On one hand, the photometric loss $\mathcal{L}_{flow}^{prior}$ facilitates the network to learn the optical flow from the clean domain. On the other hand, the consistency losses $\mathcal{L}_{flow}^{consis}$, $\mathcal{L}_{feat}^{align}$ benefit the network to transfer the motion from the clean domain to the degraded domain.

Boundary Contrastive Adaptation

The motion adaptation MCA could offer a preliminary optical flow under adverse weather. However, the motion bound-

ary may be over-smooth and suffers from the perturbation caused by degradation. To further refine the motion boundary, we propose a boundary contrastive adaptation to distinguish the image motion boundary from that of the erroneous caused by the adverse degradation, such as rain streaks.

In Eq. (4), the $P_d(F(D))$ is the prior knowledge about the degradation flow $F(D)$ in a transform domain, which could better facilitate us to decouple the optical flow $F(D)$ from $F(Y)$. Thus, the key is how to properly define the loss about $P_d(F(D))$ and design the corresponding network. Since warp error measures the error of motion boundary, we assume the sparsity transform P_d as the warp error transform:

$$\mathbf{w}(i) = \mathbf{I}_1(i) - \mathbf{I}_2(i + F(i)), \quad (10)$$

where \mathbf{w} is warp error operator, $\mathbf{I}_1, \mathbf{I}_2$ are adjacent frames, F is optical flow, and i is pixel index. Thus the $P_d(F(D))$ can be expressed as $S(\mathbf{w}_d)$ where $S(\cdot)$ is a sparsity function, and \mathbf{w}_d is the warp error about the degradation D .

According to the optical flow decomposition model Eq. (2), we can further exploit the relationship among warp errors of clean images, degraded images, and degradation as:

$$\begin{aligned} \mathbf{w}_y(i) - \mathbf{w}_x(i) = & (Y_1(i) - X_1(i)) - (Y_2(i + F(i)) - X_2(i + F(i))) \\ = & D_1(i) - D_2(i + F(i)) = \mathbf{w}_d(i), \end{aligned} \quad (11)$$

where $\mathbf{w}_x, \mathbf{w}_y$ and \mathbf{w}_d mean the warp errors of each component. Eq. (11) supports our observation that the warp errors between clean domain and degraded domain are inconsistent with \mathbf{w}_d caused by adverse weather components. Therefore, instead of directly modeling \mathbf{w}_d , we can indirectly model the warp error relationship between clean domain and degraded domain, namely $(\mathbf{w}_y - \mathbf{w}_x)$. We take $S(\cdot)$ as exponential function $\exp(\cdot)$, thus $P_d(F(D))$ can be deduced as:

$$P_d(F(D)) = \exp(\mathbf{w}_y - \mathbf{w}_x) = \exp(\mathbf{w}_y) \setminus \exp(\mathbf{w}_x). \quad (12)$$

That is to say, $P_d(F(D))$ is to model the boundary inconsistency between clean and degraded domains. Interestingly, we observe that Eq. (12) shares similar formulation with the contrastive learning (Chen et al. 2020). Both of them are to measure the exclusive relationship between two domains, motivating us to employ the contrastive for boundary adaptation.

In Fig. 2, there are four warp error maps including both the clean and degraded from two different scenes. We design the intra- and inter-scene boundary contrastive adaptation. The former intra-scene adaptation is to learn discriminative motion boundary representation, while the latter inter-scene adaptation is to align boundary manifold variation for different scenes. Both the intra- and inter-scene boundary contrastive adaptations benefit us to better distinguish the clean motion boundary from the degradation motion boundary.

In contrastive adaptation, we propose an edge-aware sampling strategy by extracting the salient edges (Dollár and Zitnick 2013) [green patches in Fig. 2] from the clean domain as the positive samples, in order to get rid of the most meaningless region without zero values. As for the negative samples, we propose a local entropy-aware sampling strategy by extracting the salient edges [pink patches in Fig. 2] from the degraded domain. We sort the local entropy (degree of degradation) in descending order and choose the top N patches. Note that, the negative patches would be sampled with different locations from the positives, so as to exclude




Figure 3: Visual comparison of optical flows on real GOF dataset.

Method		RobustFlow	Selflow	Selflow			ARFlow	ARFlow			HMBA-Flow
				DDN+	JORDER-E+	RLNet+		DDN+	JORDER-E+	RLNet+	
LR-KITTI	EPE	22.42	9.07	7.67	6.98	7.13	7.10	7.44	7.65	7.90	5.86
	F1-all	80.13%	44.87%	44.42%	42.08%	42.14%	43.47%	44.68%	43.86%	44.95%	33.61%
HR-KITTI	EPE	23.40	10.78	7.82	7.11	7.04	8.89	7.74	6.95	7.77	6.47
	F1-all	81.12%	52.18%	45.43%	41.78%	41.35%	48.98%	45.05%	43.55%	45.89%	38.19%

Table 1: Quantitative results on synthetic Heavy (HR-KITTI) and Light Rain-KITTI2015 (LR-KITTI) datasets.

the clean image boundary. The advantage of the proposed sampling strategy over the conventional random sampling strategy is analyzed in the discussion.

Therefore, we denoted $f_{P_j^x} = E_G(w_{x_j})$, $f_{P_j^y} = E_G(w_{y_j})$ and $f_{N_i} = E_G(w_{y_i})$ for the output positive and negative features. In the boundary adaptation step, the intra-scene contrastive adaptation loss is as follows:

$$\mathcal{L}_{errIntra}^{contra} = \sum_{j=1}^N \frac{\exp(f_{P_j^x} \cdot f_{P_j^y} / \tau)}{\exp(f_{P_j^x} \cdot f_{P_j^y} / \tau) + \sum_{i=1}^N \exp(f_{N_i} \cdot f_{P_j^x} / \tau)}, \quad (13)$$

The intra-scene contrastive loss constrains the clean warp error patches w_{x_j} at location j to be positive with the corresponding degraded patches w_{y_j} in comparison to other warp error patches w_{y_i} , so as to force motion boundary of the degraded domain close to that of the clean domain in the same scene and get rid of degradation. We further enforce the inter-scene contrastive adaptation loss to align the distribution of motion boundary in the clean domain of different scenes:

$$\mathcal{L}_{errInter}^{contra} = \frac{1}{N} \sum_{k=1}^N \sum_{j=1}^N \frac{\exp(f_{P_j^x} \cdot f_{P_k^y} / \tau)}{\exp(f_{P_j^x} \cdot f_{P_k^y} / \tau) + \sum_{i=1}^{2N} \exp(f_{N_i} \cdot f_{P_j^x} / \tau)}, \quad (14)$$

where N is the positive/negative sample numbers, τ denotes the scale parameter. The total loss of our framework strictly mimics the optimization function Eq. (4) as follows:

$$\mathcal{L} = \alpha \mathcal{L}_{img}^{tran} + \beta \mathcal{L}_{img}^{prior} + \gamma_1 \mathcal{L}_{feat}^{align} + \gamma_2 \mathcal{L}_{flow}^{consis} + \delta \mathcal{L}_{flow}^{prior} + \rho_1 \mathcal{L}_{errIntra}^{contra} + \rho_2 \mathcal{L}_{errInter}^{contra}, \quad (15)$$

Implementation Detail

We empirically set the parameters $\{\alpha, \beta, \gamma_1, \gamma_2, \delta, \rho_1, \rho_2\} = \{1, 1, 1, 1, 0.1, 0.1\}$. Our method is implemented on the Tensorflow platform with two NVIDIA RTX 2080Ti GPUs within three days. Note that, the proposed network contains image translation generators and discriminators, two unshared optical flow encoders and one shared decoder, and one shared contrastive encoder. We first train the image translation generators and discriminators via \mathcal{L}_{img}^{tran} and $\mathcal{L}_{img}^{prior}$. Then we update the optical flow encoders and decoder via

Method	RobustFlow	PWC-Net	Selflow	GyroFlow	Ours	
Rain	EPE	7.10	2.85	1.78	1.07	0.91
	F1-all	53.35%	21.39%	16.21%	9.94%	8.07%
Fog	EPE	12.25	5.99	1.97	0.95	0.84
	F1-all	80.93%	41.02%	18.35%	9.13%	7.46%

Table 2: Quantitative results on real GOF dataset.

$\mathcal{L}_{feat}^{align}$, $\mathcal{L}_{flow}^{consis}$ and $\mathcal{L}_{flow}^{prior}$ with 200k iterators and 0.0002 learning rate, where we initialize the optical flow network of clean domain with clean KITTI2015 dataset (Menze and Geiger 2015) via $\mathcal{L}_{flow}^{prior}$ only. After that, we learn the whole network with the contrastive losses $\mathcal{L}_{errIntra}^{contra}$ and $\mathcal{L}_{errInter}^{contra}$. At the test stage, the final model only needs the optical flow encoder E_y and decoder D which takes real degraded images as input and outputs corresponding optical flow.

Experiments

Experiment Setup

Datasets. We validate the performance on one synthetic and two real degraded datasets. We simulate the synthetic rain with different densities of rain (e.g., heavy rain and light rain) on KITTI2015 by the software *Adobe After Effects*. GOF (Li, Luo, and Liu 2021) is a real dataset with 1000 images for training and 120 images for testing. Another real adverse weather dataset is collected from the *Youtube*. We select 1200 adverse weather images including e.g., rain, fog, and snow for training and 200 images for testing. We choose EPE (Dosovitskiy et al. 2015) and F1-all (Menze and Geiger 2015) for the quantitative evaluation.

Comparison Methods. Two competing methods weakly-supervised GyroFlow (Li, Luo, and Liu 2021) and optimization-based RobustFlow (Li, Tan, and Cheong 2018) are chosen which are designed for rainy scene optical flow.



Figure 4: Comparison of optical flows on real adverse weather dataset. HMBA-FlowNet can estimate motion-smooth and boundary-sharpen optical flow under adverse weather. Moreover, it can generalize well for real various adverse weather.

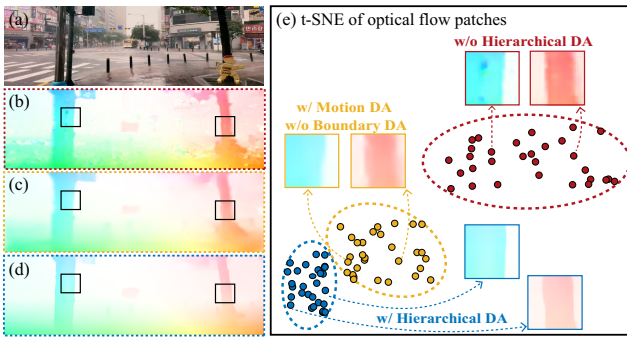


Figure 5: Effectiveness of hierarchical adaptation architecture. (a) Rainy image. (b)-(d) Optical flows without DA, with MCA only, with hierarchical MCA and BCA, respectively. (e) The t-SNE visualization of each adaptation strategy.

Moreover, we select several state-of-the-art supervised (PWC-Net (Sun et al. 2018)) and unsupervised (Selflow (Liu et al. 2019b) and ARFlow (Liu et al. 2020)) clean image-based optical flow approaches. For fair comparison, the supervised methods are only trained on the synthetic dataset with optical flow labels. The unsupervised methods only have access to the degraded images for training. As for the synthetic dataset, we design two different training strategies for competing methods. The first one is that we directly train the comparison methods on rainy images. The second one is to perform the deraining first via derain approaches (*e.g.*, DDN (Fu et al. 2017b), JORDER-E (Yang et al. 2020) and RLNet (Chen and Li 2021)), and then we train the comparison methods on the deraining results (named as DDN+/JORDER-E+/RLNet+).

Experiments on Synthetic Images

In Table 1, we show the quantitative comparison on the synthetic heavy and light Rain-KITTI2015 dataset. We have two key observations. First, the proposed HMBA-FlowNet is significantly better than that of the unsupervised counterparts under different rainy conditions. Second, the pre-processing procedure deraining may not be always positive to the final optical flow. As for the light rain condition, the deraining result may be even harmful because the image deraining would

$\mathcal{L}_{flow}^{consis}$	$\mathcal{L}_{feat}^{align}$	$\mathcal{L}_{errIntra}^{contra}$	$\mathcal{L}_{errInter}^{contra}$	LR-KITTI	HR-KITTI
×	×	×	×	9.14	10.61
✓	×	×	×	6.71	6.89
✓	✓	×	×	6.65	6.89
✓	✓	✓	×	6.52	6.67
✓	✓	×	✓	6.01	6.54
✓	✓	✓	✓	5.86	6.47

Table 3: Ablation study on motion-boundary adaptation loss.

unexpectedly cause damage to image details and outweigh its removal benefits. On the contrary, the proposed method could well handle both the light and heavy rain conditions, since the motion-consistency and boundary-inconsistency knowledge would greatly relieve the need for the clean image.

Experiments on Real Images

In Fig. 3 and 4, we show the visual comparison results on real GOF and collected datasets, respectively. The supervised PWC-Net contains obvious erroneous, due to the domain gap between synthetic training and real test images. The unsupervised methods Selflow and ARFlow cannot work well, in which there are artifacts because degradation has violated the basic assumptions of optical flow. The hardware-assisted GyroFlow can predict the accurate motion of the background, but the boundary of the foreground moving object is not clear in Fig. 3(e), since it heavily relies on the ego-motion of the camera captured by the gyroscope data yet less effective for the independent foreground object motion. The proposed HMBA-FlowNet can remove erroneous outliers and obtain satisfactory results under real adverse weather in Fig. 3 and 4(f). In Table 2, the quantitative results on GOF dataset further verify the superiority of the proposed method.

Ablation Study

Effectiveness of Hierarchical Adaptation Architecture.

To illustrate the effectiveness of the hierarchical adaptation architecture, in Fig. 5, we show the optical flow estimation of different architectures and visualize their low-dimensional distributions via t-SNE. In Fig. 5(b), we can observe that there are obvious artifacts in the optical flow without hierarchical

Domain	Clean	Degraded	LR-KITTI	HR-KITTI
Sampling strategy	Random	Random	6.59	5.74
	Random	Entropy-aware	6.43	6.72
	Edge-aware	Random	5.98	6.52
	Edge-aware	Entropy-aware	5.86	6.47
Weight strategy	Optical flow encoder		LR-KITTI	HR-KITTI
	Shared		7.32	7.60
	Unshared (ours)		5.86	6.47
Backbone strategy	Image translation network		LR-KITTI	HR-KITTI
	PatchGAN		6.67	7.03
	CycleGAN (ours)		5.86	6.47

Table 4: Discussion on strategies of different modules.

adaptation. With motion adaptation only [Fig. 5(c)], most of the outliers caused by degradation are removed. With the additional boundary adaptation [Fig. 5(d)], the flow boundary is significantly clearer. We can conclude that the hierarchical domain adaptation could improve adverse weather optical flow. Moreover, in Fig. 5(e), we visualize their corresponding t-SNE distributions. The red, yellow, and blue denote the without hierarchical adaptation, with motion adaption yet without boundary adaptation, and with hierarchical motion and boundary adaptation. The red circle is scattered, the yellow circle is gradually focused, and the blue circle is most concentrated, illustrating that hierarchical adaptation could greatly improve optical flow discriminative representation.

Effectiveness of Motion-Boundary Adaptation Losses. In Table 3, we show how motion-boundary adaptation losses contribute to the final result. $\mathcal{L}_{feat}^{align}$ and $\mathcal{L}_{flow}^{consis}$ aim to transfer optical flow knowledge from the clean domain to the degraded domain via the consistency loss. The goal of $\mathcal{L}_{errIntra}^{contra}$ and $\mathcal{L}_{errInter}^{contra}$ is to distinguish the motion boundary of image structure from that of the degradation errors, and also align the motion boundary distribution of different scenes. We can observe that motion consistency losses make a major contribution to the optical flow results and boundary contrastive losses further improve the final results.

Discussion

Contrastive Sampling Strategies. The positive and negative sampling strategy is of great importance to the quality of the discriminative feature in contrastive learning. In Table 4, we show the advantage of the proposed edge-aware and entropy-aware sampling strategies over the conventional random sampling. Compared with random sampling, the edge-aware sampling strategy in the clean domain could greatly improve the optical flow. The main reason is most random sample patches in the warp error map are the meaningless region with zero values. On the contrary, the sharp edge-aware sampling would guarantee us abundant discriminative information so as to better reflect the alignment of the optical flow boundary. Similarly, the entropy-aware sampling could also select the informative highly-informative boundary of warp error caused by degradations, facilitating to differ the flow boundary of the image structure from that of degradations.

Shared v.s. Unshared Weight. To study the influence of flow




Figure 6: The proposed method still faces great challenging under dense foggy, especially for objects in distant regions.

encoders, in Table 4, we design the shared- and unshared-weight for the optical flow encoders. We can observe that the performance of the unshared-weight strategy is better than that of the shared-weight strategy. Imagine that if the two flow encoders are exactly shared, the learned features are hard to be aligned into the common space due to the different inputs of the two domains. The optical flow of clean domain is learned via the photometric loss, and transferred to the degraded domain via the consistency loss. Thus, it is natural to enforce different learning weights for the two encoders.

Choice of the Image Translation Network. The key to the proposed method is to build the relationship between the clean and degraded domains. Thus, the image translation between the two domains also plays an important role. We have tested the classical PatchGAN (Isola et al. 2017) and also the CycleGAN (Zhu et al. 2017) as the baseline, as shown in Table 4. It is shown that the CycleGAN is slightly better than the PatchGAN. Thus, we choose CycleGAN as our image translation network. Note that, the proposed method utilizes the self-supervised motion consistency and boundary inconsistency relationship between both domains, which significantly relieves the requirement of absolute clean images.

Limitation. The HMBA-FlowNet may fail for distant objects under dense fog. In Fig. 6, although our method can well estimate the optical flow of the nearby object, for the distant object our method may suffer from challenges. In this work, we simply assume the adverse artifacts obey a linear additive model via Eq. (1). However, fog is a non-uniform degradation highly related to scene depth, in which degraded optical flow at the nearby distance can be approximated, but flow in the distant area does not satisfy the degradation model. In the future, we will incorporate depth into our degradation model.

Conclusion

In this work, we propose the first unsupervised hierarchical motion-boundary domain adaptation framework for adverse weather optical flow. We mathematically reveal that optical flow consistency and warp error inconsistency between clean and degraded domains can be well utilized for optical flow estimation. Thus, we transfer knowledge from clean domain to degraded domain by learning a motion-invariance common space with the consistency constraint. Moreover, we refine the motion boundary via boundary contrastive adaptation to effectively distinguish the motion boundary from the warp error caused by degradation. The hierarchy motion-boundary adaptation benefits us to handle the various adverse weather optical flow. We have conducted extensive experiments to verify the superiority of the proposed HMBA-FlowNet.

Acknowledgments

This work was supported in part by the National Natural Science Foundation of China under Grant 61971460, in part by JCJQ Program under Grant 2021-JCJQ-JJ-0060, and in part by the National Natural Science Foundation of China under Grant 62101294.

References

- Chen, C.; and Li, H. 2021. Robust Representation Learning with Feedback for Single Image Deraining. In *IEEE Conf. Comput. Vis. Pattern Recog.*, 7738–7747.
- Chen, S.; Jia, X.; He, J.; Shi, Y.; and Liu, J. 2021a. Semi-supervised domain adaptation based on dual-level domain mixing for semantic segmentation. In *IEEE Conf. Comput. Vis. Pattern Recog.*, 11018–11027.
- Chen, S.; Xie, G.; Liu, Y.; Peng, Q.; Sun, B.; Li, H.; You, X.; and Shao, L. 2021b. Hsva: Hierarchical semantic-visual adaptation for zero-shot learning. *Adv. Neural Inform. Process. Syst.*, 34: 16622–16634.
- Chen, T.; Kornblith, S.; Norouzi, M.; and Hinton, G. 2020. A simple framework for contrastive learning of visual representations. In *Int. Conf. on Machine Learning*, 1597–1607.
- Dollár, P.; and Zitnick, C. L. 2013. Structured forests for fast edge detection. In *Int. Conf. Comput. Vis.*, 1841–1848.
- Dosovitskiy, A.; Fischer, P.; Ilg, E.; Hausser, P.; Hazirbas, C.; Golkov, V.; Van Der Smagt, P.; Cremers, D.; and Brox, T. 2015. Flownet: Learning optical flow with convolutional networks. In *Int. Conf. Comput. Vis.*, 2758–2766.
- Fan, H.; Lin, L.; Yang, F.; Chu, P.; Deng, G.; Yu, S.; Bai, H.; Xu, Y.; Liao, C.; and Ling, H. 2019. Lasot: A high-quality benchmark for large-scale single object tracking. In *IEEE Conf. Comput. Vis. Pattern Recog.*, 5374–5383.
- Fu, X.; Huang, J.; Ding, X.; Liao, Y.; and Paisley, J. 2017a. Clearing the skies: A deep network architecture for single-image rain removal. *IEEE Trans. Image Process.*, 26(6): 2944–2956.
- Fu, X.; Huang, J.; Zeng, D.; Huang, Y.; Ding, X.; and Paisley, J. 2017b. Removing rain from single images via a deep detail network. In *IEEE Conf. Comput. Vis. Pattern Recog.*, 3855–3863.
- Gao, H.; Guo, J.; Wang, G.; and Zhang, Q. 2022. Cross-Domain Correlation Distillation for Unsupervised Domain Adaptation in Nighttime Semantic Segmentation. In *IEEE Conf. Comput. Vis. Pattern Recog.*, 9913–9923.
- Goodfellow, I.; Pouget-Abadie, J.; Mirza, M.; Xu, B.; Warde-Farley, D.; Ozair, S.; Courville, A.; and Bengio, Y. 2014. Generative adversarial nets. *Adv. Neural Inform. Process. Syst.*, 27.
- Isola, P.; Zhu, J.-Y.; Zhou, T.; and Efros, A. A. 2017. Image-to-image translation with conditional adversarial networks. In *IEEE Conf. Comput. Vis. Pattern Recog.*, 1125–1134.
- Li, H.; Luo, K.; and Liu, S. 2021. GyroFlow: Gyroscope-Guided Unsupervised Optical Flow Learning. 12869–12878.
- Li, R.; Tan, R. T.; and Cheong, L.-F. 2018. Robust optical flow in rainy scenes. In *Eur. Conf. Comput. Vis.*, 288–304.
- Li, R.; Tan, R. T.; Cheong, L.-F.; Aviles-Rivero, A. I.; Fan, Q.; and Schonlieb, C.-B. 2019. Rainflow: Optical flow under rain streaks and rain veiling effect. In *Int. Conf. Comput. Vis.*, 7304–7313.
- Liu, L.; Zhang, J.; He, R.; Liu, Y.; Wang, Y.; Tai, Y.; Luo, D.; Wang, C.; Li, J.; and Huang, F. 2020. Learning by analogy: Reliable supervision from transformations for unsupervised optical flow estimation. In *IEEE Conf. Comput. Vis. Pattern Recog.*, 6489–6498.
- Liu, P.; King, I.; Lyu, M. R.; and Xu, J. 2019a. DdfLOW: Learning optical flow with unlabeled data distillation. In *AAAI Conf. on Arti. Intell.*, 8770–8777.
- Liu, P.; Lyu, M.; King, I.; and Xu, J. 2019b. Selfflow: Self-supervised learning of optical flow. In *IEEE Conf. Comput. Vis. Pattern Recog.*, 4571–4580.
- Liu, X.; Ma, Y.; Shi, Z.; and Chen, J. 2019c. Griddehazenet: Attention-based multi-scale network for image dehazing. In *Int. Conf. Comput. Vis.*, 7314–7323.
- Luo, K.; Wang, C.; Liu, S.; Fan, H.; Wang, J.; and Sun, J. 2021. Upflow: Upsampling pyramid for unsupervised optical flow learning. In *IEEE Conf. Comput. Vis. Pattern Recog.*, 1045–1054.
- Lv, Z.; Kim, K.; Troccoli, A.; Sun, D.; Rehg, J. M.; and Kautz, J. 2018. Learning rigidity in dynamic scenes with a moving camera for 3d motion field estimation. In *Eur. Conf. Comput. Vis.*, 468–484.
- Ma, X.; Wang, Z.; Zhan, Y.; Zheng, Y.; Wang, Z.; Dai, D.; and Lin, C.-W. 2022. Both style and fog matter: Cumulative domain adaptation for semantic foggy scene understanding. In *IEEE Conf. Comput. Vis. Pattern Recog.*, 18922–18931.
- Menze, M.; and Geiger, A. 2015. Object scene flow for autonomous vehicles. In *IEEE Conf. Comput. Vis. Pattern Recog.*, 3061–3070.
- Murphy, K. P. 2012. *Machine learning: a probabilistic perspective*. MIT press.
- Ren, W.; Ma, L.; Zhang, J.; Pan, J.; Cao, X.; Liu, W.; and Yang, M.-H. 2018. Gated fusion network for single image dehazing. In *IEEE Conf. Comput. Vis. Pattern Recog.*, 3253–3261.
- Ren, Z.; Yan, J.; Ni, B.; Liu, B.; Yang, X.; and Zha, H. 2017. Unsupervised deep learning for optical flow estimation. In *AAAI Conf. on Arti. Intell.*, 1495–1501.
- Sun, D.; Yang, X.; Liu, M.-Y.; and Kautz, J. 2018. Pwcnet: Cnns for optical flow using pyramid, warping, and cost volume. In *IEEE Conf. Comput. Vis. Pattern Recog.*, 8934–8943.
- Sun, P.; Kretschmar, H.; Dotiwalla, X.; Chouard, A.; Patnaik, V.; Tsui, P.; Guo, J.; Zhou, Y.; Chai, Y.; Caine, B.; et al. 2020. Scalability in perception for autonomous driving: Waymo open dataset. In *IEEE Conf. Comput. Vis. Pattern Recog.*, 2446–2454.
- Xie, E.; Ding, J.; Wang, W.; Zhan, X.; Xu, H.; Sun, P.; Li, Z.; and Luo, P. 2021. Detco: Unsupervised contrastive learning for object detection. In *Int. Conf. Comput. Vis.*, 8392–8401.

- Yan, W.; Sharma, A.; and Tan, R. T. 2020. Optical flow in dense foggy scenes using semi-supervised learning. In *IEEE Conf. Comput. Vis. Pattern Recog.*, 13259–13268.
- Yang, W.; Tan, R. T.; Feng, J.; Guo, Z.; Yan, S.; and Liu, J. 2020. Joint Rain Detection and Removal from a Single Image with Contextualized Deep Networks. *IEEE Trans. Pattern Anal. Mach. Intell.*, 42(6): 1377–1393.
- Yang, W.; Tan, R. T.; Feng, J.; Liu, J.; Guo, Z.; and Yan, S. 2017. Deep joint rain detection and removal from a single image. In *IEEE Conf. Comput. Vis. Pattern Recog.*, 1357–1366.
- Yu, J. J.; Harley, A. W.; and Derpanis, K. G. 2016. Back to basics: Unsupervised learning of optical flow via brightness constancy and motion smoothness. In *Eur. Conf. Comput. Vis.*, 3–10. Springer.
- Zhang, H.; and Patel, V. M. 2018. Density-aware single image de-raining using a multi-stream dense network. In *IEEE Conf. Comput. Vis. Pattern Recog.*, 695–704.
- Zhu, J.-Y.; Park, T.; Isola, P.; and Efros, A. A. 2017. Unpaired image-to-image translation using cycle-consistent adversarial networks. In *Int. Conf. Comput. Vis.*, 2223–2232.