CoT-AMFlow: Adaptive Modulation Network with Co-Teaching Strategy for Unsupervised Optical Flow Estimation

Hengli Wang
ECE Department
HKUST
hwangdf@connect.ust.hk

Rui Fan
CSE Department
UC San Diego
rui.fan@ieee.org

Ming Liu
ECE Department
HKUST
eelium@ust.hk

Abstract: The interpretation of ego motion and scene change is a fundamental task for mobile robots. Optical flow information can be employed to estimate motion in the surroundings. Recently, unsupervised optical flow estimation has become a research hotspot. However, unsupervised approaches are often easy to be unreliable on partially occluded or texture-less regions. To deal with this problem, we propose CoT-AMFlow in this paper, an unsupervised optical flow estimation approach. In terms of the network architecture, we develop an adaptive modulation network that employs two novel module types, flow modulation modules (FMMs) and cost volume modulation modules (CMMs), to remove outliers in challenging regions. As for the training paradigm, we adopt a co-teaching strategy, where two networks simultaneously teach each other about challenging regions to further improve accuracy. Experimental results on the MPI Sintel, KITTI Flow and Middlebury Flow benchmarks demonstrate that our CoT-AMFlow outperforms all other state-of-the-art unsupervised approaches, while still running in real time. Our project page is available at <https://sites.google.com/view/cot-amflow>.

Keywords: optical flow, unsupervised learning, co-teaching strategy.

1 Introduction

Mobile robots typically operate in complex environments that are inherently dynamic [1]. Therefore, it is important for such autonomous systems to be conscious of dynamic objects in their surroundings. Optical flow describes pixel-level correspondence between two ordered images, and can be regarded as a useful representation for dynamic object detection. Therefore, many approaches for mobile robot tasks, such as SLAM [2], dynamic object detection [3] and robot navigation [4], incorporate optical flow information to improve their performance.

With the development of deep learning technology, deep neural networks have presented highly compelling results for optical flow estimation [5, 6, 7]. These networks typically excel at learning optical flow estimation from large amounts of data along with hand-labeled ground truth. However, this data labeling process can be extremely time-consuming and labor-intensive. Recent unsupervised optical flow estimation approaches have attracted much attention, because their advantage in not requiring ground truth enables them to be easily deployed in real-world applications [8, 9, 10, 11, 12]. However, their performance in challenging regions, such as partially occluded or texture-less regions, is often unsatisfactory [10, 13]. The underlying cause of this performance degradation is threefold: 1) The popular coarse-to-fine framework [12, 13] is often sensitive to noises in the flow initialization from the preceding pyramid level, and the challenging regions can introduce errors in the flow estimations, which in turn propagate to subsequent levels. 2) The commonly used cost volume [10, 11] for establishing feature correspondence can contain many outliers due to the ambiguous correspondence in challenging regions. However, most existing networks directly send the noisy cost volume to the following flow estimation layers without explicitly alleviating the impact of outliers. 3) Many training strategies have been proposed to improve accuracy in chal-

Figure 1: An overview of our CoT-AMFlow. We integrate self-supervision into a co-teaching framework, where two AMFlows with different initializations teach each other about challenging regions to improve stability against outliers and further enhance the accuracy of flow estimation.

lenging regions for unsupervised optical flow estimation, such as occlusion reasoning [9, 10] and self-supervision [11, 12, 13]. These strategies generally train a single network to provide prior information. However, the prior information is not accurate enough because a single network can be easily disturbed by outliers if the ground truth is inaccessible. Also, the inaccurate prior information can further lead to significant performance degradation.

To overcome these limitations, we propose CoT-AMFlow, which comprises adaptive modulation networks, named AMFlows, that learn optical flow estimation in an unsupervised way with a co-teaching strategy. The overview of our proposed CoT-AMFlow is illustrated in Fig. 1, and we leverage three novel techniques to improve the flow accuracy, as follows:

- We apply *flow modulation modules* (FMMs) in our AMFlow to refine the flow initialization from the preceding pyramid level using local flow consistency, which can address the issue of accumulated errors.
- We present *cost volume modulation modules* (CMMs) in our AMFlow to explicitly reduce outliers in the cost volume using a flexible and efficient sparse point-based scheme.
- We adopt a *co-teaching* strategy, where two AMFlows with different initializations simultaneously teach each other about challenging regions to improve robustness against outliers.

We conduct extensive experiments on the MPI Sintel [14], KITTI Flow 2012 [15], KITTI Flow 2015 [16] and Middlebury Flow [17] benchmarks. Experimental results show that our CoT-AMFlow outperforms all other unsupervised approaches, while still running in real time.

2 Related Work

2.1 Optical Flow Estimation

Traditional approaches typically estimate optical flow by minimizing a global energy that measures both brightness consistency and spatial smoothness [18, 19, 20]. With recent development in deep learning technology, supervised approaches using convolutional neural networks (CNNs) have been extensively applied in optical flow estimation, and the achieved results are very promising. FlowNet [5] was the first end-to-end deep neural network for optical flow estimation. It employs a correlation layer to compute feature correspondence. Later on, PWC-Net [6] and LiteFlowNet [7] presented a pyramid architecture, which consists of feature warping layers, cost volumes and flow estimation layers. Such an architecture can achieve remarkable flow accuracy and high efficiency simultaneously. Their subsequent versions [21, 22] also made incremental improvements. Unsupervised approaches generally adopt similar network architectures to supervised approaches, and

Figure 2: An illustration of our AMFlow, which uses FMMs and CMMs to refine flow initializations and remove outliers in cost volumes, respectively.

focus more on training strategies. However, existing network architectures do not explicitly address the issues of noisy flow initializations and outliers in the cost volume, as previously mentioned. Therefore, we develop the FMMs and CMMs in our AMFlow to overcome these limitations.

Among the training strategies for unsupervised approaches, DSTFlow [8] first presented a photometric loss and a smoothness loss for unsupervised training. Additionally, some approaches train a single network to perform occlusion reasoning for accuracy improvement [9, 10]. Self-supervision [11, 12] is also an important strategy for unsupervised training. It first trains a single network to generate flow labels, and then conducts data augmentation to make flow estimations more challenging. The augmented samples are further employed as supervision to train another network. One variant of self-supervision is to train only one network with a two-forward process [13]. However, training a single network to provide flow labels is likely to be unreliable due to the disturbance of outliers and the lack of ground-truth supervision. To address this issue, we integrate self-supervision into a co-teaching framework, where two networks simultaneously teach each other about challenging regions to improve stability against outliers.

2.2 Co-Teaching Strategy

The co-teaching strategy was first proposed for the image classification task with extremely noisy labels [23]. Since then, many researchers have resorted to this strategy for various specific robust training tasks, such as face recognition [24] and object detection [25]. The main difference between previous studies and our approach is that they focus on the task of supervised learning with noisy labels, while we focus on the task of unsupervised learning. Moreover, the noises in their tasks exist at image level (noisy image classification labels), while the outliers in our task exist at pixel level (inaccurate flow estimation pixels in challenging regions).

3 Methodology

3.1 AMFlow

In this subsection, we first introduce the overall architecture of our AMFlow, and then present our FMM and CMM. Since we use many notations, we suggest readers refer to the glossary provided in

the appendix for better understanding. Fig. 2 illustrates an overview of our proposed AMFlow, which follows the pipeline of PWC-Net [6]. Different pyramid levels of feature maps are first extracted hierarchically from the input images $\mathbf{I}_1$ and $\mathbf{I}_2$ using a siamese feature pyramid network, and then are sent to the coarse-to-fine flow decoder. Here, we take level l as an example to introduce our flow decoder, for simplicity. First, the upsampled flow estimation $\bar{\mathbf{F}}_{12}^{l+1}$ at level $l+1$ is processed by our FMM for refinement, and the generated modulated flow $\hat{\mathbf{F}}_{12}^{l+1}$ is employed to align the feature map $\mathbf{x}_2^l$ with the feature map $\mathbf{x}_1^l$. A correlation operation is then employed to compute the cost volume $\mathbf{C}^l$, which is then processed by our CMM to remove outliers. After getting the modulated cost volume $\hat{\mathbf{C}}^l$, we take it as input and employ the same flow estimation layer as PWC-Net [6] to estimate the flow residual, which is subsequently added with $\hat{\mathbf{F}}_{12}^{l+1}$ to obtain the flow estimation $\mathbf{F}_{12}^l$ at level l . This process iterates and the flow estimations at different scales are generated.

Flow Modulation Module (FMM). In the coarse-to-fine framework, a flow estimation from the preceding level is adopted as a flow initialization at the current level. Therefore, the inaccurate flow estimations in challenging regions can propagate to subsequent levels and cause significant performance degradation. Our FMM is developed to address this problem based on the concept of local flow consistency [26].

Our FMM is based on the assumption that the neighboring pixels with similar feature maps should have similar optical flows. Therefore, for a pixel $\mathbf{p}$ with an inaccurate flow estimation $\mathbf{F}(\mathbf{p})$, we will look for another pixel $\mathbf{q}$ around $\mathbf{p}$, which has a similar feature map to $\mathbf{p}$ and an accurate flow estimation $\mathbf{F}(\mathbf{q})$. Then, we replace $\mathbf{F}(\mathbf{p})$ with $\mathbf{F}(\mathbf{q})$.

To this end, we first compute a confidence map $\mathbf{M}^l$ based on the upsampled flow estimation $\bar{\mathbf{F}}_{12}^{l+1}$ and the downsampled input images $\mathbf{I}_1^l$ and $\mathbf{I}_2^l$, as illustrated in Fig. 2. The confidence computing operation is defined as follows:

$$\mathbf{M}^l = \exp \left(- \left| \mathcal{B} \left(\mathbf{I}_1^l, \omega \left(\mathbf{I}_2^l, \bar{\mathbf{F}}_{12}^{l+1} \right) \right) \right| \right), \quad (1)$$

where $\mathcal{B}(\cdot, \cdot)$ denotes the function for measuring the photometric difference [13], and $\omega(\mathbf{I}, \mathbf{F})$ denotes the warping operation of image $\mathbf{I}$ based on flow $\mathbf{F}$. Then, we use a self-correlation layer to compute a self-cost volume $\mathbf{C}_s^l$, which measures the similarity between each pixel in the feature map $\mathbf{x}_1^l$ and its neighboring pixels. The adopted self-correlation layer is identical to the correlation layer used in the above-mentioned flow decoder, except that it only takes one feature map as input. We further concatenate $\mathbf{M}^l$ with $\mathbf{C}_s^l$, and send the concatenation to several convolution layers to obtain a displacement map $\mathbf{D}^l$. Finally, we warp $\bar{\mathbf{F}}_{12}^{l+1}$ based on $\mathbf{D}^l$ to get the modulated flow estimation $\hat{\mathbf{F}}_{12}^{l+1}$.

Cost Volume Modulation Module (CMM). Ambiguous correspondence in challenging regions can introduce noises into the cost volume, which further influence the subsequent flow estimation layers. Our CMM is designed to reduce noises in the cost volume.

Several traditional approaches have formulated the task of denoising the cost volume as a weighted least squares problem, which obtains the following solution for level l [27, 28]:

$$\hat{\mathbf{C}}^l(\mathbf{p}, f) = \sum_{\mathbf{q} \in \mathcal{N}^l(\mathbf{p})} w^l(\mathbf{p}, \mathbf{q}) \cdot \mathbf{C}^l(\mathbf{q}, f), \quad (2)$$

where $\hat{\mathbf{C}}^l(\mathbf{p}, f)$ denotes the modulated cost at pixel $\mathbf{p}$ for flow residual candidate f ; pixel $\mathbf{q}$ belongs to the neighbors $\mathcal{N}^l(\mathbf{p})$ of $\mathbf{p}$; $w^l(\mathbf{p}, \mathbf{q})$ denotes the modulation weight; and $\mathbf{C}^l(\mathbf{q}, f)$ denotes the original cost at pixel $\mathbf{q}$ for flow residual candidate f . Note that the one-dimensional f is transformed from the original two-dimensional flow residual candidate for simplicity, which is the same as the scheme adopted in PWC-Net [6].

The intuition of our CMM is to implement (2) in deep neural networks, which is realized by a flexible and efficient sparse point-based scheme based on deformable convolution [29]:

$$\hat{\mathbf{C}}^l(\mathbf{p}, f) = \sum_{k=1}^K w_k^l \cdot \mathbf{C}^l \left(\mathbf{p} + \mathbf{p}_k + \Delta \mathbf{p}_k^l, f \right) \cdot \Delta m_k^l, \quad (3)$$

where K denotes the number of sampling points; w_k^l denotes the modulation weight for the k -th point; and $\mathbf{p}_k$ is the fixed offset of the original convolution layer to $\mathbf{p}$. To make the modulation

Algorithm 1: Co-Teaching Strategy

Input: Θ^A and Θ^B , learning rate η , constant threshold τ , epoch T_k and $T_{\max}$, iteration $N_{\max}$.

Output: Θ^A and Θ^B .

```
1 for  $T = 1 \rightarrow T_{\max}$  do
2   Shuffle training set  $\mathcal{D}$ 
3   for  $N = 1 \rightarrow N_{\max}$  do
4     Forward individually to obtain  $\mathbf{F}_{12}^i, \mathbf{O}_{12}^i, \tilde{\mathbf{F}}_{12}^i, \tilde{\mathbf{F}}_{12}^{i*}$  and  $\tilde{\mathbf{O}}_{12}^i, i \in \{A, B\}$ 
5     Set  $\mathbf{O}_{12}^i$  ( $\mathbf{O}_{12}^i > \mathcal{R}(T)$ ) = 1,  $i \in \{A, B\}$   $\triangleright$  Filter out pixels with high occlusion
        probability
6     Compute  $\mathcal{L}^A = \mathcal{L}_{\text{ph}}^A(\mathbf{I}_1, \mathbf{I}_2, \mathbf{F}_{12}^A, \mathbf{O}_{12}^A) + \lambda_1 \cdot \mathcal{L}_{\text{sm}}^A(\mathbf{I}_1, \mathbf{F}_{12}^A) + \lambda_2 \cdot \mathcal{L}_{\text{ss}}^A(\tilde{\mathbf{F}}_{12}^A, \tilde{\mathbf{F}}_{12}^{A*}, \tilde{\mathbf{O}}_{12}^A)$ 
7     Compute  $\mathcal{L}^B = \mathcal{L}_{\text{ph}}^B(\mathbf{I}_1, \mathbf{I}_2, \mathbf{F}_{12}^B, \mathbf{O}_{12}^B) + \lambda_1 \cdot \mathcal{L}_{\text{sm}}^B(\mathbf{I}_1, \mathbf{F}_{12}^B) + \lambda_2 \cdot \mathcal{L}_{\text{ss}}^B(\tilde{\mathbf{F}}_{12}^B, \tilde{\mathbf{F}}_{12}^{B*}, \tilde{\mathbf{O}}_{12}^B)$ 
8     Update  $\Theta^i = \Theta^i - \eta \nabla \mathcal{L}^i, i \in \{A, B\}$ 
9   end
10  Update  $\mathcal{R}(T) = 1 - \tau \cdot \min \left\{ \frac{T}{T_k}, 1 \right\}$ 
11 end
```

scheme more flexible, we also employ a separate convolutional layer on $\mathbf{C}^l$ to learn an additional offset $\Delta \mathbf{p}_k^l$ and a spatial-variant weight Δm_k^l . These two terms can effectively and efficiently help remove outliers in challenging regions.

3.2 Loss Function

We employ three common loss functions, 1) photometric loss $\mathcal{L}_{\text{ph}}$, 2) smoothness loss $\mathcal{L}_{\text{sm}}$ and 3) self-supervision loss $\mathcal{L}_{\text{ss}}$, to train our CoT-AMFlow, as illustrated in Fig. 1. For each network, the forward flow $\mathbf{F}_{12}$ and backward flow $\mathbf{F}_{21}$ can be obtained given the input images $\mathbf{I}_1$ and $\mathbf{I}_2$. Then, we can compute an occlusion map $\mathbf{O}_{12}$ with the range between 0 and 1 [10], where a higher value indicates that the corresponding pixel is more likely to be occluded, and vice versa. Based on these notations, we first introduce our adopted photometric loss [10] as follows:

$$\mathcal{L}_{\text{ph}}(\mathbf{I}_1, \mathbf{I}_2, \mathbf{F}_{12}, \mathbf{O}_{12}) = \frac{\sum_{\mathbf{p}} \psi(\mathcal{B}(\mathbf{I}_1, \omega(\mathbf{I}_2, \mathbf{F}_{12}))) \odot (1 - \mathcal{S}(\mathbf{O}_{12}))}{\sum_{\mathbf{p}} (1 - \mathcal{S}(\mathbf{O}_{12}))}, \quad (4)$$

where $\psi(x) = \sqrt{x^2 + 0.001^2}$ is the generalized Charbonnier penalty function [30]; $\mathcal{S}(\cdot)$ stands for the stop-gradient; and $\odot$ denotes element-wise multiplication. (4) shows that occluded regions have little impact on $\mathcal{L}_{\text{ph}}$, since there does not exist correspondence in these regions. Moreover, we stop the gradient at the occlusion maps to avoid a trivial solution. Then, the following formulation shows our utilized second-order edge-aware smoothness loss [31]:

$$\mathcal{L}_{\text{sm}}(\mathbf{I}_1, \mathbf{F}_{12}) = \frac{1}{N_{\mathbf{p}}} \sum_{\mathbf{p}} \sum_{d \in \{x, y\}} \exp \left(-50 \sum_c \left| \frac{\partial \mathbf{I}_1}{\partial d} \right| \right) \cdot \left| \frac{\partial^2 \mathbf{F}_{12}}{\partial d^2} \right|, \quad (5)$$

where c denotes the color channel and $N_{\mathbf{p}}$ is the total number of pixels. We also adopt a self-supervision scheme [13]. Specifically, we first conduct transformations $\mathbf{T}_{\theta}^{\text{img}}$, $\mathbf{T}_{\theta}^{\text{flo}}$ and $\mathbf{T}_{\theta}^{\text{occ}}$ on $(\mathbf{I}_1, \mathbf{I}_2)$, $\mathbf{F}_{12}$ and $\mathbf{O}_{12}$ respectively to construct augmented samples $\tilde{\mathbf{I}}_1, \tilde{\mathbf{I}}_2, \tilde{\mathbf{F}}_{12}$ and $\tilde{\mathbf{O}}_{12}$. The transformations include spatial, occlusion and appearance transformation [13]. We also obtain a flow prediction $\tilde{\mathbf{F}}_{12}^*$ based on $\tilde{\mathbf{I}}_1$ and $\tilde{\mathbf{I}}_2$. Then, our self-supervision loss is shown as follows [11]:

$$\mathcal{L}_{\text{ss}}(\tilde{\mathbf{F}}_{12}, \tilde{\mathbf{F}}_{12}^*, \tilde{\mathbf{O}}_{12}) = \frac{\sum_{\mathbf{p}} \psi(\|\mathcal{S}(\tilde{\mathbf{F}}_{12}) - \tilde{\mathbf{F}}_{12}^*\|_2) \odot \mathcal{S}(\tilde{\mathbf{O}}_{12})}{\sum_{\mathbf{p}} \mathcal{S}(\tilde{\mathbf{O}}_{12})}, \quad (6)$$

where $\|\cdot\|_2$ denotes the L2 norm. Note that, different from $\mathbf{O}_{12}$, $\tilde{\mathbf{O}}_{12}$ measures the occlusion relationship between $\tilde{\mathbf{F}}_{12}$ and $\tilde{\mathbf{F}}_{12}^*$. A higher value in $\tilde{\mathbf{O}}_{12}$ indicates that the corresponding pixel is less likely to be occluded in $\tilde{\mathbf{F}}_{12}$ but more likely to be occluded in $\tilde{\mathbf{F}}_{12}^*$ [11]. Therefore, (6) shows that $\mathcal{L}_{\text{ss}}$ helps improve the accuracy of flow estimations in challenging regions.

The whole loss function for training our CoT-AMFlow is a weighted sum of the above three losses, as shown on Line 6 and 7 in Algorithm 1. The details will be introduced in Section 3.3.

3.3 Co-Teaching Strategy

Our co-teaching strategy is illustrated in Fig. 1, and the corresponding steps are shown in Algorithm 1. Specifically, we simultaneously train two networks A (with parameter Θ^A) and B (with parameter Θ^B). In each mini-batch, we first let the two networks forward individually to obtain several outputs (Line 4). Then, we filter out the pixels with a high occlusion probability by setting their value in the occlusion map as 1 (completely occluded and thus have no impact on $\mathcal{L}_{ph}$) (Line 5). The filtering threshold is controlled by $\mathcal{R}(T)$, which equals 1 at the beginning and then decreases gradually with the increase of epoch number. The key point of our co-teaching strategy is that each network uses the occlusion maps estimated by the other network to compute its own loss function (Line 6 and 7). Finally, we update the parameters of the two networks separately and also update $\mathcal{R}(T)$ (Line 8 and 10). Next, we will answer two important questions about our co-teaching strategy: 1) Why do we need a dynamic threshold $\mathcal{R}(T)$ and 2) why can swapping the occlusion maps estimated by two networks help improve the accuracy for unsupervised optical flow estimation?

To answer the first question, we know that it is meaningless to compute photometric loss on the occluded regions, and thus we adopt an occlusion-masked photometric loss. According to [32], networks will first learn easy and clear patterns, *i.e.*, unchallenging regions. However, with the number of epochs increasing, networks will gradually be affected by the inaccurately estimated occlusion maps and thus overfit on the occluded regions, which in turn will lead to more inaccurate occlusion estimations and further cause significant performance degradation. To address this, we keep more pixels in the initial epochs, *i.e.*, $\mathcal{R}(T)$ is large. Then, we gradually filter our pixels with high occlusion probability, *i.e.*, $\mathcal{R}(T)$ gradually decreases, to ensure the networks do not memorize these possible outliers.

The dynamic threshold can, however, only alleviate but not entirely avoid the adverse impact of the occluded regions. Therefore, we further adopt a scheme with two networks, which connects to the answer to our second question. The intuition is that different networks have different abilities to learn flow estimation, and correspondingly, they can generate different occlusion estimations. Therefore, swapping the occlusion maps estimated by the two networks can help them adaptively correct the inaccurate occlusion estimations. Compared with most existing approaches that directly transfer errors back to themselves, our co-teaching strategy can effectively avoid the accumulation of errors and thus improve stability against outliers for unsupervised optical flow estimation. Note that since deep neural networks are highly non-convex and a network with different initializations can lead to different local optimums, we employ two AMFlows with different initializations in our CoT-AMFlow, following [23], as illustrated in Fig. 1.

4 Experimental Results

4.1 Dataset and Implementation Details

In our experiments, we set $\lambda_1 = 2$ in our loss function. In addition, we use $\lambda_2 = 0$ for the first 40% of epochs and increase it to 0.15 linearly for the next 20% of epochs, after which it stays at a constant value. The learning rate η adopts an exponential decay scheme, with the initialization as 10^{-4} , and the Adam optimizer is used. Moreover, we set $\tau = 0.8$ and $T_k = 0.1T_{max}$ in Algorithm 1 for evaluation on public benchmarks.

We first evaluate our CoT-AMFlow on three popular optical flow benchmarks, MPI Sintel [14], KITTI Flow 2012 [15] and KITTI Flow 2015 [16]. The experimental results are shown in Section 4.2. Then, we perform a generalization evaluation on the Middlebury Flow benchmark [17], as presented in Section 4.3. We also conduct extensive ablation studies to demonstrate the superiority of 1) our selection of τ and T_k ; 2) our FMM and CMM; 3) our AMFlow over other network architectures; and 4) our co-teaching strategy over other strategies for unsupervised training. The experimental results are presented in the appendix.

Furthermore, we follow a similar training scheme to those of the previous unsupervised approaches [11, 12, 13] for fair comparison. For the MPI Sintel benchmark, we first train our model on raw movie frames and then fine-tune it on the training set. For the two KITTI Flow benchmarks, we

Table 1: Evaluation results on the MPI Sintel, KITTI Flow 2012 and KITTI Flow 2015 benchmarks. Here, we show the primary evaluation metrics used on each benchmark. For the Sintel Clean and Final benchmarks, the AEPE (px) for all pixels is presented. For the KITTI Flow 2012 and 2015, “Noc” and “All” represent the F1 (%) for non-occluded pixels and all pixels, respectively. “S” denotes supervised approaches. [†] indicates the network using more than two frames. Best results for supervised and unsupervised approaches are both shown in bold font.

Approach	S	MPI Sintel		KITTI 2012		KITTI 2015		
		Clean	Final	Noc	All	Noc	All	Time (s)
PWC-Net [6]	✓	4.39	5.04	4.22	8.10	6.12	9.60	0.03
LiteFlowNet [7]	✓	4.54	5.38	3.27	7.27	5.49	9.38	0.09
LiteFlowNet2 [22]	✓	3.48	4.69	2.63	6.16	4.42	7.62	0.05
MaskFlowNet [33]	✓	2.52	4.17	2.07	4.82	3.92	6.11	0.06
RAFT [34]	✓	1.61	2.86	–	–	3.07	5.10	0.20
UnFlow [9]	–	9.38	10.22	4.28	8.42	7.46	11.11	0.12
DDFlow [11]	–	6.18	7.40	4.57	8.86	9.55	14.29	0.06
SelfFlow [†] [12]	–	6.56	6.57	4.31	7.68	9.65	14.19	0.09
ARFlow [13]	–	4.78	5.89	4.71	8.49	8.91	11.80	0.01
ARFlow-mv [†] [13]	–	4.49	5.67	4.56	7.53	8.97	11.79	0.02
UFlow [35]	–	5.21	6.50	4.26	7.91	8.41	11.13	0.04
CoT-AMFlow (Ours)	–	3.96	5.14	3.50	8.26	6.28	10.34	0.06

first employ the KITTI raw dataset to pre-train our model and then fine-tune it using multi-view extension data. Additionally, we adopt two standard evaluation metrics, the average end-point error (AEPE) and the percentage of erroneous pixels (F1) [14, 15, 16, 17].

4.2 Performance on Public Benchmarks

According to the online leaderboards of the MPI Sintel¹, KITTI Flow 2012² and KITTI Flow 2015³ benchmarks, as shown in Table 1, our CoT-AMFlow outperforms all other unsupervised optical flow estimation approaches. We can clearly observe that our approach is significantly ahead of other unsupervised approaches, especially on the MPI Sintel benchmark, where an AEPE improvement of 0.53px–5.42px is achieved on the Sintel Clean benchmark. We also use the KITTI Flow 2015 benchmark to record the average inference time of our CoT-AMFlow. The results in Table 1 show that our approach can still run in real time with the state-of-the-art performance. One exciting fact is that our unsupervised CoT-AMFlow can achieve considerable performance when compared with supervised approaches. Specifically, on the MPI Sintel Clean benchmark, our CoT-AMFlow outperforms some classic networks such as PWC-Net [6] and LiteFlowNet [7], while achieving only a slightly inferior performance compared with LiteFlowNet2 [22], which demonstrates the effectiveness of our adaptive modulation network and co-teaching strategy. Fig. 3 illustrates examples of the three public benchmarks, where we can obviously see that our CoT-AMFlow yields more robust and accurate results.

4.3 Generalization Analysis across Datasets

We employ the CoT-AMFlow trained on the MPI Sintel benchmark directly on the Middlebury Flow benchmark to test the generalization ability of our approach. Table 2 shows the online leaderboard of the Middlebury Flow benchmark⁴. Note that our CoT-AMFlow has not been fine-tuned on the benchmark. We can observe that our CoT-AMFlow significantly outperforms the unsupervised UnFlow [9] and even presents superior performance over supervised approaches such as PWC-Net [6] and LiteFlowNet [7]. The results strongly verify that our CoT-AMFlow has an excellent generalization ability.

¹<http://sintel.is.tue.mpg.de/results>

²http://www.cvlibs.net/datasets/kitti/eval_stereo_flow.php?benchmark=flow

³http://cvlibs.net/datasets/kitti/eval_scene_flow.php?benchmark=flow

⁴<https://vision.middlebury.edu/flow/eval/results/results-e1.php>

Table 2: Evaluation results on the Middlebury Flow benchmark. “S” denotes supervised approaches. Note that our CoT-AMFlow has not been fine-tuned on the benchmark. Best results for supervised and unsupervised approaches are both shown in bold font.

Metric	PWC-Net [6]	LiteFlowNet [7]	UnFlow [9]	CoT-AMFlow (Ours)
S	✓	✓	–	–
AEPE (px)	0.33	0.40	0.76	0.26

Figure 3: Examples of the MPI Sintel Clean, KITTI Flow 2012 and KITTI Flow 2015 benchmarks, where rows (a) and (b) on columns (1)–(3) show the flow estimations and the corresponding error maps of (1) ARFlow-mv [13], (2) SelFlow [12] and (3) our CoT-AMFlow, respectively. Significantly improved regions are highlighted with green dashed boxes.

5 Conclusion

In this paper, we proposed CoT-AMFlow, an adaptive modulation network with a co-teaching strategy for unsupervised optical flow estimation. Our CoT-AMFlow presents three major contributions: 1) a flow modulation module (FMM), which can refine the flow initialization from the preceding pyramid level to address the issue of accumulated errors; 2) a cost volume modulation module (CMM), which can explicitly reduce outliers in the cost volume to improve the accuracy of optical flow estimation; and 3) a co-teaching strategy for unsupervised training, which employs two networks to teach each other about challenging regions to improve robustness against outliers for unsupervised optical flow estimation. Extensive experiments have demonstrated that our CoT-AMFlow achieves the state-of-the-art performance for unsupervised optical flow estimation with an impressive generalization ability, while still running in real time. We believe that our CoT-AMFlow can be directly used in many mobile robot tasks, such as SLAM and robot navigation, to improve their performance. It is also promising to employ the co-teaching strategy in other unsupervised tasks, such as unsupervised disparity or scene flow estimation.

Acknowledgments

We thank the anonymous reviewers for their useful comments. This work was supported by the National Natural Science Foundation of China, under grant No. U1713211, Collaborative Research Fund by Research Grants Council Hong Kong, under Project No. C4063-18G, and HKUST-SJTU Joint Research Collaboration Fund, under project SJTU20EG03, awarded to Prof. Ming Liu.

References

- [1] A. K. Ushani and R. M. Eustice. Feature learning for scene flow estimation from lidar. In *Conf. Robot Learn. (CoRL)*, pages 283–292, 2018.
- [2] T. Zhang, H. Zhang, Y. Li, Y. Nakamura, and L. Zhang. Flowfusion: Dynamic dense RGB-D SLAM based on optical flow. In *2019 Int. Conf. Robot. Automat. (ICRA)*. IEEE, 2019.
- [3] H. Wang, Y. Liu, H. Huang, Y. Pan, W. Yu, J. Jiang, D. Lyu, M. Bocus J., M. Liu, I. Pitas, and R. Fan. ATG-PVD: Ticketing parking violations on a drone. 2020.
- [4] K. Lee, J. Gibson, and E. A. Theodorou. Aggressive perception-aware navigation using deep optical flow dynamics and PixelMPC. *IEEE Robot. Automat. Lett.*, 5(2):1207–1214, 2020.
- [5] A. Dosovitskiy, P. Fischer, E. Ilg, P. Hausser, C. Hazirbas, V. Golkov, P. Van Der Smagt, D. Cremers, and T. Brox. FlowNet: Learning optical flow with convolutional networks. In *Proc. IEEE Conf. Comput. Vision Pattern Recognit. (CVPR)*, pages 2758–2766, 2015.
- [6] D. Sun, X. Yang, M.-Y. Liu, and J. Kautz. PWC-Net: CNNs for optical flow using pyramid, warping, and cost volume. In *Proc. IEEE Conf. Comput. Vision Pattern Recognit. (CVPR)*, pages 8934–8943, 2018.
- [7] T.-W. Hui, X. Tang, and C. Change Loy. LiteflowNet: A lightweight convolutional neural network for optical flow estimation. In *Proc. IEEE Conf. Comput. Vision Pattern Recognit. (CVPR)*, pages 8981–8989, 2018.
- [8] Z. Ren, J. Yan, B. Ni, B. Liu, X. Yang, and H. Zha. Unsupervised deep learning for optical flow estimation. In *Thirty-First AAAI Conf. Artif. Intell.*, 2017.
- [9] S. Meister, J. Hur, and S. Roth. Unflow: Unsupervised learning of optical flow with a bidirectional census loss. In *Thirty-Second AAAI Conf. Artif. Intell.*, 2018.
- [10] Y. Wang, Y. Yang, Z. Yang, L. Zhao, P. Wang, and W. Xu. Occlusion aware unsupervised learning of optical flow. In *Proc. IEEE Conf. Comput. Vision Pattern Recognit. (CVPR)*, pages 4884–4893, 2018.
- [11] P. Liu, I. King, M. R. Lyu, and J. Xu. DDFlow: Learning optical flow with unlabeled data distillation. In *Proceedings of the AAAI Conf. Artif. Intell.*, volume 33, pages 8770–8777, 2019.
- [12] P. Liu, M. Lyu, I. King, and J. Xu. Selfflow: Self-supervised learning of optical flow. In *Proc. IEEE Conf. Comput. Vision Pattern Recognit. (CVPR)*, pages 4571–4580, 2019.
- [13] L. Liu, J. Zhang, R. He, Y. Liu, Y. Wang, Y. Tai, D. Luo, C. Wang, J. Li, and F. Huang. Learning by analogy: Reliable supervision from transformations for unsupervised optical flow estimation. In *Proc. IEEE Conf. Comput. Vision Pattern Recognit. (CVPR)*, pages 6489–6498, 2020.
- [14] D. J. Butler, J. Wulff, G. B. Stanley, and M. J. Black. A naturalistic open source movie for optical flow evaluation. In A. Fitzgibbon et al. (Eds.), editor, *Proc. Eur. Conf. Comput. Vision (ECCV)*, Part IV, LNCS 7577, pages 611–625. Springer-Verlag, Oct. 2012.
- [15] A. Geiger, P. Lenz, and R. Urtasun. Are we ready for autonomous driving? the KITTI vision benchmark suite. In *Proc. IEEE Conf. Comput. Vision Pattern Recognit. (CVPR)*, 2012.
- [16] M. Menze and A. Geiger. Object scene flow for autonomous vehicles. In *Proc. IEEE Conf. Comput. Vision Pattern Recognit. (CVPR)*, 2015.
- [17] S. Baker, D. Scharstein, J. Lewis, S. Roth, M. J. Black, and R. Szeliski. A database and evaluation methodology for optical flow. *Inter. J. Comput. Vision*, 92(1):1–31, 2011.
- [18] B. K. Horn and B. G. Schunck. Determining optical flow. In *Techn. Appl. Image Understanding*, volume 281, pages 319–331. International Society for Optics and Photonics, 1981.

- [19] E. Mémin and P. Pérez. Dense estimation and object-based segmentation of the optical flow with robust techniques. *IEEE Trans. Image Process.*, 7(5):703–719, 1998.
- [20] T. Brox, A. Bruhn, N. Papenberg, and J. Weickert. High accuracy optical flow estimation based on a theory for warping. In *Eur. Conf. Comput. Vision (ECCV)*, pages 25–36. Springer, 2004.
- [21] Z. Ren, O. Gallo, D. Sun, M.-H. Yang, E. Sudderth, and J. Kautz. A fusion approach for multi-frame optical flow estimation. In *2019 IEEE Winter Conf. Appl. Comput. Vision (WACV)*, pages 2077–2086. IEEE, 2019.
- [22] T. Hui, X. Tang, and C. C. Loy. A lightweight optical flow CNN - revisiting data fidelity and regularization. *IEEE Trans. Pattern Anal. Mach. Intell.*, pages 1–1, 2020.
- [23] B. Han, Q. Yao, X. Yu, G. Niu, M. Xu, W. Hu, I. Tsang, and M. Sugiyama. Co-teaching: Robust training of deep neural networks with extremely noisy labels. In *Adv. Neural Inf. Process. Syst. (NIPS)*, pages 8527–8537, 2018.
- [24] F. Yang, K. Li, Z. Zhong, Z. Luo, X. Sun, H. Cheng, X. Guo, F. Huang, R. Ji, and S. Li. Asymmetric co-teaching for unsupervised cross-domain person re-identification. In *AAAI*, pages 12597–12604, 2020.
- [25] S. Chadwick and P. Newman. Training object detectors with noisy data. In *2019 IEEE Intell. Veh. Symp. (IV)*, pages 1319–1325. IEEE, 2019.
- [26] H. Zimmer, A. Bruhn, and J. Weickert. Optic flow in harmony. *Inter. J. Comput. Vision*, 93(3): 368–388, 2011.
- [27] K.-J. Yoon and I. S. Kweon. Adaptive support-weight approach for correspondence search. *IEEE Trans. Pattern Anal. Mach. Intell.*, 28(4):650–656, 2006.
- [28] A. Hosni, C. Rhemann, M. Bleyer, C. Rother, and M. Gelautz. Fast cost-volume filtering for visual correspondence and beyond. *IEEE Trans. Pattern Anal. Mach. Intell.*, 35(2):504–511, 2012.
- [29] X. Zhu, H. Hu, S. Lin, and J. Dai. Deformable convnets v2: More deformable, better results. In *Proc. IEEE Conf. Comput. Vision Pattern Recognit. (CVPR)*, pages 9308–9316, 2019.
- [30] D. Sun, S. Roth, and M. J. Black. Secrets of optical flow estimation and their principles. In *2010 IEEE Comput. Soc. Conf. Comput. Vision Pattern Recognit.*, pages 2432–2439. IEEE, 2010.
- [31] C. Tomasi and R. Manduchi. Bilateral filtering for gray and color images. In *Sixth Inter. Conf. Comput. Vision (IEEE Cat. No. 98CH36271)*, pages 839–846. IEEE, 1998.
- [32] D. Arpit, S. K. Jastrzebski, N. Ballas, D. Krueger, E. Bengio, M. S. Kanwal, T. Maharaj, A. Fischer, A. C. Courville, Y. Bengio, et al. A closer look at memorization in deep networks. In *ICML*, 2017.
- [33] S. Zhao, Y. Sheng, Y. Dong, E. I. Chang, Y. Xu, et al. Maskflownet: Asymmetric feature matching with learnable occlusion mask. In *Proc. IEEE Conf. Comput. Vision Pattern Recognit. (CVPR)*, pages 6278–6287, 2020.
- [34] Z. Teed and J. Deng. Raft: Recurrent all-pairs field transforms for optical flow. In *Eur. Conf. Comput. Vision (ECCV)*, 2020.
- [35] R. Jonschkowski, A. Stone, J. T. Barron, A. Gordon, K. Konolige, and A. Angelova. What matters in unsupervised optical flow. In *Eur. Conf. Comput. Vision (ECCV)*, 2020.

Appendix

This appendix first provides a glossary of notations used in the paper to help readers better follow, as presented in Section A. Then, we present the experimental results of our ablation studies. Specifically, we first split the original MPI Sintel training set [14] into a new training set and a validation set. Then, all models in our ablation studies are trained on the new training set and evaluated on the validation set. In addition, we adopt the average end-point error (AEPE) [14] as the evaluation metric.

In our ablation studies, we first explore the impact of different T_k and τ in the proposed co-teaching strategy on the performance, as presented in Section B. Then, we verify the effectiveness of our flow modulation module (FMM) and cost volume modulation module (CMM) in Section C. We also demonstrate the superiority of our AMFlow over other network architectures and the superiority of our co-teaching (CoT) strategy over other strategies for unsupervised training, as presented in Section D and Section E, respectively.

A Glossary of Notations

The glossary of notations used in the paper is presented in Table 3.

B Impact of Different T_k and τ

In our co-teaching strategy, T_k and τ controls the filtering speed and filtering range of the pixels with high occlusion probability, respectively. We consider three values of T_k , $T_k = \{0.05 \cdot T_{max}, 0.10 \cdot T_{max}, 0.15 \cdot T_{max}\}$ and five values of τ , $\tau = \{0.70, 0.75, 0.80, 0.85, 0.90\}$. We also test the training schemes that adopt a constant τ . The results of our CoT-AMFlow are shown in Table 4. We can observe that the dynamic threshold scheme can effectively improve the performance and our CoT-AMFlow is robust on different choices of T_k . Moreover, τ has a significant impact on the performance. Specifically, a higher τ indicates that more pixels will be filtered out. We can see that the performance can be improved when τ increases. However, when too many pixels are filtered out, *i.e.*, $\tau = 0.85$ or $\tau = 0.90$, the performance can deteriorate because the networks cannot get sufficient training data. Note that we set $T_k = 0.10 \cdot T_{max}$ and $\tau = 0.80$ in the rest of our ablation studies.

C Effectiveness of Our FMM and CMM

Table 5 shows the evaluation results of variants of our CoT-AMFlow with some of the proposed modules disabled. We can observe that our FMM and CMM can effectively improve the optical flow accuracy, especially for the pixels with large movements. This is because our FMM can refine the flow initialization from the preceding pyramid level to address the issue of accumulated errors by using local flow consistency, while our CMM can explicitly reduce outliers in the cost volume to improve the accuracy of optical flow estimation by using a flexible and efficient sparse point-based scheme. In addition, the best performance is achieved by integrating our FMM and CMM, which demonstrates the effectiveness of our proposed modules.

D Superiority of Our AMFlow over Other Network Architectures

To further demonstrate the superiority of our AMFlow over other network architectures, we compare the performance of different combinations of unsupervised network architectures and unsupervised training strategies. The results are shown in Table 6. From rows a)–d), we can observe that for each existing unsupervised approach, the performance can be significantly improved when the network architecture is changed from the original one to our AMFlow, which strongly demonstrates the effectiveness of our architecture. The reason why our AMFlow performs better is that it can address the issues of accumulated errors and reduce outliers in the cost volume to improve the optical flow accuracy by using our FMMs and CMMs. Moreover, from row e), we can see that, compared with other network architectures, our AMFlow achieves the best performance when equipped with the same training strategy, *i.e.*, our co-teaching strategy, which further demonstrates the superiority of our AMFlow over other network architectures.

Table 3: A glossary of notations used in the paper.

Notation	Meaning
Section 3.1	
$\mathbf{I}_1$ and $\mathbf{I}_2$	The input images
$\mathbf{I}_1^l$ and $\mathbf{I}_2^l$	The downsampled input images at level l
$\mathbf{x}_1^l$ and $\mathbf{x}_2^l$	The feature maps of input images at level l
$\mathbf{F}_{12}^l$	The forward flow estimation at level l
$\bar{\mathbf{F}}_{12}^{l+1}$	The upsampled forward flow estimation at level $l + 1$
$\hat{\mathbf{F}}_{12}^{l+1}$	The modulated forward flow generated via our FMM at level $l + 1$
$\mathbf{M}^l$	The confidence map used in our FMM at level l
$\mathbf{C}_s^l$	The self-cost volume used in our FMM at level l
$\mathbf{D}^l$	The displacement map used in our FMM at level l
$\mathbf{C}^l$	The cost volume at level l
$\hat{\mathbf{C}}^l$	The modulated cost volume generated via our CMM at level l
Section 3.2 and 3.3	
$\mathbf{I}_1$ and $\mathbf{I}_2$	The input images
$\mathbf{F}_{12}$	The forward flow estimation
$\mathbf{F}_{21}$	The backward flow estimation
$\mathbf{O}_{12}$	The occlusion map
$\mathbf{T}_{\theta}^{\text{img}}, \mathbf{T}_{\theta}^{\text{flo}}$ and $\mathbf{T}_{\theta}^{\text{occ}}$	The transformations employed on $\mathbf{I}_1, \mathbf{I}_2, \mathbf{F}_{12}$ and $\mathbf{O}_{12}$, respectively [13]
$\tilde{\mathbf{I}}_1, \tilde{\mathbf{I}}_2, \tilde{\mathbf{F}}_{12}$ and $\tilde{\mathbf{O}}_{12}$	The samples augmented via the above-mentioned transformations
$\tilde{\mathbf{F}}_{12}^*$	The forward flow prediction based on $\tilde{\mathbf{I}}_1$ and $\tilde{\mathbf{I}}_2$

Table 4: AEPE (px) results of our CoT-AMFlow with different T_k and τ in the proposed co-teaching strategy. The best result is shown in bold font.

	$\tau = 0.70$	$\tau = 0.75$	$\tau = 0.80$	$\tau = 0.85$	$\tau = 0.90$
Constant τ	4.31	4.10	3.95	4.34	4.89
$T_k = 0.05 \cdot T_{\max}$	4.22	3.98	3.83	4.16	4.65
$T_k = 0.10 \cdot T_{\max}$	4.27	4.05	3.79	4.02	4.64
$T_k = 0.15 \cdot T_{\max}$	4.29	3.92	3.85	4.13	4.51

Table 5: AEPE (px) results of variants of our CoT-AMFlow with some of the proposed modules disabled, where ‘‘All’’ denotes the AEPE over all pixels, and ‘‘s0 – 10’’, ‘‘s10 – 40’’ and ‘‘s40+’’ denote the AEPE over pixels that move less than 10 pixels, between 10 and 40 pixels and more than 40 pixels, respectively. Best results are shown in bold font.

FMM	CMM	All	s0 – 10	s10 – 40	s40+
–	–	4.73	0.82	2.46	29.75
✓	–	4.12	0.79	2.32	25.12
–	✓	4.23	0.73	2.23	26.49
✓	✓	3.79	0.76	2.07	23.10

E Superiority of Our Co-Teaching Strategy over Other Strategies for Unsupervised Training

From columns 1)–4) in Table 6, we can observe that for each existing unsupervised approach, the performance can be significantly improved when the training strategy is changed from the original one to our co-teaching strategy, which strongly demonstrates the effectiveness of our strategy. The reason why our co-teaching strategy performs better is that it can improve robustness against outliers for unsupervised optical flow estimation by employing two networks to teach each other about challenging regions simultaneously. Moreover, from column 5), we can see that, compared with

Table 6: AEPE (px) results of different combinations of unsupervised network architectures and unsupervised training strategies. Note that XXXNet and XXXStrat denote the corresponding network architecture and unsupervised training strategy used in XXX, respectively. [†] indicates the network using more than two frames. The best result is shown in bold font.

Network \ Strategy	1)UnFlow-Net [9]	2)DDFlow-Net [11]	3)SelFlow-Net [†] [12]	4)ARFlow-Net [13]	5)AMFlow (Ours)
a) UnFlowStrat [9]	8.87	–	–	–	6.61
b) DDFlowStrat [11]	–	5.95	–	–	5.59
c) SelFlowStrat [12]	–	–	5.22	–	4.98
d) ARFlowStrat [13]	–	–	–	4.67	4.36
e) Co-Teaching (Ours)	5.65	4.73	3.94	4.29	3.79

other training strategies, our co-teaching strategy achieves the best performance when employed in the same network architecture, *i.e.*, our AMFlow, which further demonstrates the superiority of our co-teaching strategy over other strategies for unsupervised training.