

# Semi-Supervised Learning of Optical Flow by Flow Supervisor

Woobin Im, Sebin Lee, and Sung-Eui Yoon

Korea Advanced Institute of Science and Technology (KAIST),  
{iwbn,seb.lee,sungeui}@kaist.ac.kr

**Abstract.** A training pipeline for optical flow CNNs consists of a pre-training stage on a synthetic dataset followed by a fine tuning stage on a target dataset. However, obtaining ground truth flows from a target video requires a tremendous effort. This paper proposes a practical fine tuning method to adapt a pretrained model to a target dataset without ground truth flows, which has not been explored extensively. Specifically, we propose a flow supervisor for self-supervision, which consists of parameter separation and a student output connection. This design is aimed at stable convergence and better accuracy over conventional self-supervision methods which are unstable on the fine tuning task. Experimental results show the effectiveness of our method compared to different self-supervision methods for semi-supervised learning. In addition, we achieve meaningful improvements over state-of-the-art optical flow models on Sintel and KITTI benchmarks by exploiting additional unlabeled datasets. Code is available at <https://github.com/iwbn/flow-supervisor>.

## 1 Introduction

Optical flow describes the pixel-level displacement in two images, and is a fundamental step for various motion understanding tasks in computer vision. Recently, supervised deep learning methods have shown remarkable performance in terms of overcoming challenges – such as motion blur, change of brightness and color, deformation, and occlusion – and predicting more accurate flows. The key to success is end-to-end learning from large-scale data. For optical flow learning, large-scale datasets have been released [7,14,4] and deep architectures have been advanced [7,32,34].

Building a good optical flow model on a target dataset is critical; most training data are synthetic, and it requires tremendous effort to label random video frames by pixel-wise dense correspondences. Thus, to obtain a good model on a target dataset, synthetic training set generation [31,24] and GAN-based adaptation [20,36] have been studied. In addition, unsupervised loss functions [25,30] – used without ground truth – have been proposed. However, generating a synthetic dataset for a target domain is often computationally expensive or confined to a specific domain. Moreover, unsupervised losses do not reach the state-of-the-arts, compared to supervised methods. Therefore, there have been needs for a simpler, general, and high-performance method to build a better model on a target dataset.

In this paper, we propose a fine tuning strategy for semi-supervised optical flow learning, which helps to build a better model on unlabeled or partly-labeled target datasets. Fig. 1 demonstrates the concept of our approach. Our method is a fine tuning method, where the pretrained network is adapted to the unlabeled target data. In the fine tuning stage, we use a labeled dataset with an unlabeled target dataset, further reducing errors on the target dataset.

To build our method, we investigate unsupervised and self-supervised approaches, where a network learns optical flow by unlabeled samples. However, the existing unsupervised loss does not show better performance in the fine tuning stage (Fig. 1). Moreover, self-supervision methods tend to show highly unstable behavior or lower performance in the fine tuning stage. To address the issue, we propose our flow supervisor with two strategies: the parameter separation and passing student outputs, which are effective for higher performance and stable semi-supervised learning. As shown in Fig. 1, our fine tuning method clearly reduces the error of the pretrained model, even without a label of the target dataset.

To summarize, we propose a semi-supervised fine tuning strategy to improve an optical flow network on a target dataset, which has not been explored extensively. Our strategy is distinguished by the flow supervisor module, designed with the parameter separation and passing student outputs. We show the effectiveness of our method by comparing it with alternative self-supervision methods, and confirm that our approach stabilizes the learning process and results in better accuracy. In addition, we test our method on Sintel and KITTI benchmarks and achieve meaningful improvements over the state-of-the-arts by exploiting additional unlabeled data.

## 2 Related Work

**Supervised optical flow** has been studied with the development of datasets for optical flow learning [7,14,31] and the advances of deep network architectures [7,32,34]. Due to the high annotation cost and label ambiguity in a raw video, synthetic datasets have been made, where optical flow fields are generated

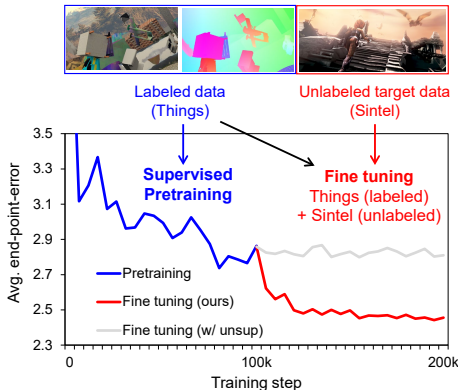


Fig. 1: **End-point-error on Sintel.** Our method is used to adapt a pretrained model to a target domain without a target domain label; it is designed to overcome an unstable convergence and low accuracy in traditional methods. For instance, our method outperforms the unsupervised loss (Eq. 2) in fine tuning, which makes our method favorably better

together with images [24,7,14,31]. Along with the datasets, network architectures for optical flow have been significantly improved by the cost-volume [7,32,37], warping [14,32], and refinement scheme [13,34]. Although generalization has been improved thanks to the synthetic datasets and the network architectures, it is still difficult to achieve better performance while being blind to a target domain [27].

**Unsupervised optical flow** is another stream of optical flow research, where optical flow is learned without an expensive labeling or label generation process [25,15,18,30]. With the advanced deep architectures from supervised optical flow, unsupervised optical flow studies have focused on designing loss functions. Previous work has mainly focused on the fully unsupervised setting, while we explore semi-supervised optical flow where labeled data is accessible in addition to unlabeled target domain data.

**Semi-supervised optical flow** has been studied to utilize unlabeled target domain data with existing labeled synthetic data [20]. The experimental setting – similar to unsupervised domain adaptation [9] – is designed since it is relatively easier to exploit synthetic training datasets than annotating images of a target dataset. A simple baseline method for a target domain would be using a traditional unsupervised loss [25] and supervised learning, which, unfortunately, has shown inferior performance than the supervision-only training [20,27] (Table 1). Thus, reducing the domain gap [20] and stabilizing unsupervised loss gradients [27] have been proposed.

In this work, we introduce the flow supervisor, which consists of the separate parameters and the student output connection. We found our design scheme is superior to the baseline designs in terms of training stability and performance in the semi-supervised setting.

**Knowledge distillation and self-supervision** in neural networks are proposed to train a network under the guidance of a teacher network [12,29] or itself [6,10]. Interestingly, the technique can build a better student network even when the same network architecture is used for both student and teacher, i.e., self-distillation [12,29]. In the optical flow field, knowledge distillation and self-supervision have been studied actively in the context of unsupervised optical flow [22,21,30]. Generally, these methods can be interpreted as learning using privileged information [35], in that a student usually sees a limited view, e.g., cropped images, while a teacher is given a privileged view, e.g., full images. In this work, we investigate the effectiveness of self-supervision in terms of semi-supervised optical flow, which has not been explored in previous work. In addition, we found that the traditional self-supervision method for optical flow [30] tends to make a loss diverge in the semi-supervised setting.

Thus, we propose a novel self-supervision method to ameliorate the unstable convergence; we found that the parameter separation of a teacher model and passing student output are the key components to make the training stable. By applying our method, a network can successfully exploit the unlabeled target data in the semi-supervised setting. In addition, we show that our method can address the lack of a target labeled dataset, e.g., 200 labeled pairs for KITTI, by the ability to utilize unlabeled samples.

### 3 Approach

#### 3.1 Preliminaries on Deep Optical Flow

Deep optical flow estimation is defined with an optical flow estimator  $f_\theta(\mathbf{x}_1, \mathbf{x}_2)$ , which predicts optical flow  $\hat{\mathbf{y}} \in \mathbb{R}^{H \times W \times 2}$  from two images  $\mathbf{x}_1, \mathbf{x}_2 \in \mathbb{R}^{H \times W \times C}$ , where  $H$  is height,  $W$  is width, and  $C$  is channel.

In supervised learning, we train  $f_\theta$  by minimizing the supervised loss:

$$\mathcal{L}_{\text{sup}}(\theta) = \mathbb{E}_{(\mathbf{x}_1, \mathbf{x}_2, \mathbf{y}) \sim p_s} [\ell_{\text{sup}}(f_\theta(\mathbf{x}_1, \mathbf{x}_2), \mathbf{y})], \quad (1)$$

where  $p_s$  is the labeled data distribution,  $\mathbf{y}$  is the ground truth optical flow, and  $\ell_{\text{sup}}$  is  $L_1$ ,  $L_2$  [32] or the generalized Charbonnier loss [25].

On the other hand, an unsupervised method defines a loss function with a differentiable target function  $\ell_{\text{unsup}}(\cdot)$  which can be computed without a ground truth  $\mathbf{y}$ , resulting in the unsupervised loss:

$$\mathcal{L}_{\text{unsup}}(\theta) = \mathbb{E}_{(\mathbf{x}_1, \mathbf{x}_2) \sim p_u} [\ell_{\text{unsup}}(f_\theta(\mathbf{x}_1, \mathbf{x}_2), \mathbf{x}_1, \mathbf{x}_2)], \quad (2)$$

where  $p_u$  is an unlabeled data distribution. Most commonly,  $\ell_{\text{unsup}}$  is defined with the photometric loss  $\ell_{\text{photo}} = \rho(\text{warp}(\mathbf{x}_2, \hat{\mathbf{y}}) - \mathbf{x}_1)$  where  $\rho$  is the Charbonnier loss [25] and  $\text{warp}(\cdot)$  is the differentiable backward warping operation [16].

#### 3.2 Problem Definition and Background

We train our model on labeled data and unlabeled data, which is similar to experimental settings that appear in [20,27]. For instance, FlyingThings3D (rendered scene) and KITTI (driving scene) can be considered as labeled and unlabeled datasets, respectively. We aim to build a high-performance model on a target domain, with only a labeled synthetic dataset and an unlabeled target dataset. This is a practical scenario since a synthetic dataset is relatively inexpensive and publicly available, whereas specific target domain data is rarely annotated.

We focus on designing a self-supervision method with a stable convergence and better accuracy, since it is not trivial to adopt unsupervised losses and self-supervision for semi-supervised learning. First, unsupervised loss functions often lead a network to a worse local minimum when it is naively fine tuned from supervised (pretrained) weights. Second, traditional self-supervision strategies for semi-supervised learning do not converge well. Thus, we propose a simple and effective practice for semi-supervised learning, where we utilize synthetic datasets for better performance on a target dataset. By our strategy, a network can successfully exploit the unlabeled target data for better fine tuning performance. Detailed method is discussed in the next section.

#### 3.3 Flow Supervisor

Our design for semi-supervised optical flow learning is based on self-supervised learning where a student network learns from a pseudo-label predicted by a

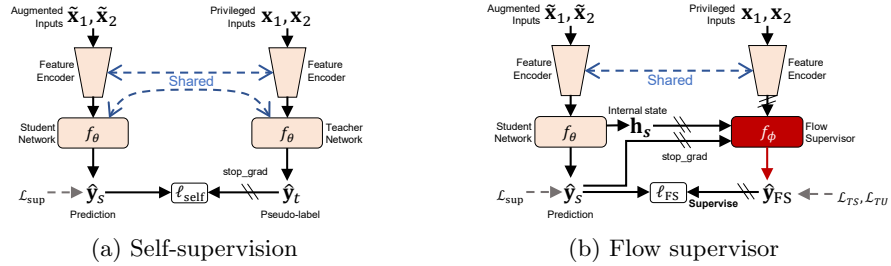


Fig. 2: (a) **Self-supervision** for optical flow is configured with a teacher network which is given privileged images as an input, i.e., full images before cropping. (b) **Flow supervisor** reviews the student flow  $\hat{\mathbf{y}}_s$  and outputs the pseudo-label  $\hat{\mathbf{y}}_{\text{FS}}$  to supervise the student without ground truth flows. We use the separate flow supervisor with parameter  $\phi$ , which improves the stability and accuracy

teacher network. This concept has been explored in the optical flow field in terms of unsupervised learning, which is not directly applicable in the semi-supervised case due to unstable convergence.

Thus, we introduce two distinctive design schemes compared to the existing self-supervision techniques for optical flow, depicted in Fig. 2. First, we introduce a supervisor parameter  $\phi$ , distinguished from the student parameter  $\theta$ , which learns to supervise the student network. Second, we add a connection from the student network to the supervisor network, which enables the supervisor network to learn conditional knowledge, i.e.,  $P(\mathbf{y}|\hat{\mathbf{y}}_s)$ , instead of predicting from scratch.

We define the student network  $f_\theta$  and the flow supervisor network  $f_\phi$ , as shown in Fig. 2b. The student network  $f_\theta$  includes a feature encoder and a flow decoder; for simplicity, we abstract the feature encoder and decoder parameters with  $\theta$ . The flow decoder has the internal feature  $\mathbf{h}_s$ , and outputs the predicted flow  $\hat{\mathbf{y}}_s$ . The student network is the optical flow network used for inference, whereas the supervisor network is only used for training; this results in no additional computational cost for the testing time. In training time, we use the flow supervisor (FS) loss function  $\mathcal{L}_{\text{FS}}$  to supervise the student flow  $\hat{\mathbf{y}}_s$  with the teacher flow  $\hat{\mathbf{y}}_{\text{FS}}$ :

$$\mathcal{L}_{\text{FS}}(\theta) = \mathbb{E}_{d_s \sim p_s, d_u \sim p_u} [\ell_{\text{sup}}(d_s) + \alpha \ell_{\text{FS}}(d_u)], \quad (3)$$

where  $\ell_{\text{FS}}(\cdot) = \rho(\hat{\mathbf{y}}_{\text{FS}} - \hat{\mathbf{y}}_s)$ ,  $\alpha$  is a hyper-parameter weight and  $(d_s, d_u)$  are sampled data from  $(p_s, p_u)$ . We use loss function  $\mathcal{L}_{\text{FS}}$  for the student network to learn from both labeled and unlabeled data.

*Separate parameters.* We empirically observed that the plain self-supervision (Fig. 2a) leads to divergence in semi-supervised optical flow learning. This undesirable behavior is also observed in siamese self-supervised learning, where preventing undesirable equilibria is important [10,6]. As our solution, we have the separate module, i.e., the flow supervisor to prevent the unstable learning

behavior. Our flow supervisor is related to the predictor module in self-supervised learning work [6], which also prevents the unstable training. We also compare our design with the mean-teacher [33,10] with exponential moving average (EMA), and a fixed teacher [39] in Fig. 3, since these designs have been widely adopted in semi-supervised learning.

*Passing student outputs.* We design our supervisor model to have input nodes for student outputs. Thus, the teacher output flow is conditioned by the student output. Specifically, our teacher model includes a residual function  $\Delta f_\phi(\cdot)$ , such that  $f_\phi(\mathbf{x}, \hat{\mathbf{y}}_s, \mathbf{h}_s) = \Delta f_\phi(\cdot) + \hat{\mathbf{y}}_s$ . We realize this concept with the residual flow decoder in the RAFT [34] architecture.

In residual learning, it has been believed and shown that learning a residual function  $\Delta f(\mathbf{x}) = f(\mathbf{x}) - \mathbf{x}$ , instead of the original function  $f(\mathbf{x})$ , is better for deeper inference [11] or domain discrepancy modeling [23]. With residual teacher function  $\Delta f_\phi$ , the flow supervisor loss is reformulated as:

$$\ell_{\text{FS}} = \rho(\hat{\mathbf{y}}_{\text{FS}} - \hat{\mathbf{y}}_s) = \rho(f_\phi(\hat{\mathbf{y}}_s) - \hat{\mathbf{y}}_s) = \rho(\Delta f_\phi(\hat{\mathbf{y}}_s)). \quad (4)$$

*Relation of  $f_\phi(\cdot)$  to a meta learner.* Meta-learning [1,8] defines how to learn by learning an update rule  $\Delta\theta^t$  within the parameter update  $\theta^{t+1} \leftarrow \theta^t + \Delta\theta^t$ . The residual function  $\Delta f_\phi(\dots, \hat{\mathbf{y}}_s)$  is regarded as a meta learner predicting update rule  $\Delta\theta^t$ , conditioned by student predictions. Assuming  $\rho$  is the square function, learning by  $\mathcal{L}_{\text{FS}}$  is equivalent to using the update rule  $\Delta\theta^t = -2 \frac{\partial \hat{\mathbf{y}}_s^T}{\partial \theta^t} \cdot \Delta f_\phi(\hat{\mathbf{y}}_s)$ ; where  $\phi$  is the parameter of our flow supervisor. That is, the learning rule is learned by the supervisor parameter  $\phi$  to supervise the student parameter  $\theta$ .

### 3.4 Supervision

Learning supervisor parameters  $\phi$  is important for supervising the student network. Basically, the supervisor learns to maximize the likelihood, e.g.,  $\log P(\mathbf{y} - \hat{\mathbf{y}}_s | \mathbf{x}_1, \mathbf{x}_2, \hat{\mathbf{y}}_s, \phi)$ , where the student output  $\hat{\mathbf{y}}_s$  is inferred from an augmented input  $\tilde{\mathbf{x}}_1, \tilde{\mathbf{x}}_2$ . First, it is natural to give the conditional knowledge from the labeled data  $p_s$ . For this purpose, we minimize  $\mathcal{L}_{\text{TS}}$  – which stands for supervised teacher loss – to train the supervisor:

$$\mathcal{L}_{\text{TS}}(\phi) = \mathbb{E}_{(\mathbf{x}_1, \mathbf{x}_2, \mathbf{y}) \sim p_s} [\ell_{\text{sup}}(f_\phi(\mathbf{x}_1, \mathbf{x}_2, f_\theta(\tilde{\mathbf{x}}_1, \tilde{\mathbf{x}}_2)), \mathbf{y})]. \quad (5)$$

In addition to using the labeled data, we found that if a labeled dataset and an unlabeled dataset are distant, e.g., Things  $\leftrightarrow$  KITTI, using the unsupervised loss is especially effective. The unsupervised teacher loss  $\mathcal{L}_{\text{TU}}$  on unlabeled data  $p_u$  is defined by:

$$\mathcal{L}_{\text{TU}}(\phi) = \mathbb{E}_{(\mathbf{x}_1, \mathbf{x}_2) \sim p_u} [\ell_{\text{unsup}}(f_\phi(\mathbf{x}_1, \mathbf{x}_2, \hat{\mathbf{y}}_s), \mathbf{x}_1, \mathbf{x}_2)]. \quad (6)$$

In the pretraining stage, we train the student model from scratch using the supervised loss  $\mathcal{L}_{\text{sup}}$ , resulting in a pretrained weight  $\theta$ . In the fine tuning stage,

we initialize  $\phi$  with  $\theta$  and jointly optimize  $\theta$  and  $\phi$  on labeled and unlabeled datasets. Formally, we use stochastic gradient descent to minimize

$$\mathcal{L}(\theta, \phi) = \mathcal{L}_{\text{FS}}(\theta) + \lambda_{\text{TS}}\mathcal{L}_{\text{TS}}(\phi) + \lambda_{\text{TU}}\mathcal{L}_{\text{TU}}(\phi), \quad (7)$$

with hyper-parameter loss weights  $\lambda_{\text{TS}}$  and  $\lambda_{\text{TU}}$ .

## 4 Experiments

### 4.1 Experimental setup

*Pretraining.* Our pretraining stage follows the original RAFT [34]. We pretrain our student network with FlyingChairs [7] and FlyingThings3D [14] with random cropping, scaling, color jittering, and block erasing. The pretraining scheme includes 100k steps for FlyingChairs and additional 100k steps for FlyingThings3D, with the same learning rate schedule of RAFT, on which we base our network. For the supervised loss (Eq. 1), we use  $L_1$  loss.

*Flow supervisor.* The flow supervisor is implemented with the GRU update module of RAFT, which performs an iterative refinement process with the output flow of the previous step. At the start of the semi-supervised training phase, we initialize the supervisor model with the pretrained weights of the GRU update module. To match the student prediction from cropped inputs to the supervisor module with full resolution, i.e., privileged, we pad the student predictions with zero. In a training and testing phase, we use 12 iterations for both the student and supervisor GRUs.

*Semi-supervised dataset.* To compare semi-supervised learning performance on Sintel [4] and KITTI [26], we follow the protocol [18], which utilizes the unlabeled portion, i.e., testing set, of each dataset for training; the difference is that we use a labeled dataset, e.g., FlyingThings3D, into training. Specifically, we use the unlabeled portion of Sintel for fine tuning on Sintel and unlabeled KITTI multiview dataset for fine tuning on KITTI; then the labeled splits are used for evaluation.

*Optimization.* We initialize our student and supervisor models with the pretrained weights, then minimize the joint loss (Eq. 7) for 100k steps. We use  $\alpha = 1.0$ ,  $\lambda_{\text{TS}} = 1.0$  and  $\lambda_{\text{TU}} = 0$  by default and  $\lambda_{\text{TU}} = 0.01$  for the Things + KITTI setting, unless otherwise stated. Detailed optimization settings and hyper-parameters are provided in the supplementary material and the code.

### 4.2 Empirical Study

**Comparison to Semi-Supervised Baselines.** In Table 1 and Fig. 3, we perform an empirical study, comparing ours with several baselines for semi-supervised learning. In the experiments, we use FlyingThings3D as the labeled data  $p_s$  and each target dataset as the unlabeled data  $p_u$ .

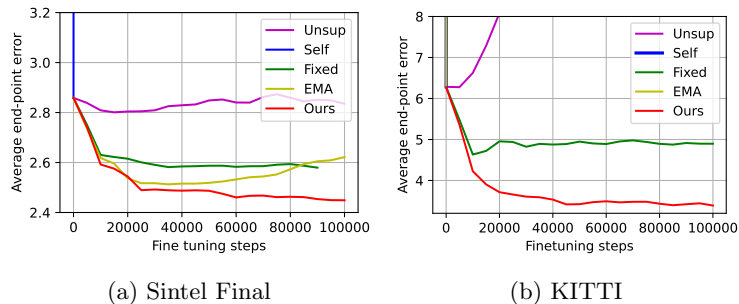


Fig. 3: **Plots comparing finetuning stage.** We use FlyingThings3D as labeled data, and each target dataset as an unlabeled data. The EPEs are measured on unseen portion of data

*Unsupervised loss.* An unsupervised loss is designed to learn optical flow without labels. Exploring the unsupervised loss is a good start for the research since we could expect a meaningful supervision signal from unlabeled images.

In the second row of Table 1, we report results by the unsupervised loss. We use the loss function in Eq. 2, which includes the census transform, full-image warping, smoothness prior, and occlusion handling as used in the prior work [30]. Unfortunately, the unsupervised loss function does not give a meaningful supervision signal when

jointly used with the supervised loss, resulting in a degenerated EPE on KITTI (5.79  $\rightarrow$  9.21). Interestingly, applying the identical loss to our supervisor, i.e.,  $\mathcal{L}_{TU}$ , results in a better EPE on KITTI (5.79  $\rightarrow$  3.35), see also Table 3.

*Self- and teacher-supervision.* We summarize the results of the self-supervision and teacher-based methods in Table 1 and Fig. 3. The baseline models include the plain self-supervision (Self), the fixed teacher (Fixed), and the EMA teacher (EMA). In the plain self-supervision, we use the plain siamese networks (Fig. 2a). The fixed teacher [39] and the EMA teacher [33] are inspired by the existing literature. The self-supervision loss used in the experiments is defined by:

$$\ell_{\text{self}} = \rho[f_{\theta}(\tilde{\mathbf{x}}_1, \tilde{\mathbf{x}}_2) - \text{stop\_grad}(f_t(\mathbf{x}_1, \mathbf{x}_2))], \quad (8)$$

where  $f_t$  is the teacher network of each baseline:  $t = \theta$  for plain self-supervision,  $t = \text{EMA}(\theta)$  for mean teacher, and  $t = \theta_{\text{pretrained}}$  for fixed teacher. The plain

Table 1: **Comparison to semi-supervised baselines.** Semi-supervised baselines and ours are compared, as well as the supervised loss only (Sup) result. We use widely-used metrics in optical flow: end-point error (EPE) and ratio of erroneous pixels (Fl)

| Method                 | Sintel      |             | KITTI       |              |
|------------------------|-------------|-------------|-------------|--------------|
|                        | Clean       | Final EPE   | Fl-all (%)  |              |
| Sup                    | 1.46        | 2.80        | 5.79        | 18.79        |
| Sup + Unsup            | 1.47        | 2.73        | 9.21        | 16.95        |
| Sup + Self             | diverged    |             |             |              |
| Sup + EMA              | 1.40        | 2.63        | diverged    |              |
| Sup + Fixed            | 1.32        | 2.58        | 4.91        | 15.92        |
| <b>Sup + FS (Ours)</b> | <b>1.30</b> | <b>2.46</b> | <b>3.35</b> | <b>11.12</b> |



Table 2: **Comparison to optical flow approaches.** We report a percentage of improvement over each baseline in the parentheses. We mark used datasets: FlyingChairs (C), FlyingThings (T), unlabeled KITTI (K), and AutoFlow (A) [31]

(a) **Comparison to SemiFlowGAN [20]**    (b) **Comparison to AutoFlow [31]**

| Data Method    | KITTI                |                       | Method                       | Sintel      |             | KITTI       |              |
|----------------|----------------------|-----------------------|------------------------------|-------------|-------------|-------------|--------------|
|                | EPE                  | Fl-all (%)            |                              | Clean       | Final       | EPE         | Fl-all (%)   |
| C SFGAN [20]   | 17.19                | 40.82                 | RAFT <sup>†</sup> (C+T) [31] | 1.68        | 2.80        | 5.92        | -            |
| C+K SFGAN [20] | 16.02 (-6.8%)        | 38.77 (-5.0%)         | RAFT (C+T)                   | 1.46        | 2.80        | 5.79        | 18.79        |
| C+T RAFT       | 5.79                 | 18.79                 | RAFT <sup>†</sup> (A) [31]   | 1.95        | 2.57        | 4.23        | -            |
| T+K RAFT + FS  | <b>3.35 (-42.1%)</b> | <b>11.12 (-40.8%)</b> | <b>RAFT + FS (Ours)</b>      | <b>1.30</b> | <b>2.46</b> | <b>3.35</b> | <b>11.12</b> |

† model implemented by [31]

self-supervision (Self) quickly diverges during the early fine tuning step for both Sintel and KITTI. In the EMA teacher (EMA) results, more stable convergence is observed for Sintel, not KITTI. We speculate that this unstable convergence is caused by the domain gap between the unlabeled data and the labeled data; there exists a wider domain difference between FlyingThings3D ( $p_s$ ) and KITTI ( $p_u$ ) than the difference between FlyingThings3D ( $p_s$ ) and Sintel ( $p_u$ ), since KITTI is a real-world dataset, while Things and Sintel are both three-dimensional rendered datasets. When the fixed teacher (Fixed) is used, we can observe more stable learning. Our method (FS), enabling the supervisor learning along with the student, shows superior semi-supervised performance on both datasets than the EMA and the fixed teacher. We further analyze our method in the following sections.

**Comparison to SemiFlowGAN [20].** We compare our method with the existing semi-supervised optical flow method [20] in Table 2a. SemiFlowGAN uses a domain adaptation-like approach by matching distributions of the error maps from each domain. Compared to SemiFlowGAN, our approach gives more direct supervision to optical flow predictions from the supervisor model. Here, we evaluate how much a supervised-only baseline, i.e., trained w/o KITTI, is able to be improved by exploiting the unlabeled target dataset, i.e. trained w/ KITTI. We can clearly observe much larger performance improvement (-6.8% vs. -42.1%) when our method is used.

**Comparison to AutoFlow [31].** In Table 2b, we compare ours with AutoFlow, where our method shows better EPEs on both Sintel and KITTI. AutoFlow is devised to train a better neural network on target datasets by learning to generate data, as opposed to ours using semi-supervised learning. For instance, the AutoFlow dataset, whose generator is optimized on Sintel dataset shows superiority over a traditional synthetic dataset, e.g., FlyingThings, in generalization on unseen target domains. On the other hand, our strategy is to utilize a synthetic dataset and an unlabeled target domain dataset by the supervision of the flow supervisor; it boosts performance on the target domain without labels.

**Ablation Study.** In Table 3, we compare ours with several alternatives.

*Parameter separation.* Self-supervision with the shared network (see Self in Fig. 3) is unsuitable for semi-supervised optical flow learning. On the other hand, having the separate parameters for the supervisor network effectively prevents the divergence. This separation can be viewed as the predictor strategy in the self-supervised learning context [6]. We also analyze the separate network as a meta-learner, which learns knowledge specifically for supervision. We have shortly discussed this aspect in Sec. 3.3.

*Passing student outputs ( $\hat{\mathbf{y}}_s, \mathbf{h}_s$ ).* Our architecture passes the student output flow and internal state to the supervisor network to enable the supervisor to learn the residual function, as described in Eq. 4. In row 5 in Table 3, we show the results when we do not pass student outputs ( $\hat{\mathbf{y}}_s, \mathbf{h}_s$ ) to the teacher. Interestingly, the Sintel results remain the same even without passing student outputs, while the KITTI results benefit from passing student outputs. We ablate the residual connection from the network, while still passing student outputs to the supervisor (w/o res). In this case (w/o res), we can observe the better KITTI results over the no connection case (off), but worse than our full design. These results indicate that conditioning supervisor with the student helps find residual function  $\Delta f_\phi(\hat{\mathbf{y}}_s) \approx \mathbf{y} - \hat{\mathbf{y}}_s$ , especially on distant domains, i.e., Things  $\leftrightarrow$  KITTI. Overall, passing with residual connection (full) performs consistently better.

*Shared encoder.* Shared encoder design results in a better flow accuracy, as shown in Table 3. Our method uses the shared encoder design (Fig. 2b). Separating the encoder results in worse EPEs for Sintel and KITTI. Not only the accuracy, but it also results in a less efficient training pipeline due to the increased number of parameters by the encoder.

*Teacher loss type.* We propose to train the supervisor network with labeled and unlabeled data. Supervised teacher (TS) loss utilizes the labeled dataset for supervisor learning. In both datasets, TS improves the baseline performance. We additionally apply the unsupervised teacher loss (+TU), which results in better EPE in KITTI but is ineffective for Sintel. This is related to the results of pure unsupervised learning [30], where the loss is shown to be more effective on KITTI than on Sintel.

Table 3: **Ablation experiments.** We underline the final settings

| Experiment             | Sintel       |             | KITTI       |             |              |
|------------------------|--------------|-------------|-------------|-------------|--------------|
|                        | Clean        | Final       | EPE         | Fl-all      |              |
| Parameter separation   | <u>on</u>    | <b>1.30</b> | <b>2.46</b> | <b>3.35</b> | <b>11.12</b> |
|                        | off          | diverged    |             | diverged    |              |
| Passing student output | <u>full</u>  | <b>1.30</b> | <b>2.46</b> | <b>3.35</b> | <b>11.12</b> |
|                        | w/o res      | 1.34        | 2.50        | 6.16        | 12.34        |
|                        | off          | <b>1.30</b> | <b>2.46</b> | 6.68        | 13.36        |
| Shared encoder         | <u>on</u>    | <b>1.30</b> | <b>2.46</b> | <b>3.35</b> | <b>11.12</b> |
|                        | off          | 1.33        | 2.58        | 3.80        | 12.03        |
| Teacher loss type      | <u>TS</u>    | <b>1.30</b> | <b>2.46</b> | 4.69        | 14.48        |
|                        | +TU          | 1.55        | 2.80        | <b>3.35</b> | <b>11.12</b> |
| Teacher input          | <u>clean</u> | <b>1.30</b> | <b>2.46</b> | <b>3.35</b> | <b>11.12</b> |
|                        | aug          | 1.33        | 2.56        | 4.17        | 11.55        |

**Virtual KITTI.** With our method, we can utilize a virtual driving dataset, VKITTI [5], for training a better model for the real KITTI dataset. In Table 4, we show results obtained by VKITTI. In the supervised case, our base network (RAFT) shows 3.64 EPE on KITTI, and SeperableFlow [38] shows 2.60 EPE. In the semi-supervised case, our method leverages the unlabeled KITTI multiview set additional to VKITTI, and we train the RAFT network for 50k steps; the resulting network performs better on KITTI.

Table 4: **KITTI results of models trained on VKITTI [5]**

| Metric     | Supervised |              | Semi-supervised |
|------------|------------|--------------|-----------------|
|            | RAFT       | SepFlow [38] | RAFT+FS (ours)  |
| EPE        | 3.64       | 2.60         | <b>2.39</b>     |
| Fl-all (%) | 8.78       | 7.74         | <b>7.63</b>     |

### Supervisor vs. student.

In Fig. 4, we show the performance of the student and the supervisor networks by training steps. We use clean inputs for both the student and the supervisor networks for evaluation. Interestingly, we observe that the student is better than the supervisor during fine tuning. In addition, we could observe EPEs of the student and the supervisor are correlated, and they are both improved by our semi-supervised training. In knowledge distillation, we can observe similar behavior [29], where a student model shows better validation accuracy than a teacher model.




Fig. 4: **Supervisor vs. student EPEs** during semi-supervised fine tuning. EPEs measured on unseen portion of data

### 4.3 Qualitative results

We provide qualitative results on KITTI (Fig. 5, 6) and Sintel (Fig. 7). Since the KITTI dataset includes sparse ground truth flows, training by the ground truth supervision often results in incorrect results, especially in boundaries and deformable objects (e.g., person on a bike). Thus, in this case, we can expect better flows by exploiting semi-supervised methods. The Sintel Final dataset includes challenging blur, fog, and lighting conditions as shown in the examples. Interestingly, the semi-supervision without a target dataset label helps improve the challenging regions, when our method is used.

In Fig. 8, we show results on DAVIS dataset [28] by the pretrained models and our fine tuned models; the dataset does not contain any optical flow ground truth. From the training set with 3,455 frames, we utilize 90% of frames for training and the rest for qualitative evaluation. Even though our fine tuning does not utilize ground truth, we can observe a clear positive effect in several challenging regions. Especially, our method improves blurry regions (e.g., moving hand of the dancer) and object boundaries.



Fig. 5: **Qualitative results on KITTI.** We visualize optical flow predicted by (b) supervised on target dataset and (c) semi-supervised w/o target label. Note that sparse ground-truth (d) is not sufficient to make a clear boundary of objects (marked with arrows), while our method shows better results




Fig. 6: **Qualitative results on KITTI testing samples.** We compare the supervised model (Sup) with the semi-supervised model (Ours). Both models exploit KITTI labels; ours utilizes additional unlabeled KITTI for fine tuning




Fig. 7: **Qualitative results on Sintel.** We visualize optical flow predicted by (b) supervised on FlyingThings3D and (c) semi-supervised without target label. Though ours trained without target labels, it successfully adapts the pretrained model to the target domain. Improved areas are marked by arrows




Fig. 8: **Qualitative results on DAVIS dataset [28].** We fine tune each pre-trained network (Sup) on Davis dataset by our semi-supervised method (Ours). Improved regions are marked with arrows. These optical flows are inferred on unseen portion of the dataset

Table 5: **Comparison to state-of-the-arts.** Data usage is abbreviated to FlyingChairs (C), FlyingThings3D (T), Sintel (S), KITTI (K), HD1K [19] (H), Sintel unlabeled ( $S^u$ ), KITTI unlabeled ( $K^u$ ), and Spring ( $Spg^u$ ). For labeled datasets, we follow the training scheme detailed in each paper. RAFT<sup>tf</sup> is our implementation in TensorFlow

| Labeled data | Unlabeled data | Method                        | Sintel (train) |             | KITTI (train) |              | Sintel (test) |             | KITTI (test) |             |
|--------------|----------------|-------------------------------|----------------|-------------|---------------|--------------|---------------|-------------|--------------|-------------|
|              |                |                               | Clean          | Final       | EPE           | Fl-all (%)   | Clean         | Final       | Fl-all (%)   |             |
| C+T          | -              | RAFT [34]                     | 1.43           | 2.71        | 5.05          | 17.4         | -             | -           | -            | -           |
|              |                | RAFT <sup>tf</sup>            | 1.46           | 2.80        | 5.79          | 18.8         | -             | -           | -            | -           |
|              |                | GMA [17]                      | <b>1.30</b>    | 2.74        | 4.69          | 17.1         | -             | -           | -            | -           |
|              |                | SeparableFlow [38]            | <b>1.30</b>    | <u>2.59</u> | <u>4.60</u>   | <u>15.9</u>  | -             | -           | -            | -           |
|              | $S^u/K^u$      | RAFT <sup>tf</sup> +FS (Ours) | <b>1.30</b>    | <b>2.46</b> | <b>3.35</b>   | <b>11.12</b> | -             | -           | -            | -           |
| C+T+S+K+H    | -              | RAFT [34]                     | (0.77)         | (1.27)      | (0.63)        | (1.5)        | 1.61          | 2.86        | 5.10         |             |
|              |                | GMA [17]                      | (0.62)         | (1.06)      | (0.57)        | (1.2)        | <b>1.39</b>   | <u>2.47</u> | 5.15         |             |
|              |                | SeparableFlow [38]            | (0.69)         | (1.10)      | (0.69)        | (1.60)       | 1.50          | 2.67        | <b>4.64</b>  |             |
|              |                | $Spg^u/K^u$                   | RAFT+FS (Ours) | (0.75)      | (1.29)        | (0.69)       | (1.75)        | 1.65        | 2.79         | <u>4.85</u> |
|              |                | GMA+FS (Ours)                 | (0.63)         | (1.05)      | (0.61)        | (1.47)       | <u>1.43</u>   | <b>2.44</b> | 4.95         |             |

#### 4.4 Comparison to State-of-the-arts

*Experimental settings.* In this experiment we compare our method to the existing supervised methods. The biggest challenge in supervised optical flow is that we do not have ample ground truth flows for our target domains, e.g., 200 pairs labeled in KITTI. Thanks to our semi-supervised method, we can train with related videos for better performance on the target datasets. For GMA, we use  $\alpha = 0.25$  (Sintel) and  $\alpha = 0.05$  (KITTI).

*Dataset configuration.* Results evaluated on the training sets of each dataset are experimented with the setting described in Sec. 4.1. On the other hand, the results on the test splits of the two benchmarks are obtained by justifiable external datasets for semi-supervised training. Though our method is able to utilize the unlabeled portion of Sintel, we bring another external set to avoid the relation to the testing samples, as follows. To assist Sintel performance, we use Spring (abbr. to  $Spg^u$ ) [2], which is the ‘open animated movie’ by Blender Institute, similar to Sintel [3]. In Spring, we use the whole frames (frame no. 1-11,138) without any modification. Additionally, we use Sintel (train) with interval two as unlabeled. For KITTI, we additionally use KITTI multiview dataset, which does not contain ground truth optical flows; we use the training split of the dataset to avoid duplicated scenes with testing samples.

*Results.* In Table 5, we report the C+T result, which is a commonly used protocol to evaluate the generality of models. Since our method is designed to use unlabeled samples, our method exploits each target dataset – which is not overlapped with each evaluation sample – without ground truth. The results indicate we could achieve better accuracy than the supervised-only approaches. Interestingly, our approach performs better than the advanced architectures, i.e., GMA [17] and SeparableFlow [38], in the C+T category.

In ‘C+T+S+K+H’, we report the test results with the external datasets. In the KITTI test result where we have access to 200 labeled samples, using additional samples results in better Fl-all (5.10  $\rightarrow$  4.85, 5.15  $\rightarrow$  4.95). For Sintel, we test on two base networks: RAFT and GMA. In two networks, our method makes improvements on Sintel Final (test). However, our method does not improve accuracy on Sintel Clean set. That is because we use the external Spring dataset, which includes various challenging effects, e.g., blur, as Sintel Final, the student model becomes robust to those effects rather than clean videos. Note that, for Sintel test, our improvement is made by the video of a different domain which is encoded by a lossy compression, like most videos on the web.

#### 4.5 Limitations

A limitation of our approach is that it depends on a supervised baseline, which is sometimes less preferable than an unsupervised approach. A handful of unsupervised optical flow researches have achieved amazing performance improvement. Especially, on KITTI (w/o target label), unsupervised methods have performed better than supervised methods, and the recent performance gap has been widened (EPE: 2.0 vs 5.0). Unfortunately, we observe that our method is not compatible with unsupervised baselines; the semi-supervised fine tuning of SMURF [30] with our method (T+K<sup>u</sup>) results in a worse EPE (2.0  $\rightarrow$  2.5) on KITTI. Thus, one of the future work lies upon developing a fine-tuning method to improve the unsupervised baselines.

Nonetheless, our work contributes to the research by ameliorating the limitation of a supervised method, suffering from worse generalization. A supervised method is often preferable than an unsupervised one for higher performance even without target labels; we show ours can be used in such cases. For example, on Sintel Final, ours shows better EPE (2.46) than the supervised one [34] (2.71) and the unsupervised method [30] (2.80).

## 5 Conclusion

We have presented a self-supervision strategy for semi-supervised optical flow, which is simple, yet effective. Our flow supervisor module supervises a student model, which is effective in the semi-supervised optical flow setting where we have no or few samples in a target domain. Our method outperforms various self-supervised baselines, shown by the empirical study. In addition, we show that our semi-supervised method can improve the state-of-the-art supervised models by exploiting additional unlabeled datasets.

**Acknowledgment.** This work was supported by the NRF (2019R1A2C3002833) and IITP (IITP-2015-0-00199, Proximity computing and its applications to autonomous vehicle, image search, and 3D printing) grants funded by the Korea government (MSIT). We thank the authors of RAFT, GMA, and SMURF for their public codes and dataset providers of Chairs, Things, Sintel, KITTI, VKITTI, and Spring for making the datasets available. Sung-Eui Yoon is a corresponding author.

## References

1. Bengio, Y., Bengio, S., Cloutier, J.: Learning a synaptic learning rule. Citeseer (1990) [6](#)
2. Blender: Spring Open Movie. <https://www.blender.org/press/spring-open-movie/> [13](#)
3. BlenderFoundation: Sintel, the durian open movie project. <http://durian.blender.org/> [13](#)
4. Butler, D.J., Wulff, J., Stanley, G.B., Black, M.J.: A naturalistic open source movie for optical flow evaluation. In: A. Fitzgibbon et al. (Eds.) (ed.) European Conf. on Computer Vision (ECCV). pp. 611–625. Part IV, LNCS 7577, Springer-Verlag (Oct 2012) [1](#), [7](#)
5. Cabon, Y., Murray, N., Humenberger, M.: Virtual kitti 2 (2020) [11](#)
6. Chen, X., He, K.: Exploring simple siamese representation learning. In: Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition. pp. 15750–15758 (2021) [3](#), [5](#), [6](#), [10](#)
7. Dosovitskiy, A., Fischer, P., Ilg, E., Hausser, P., Hazirbas, C., Golkov, V., Van Der Smagt, P., Cremers, D., Brox, T.: FlowNet: Learning optical flow with convolutional networks. In: Proceedings of the IEEE international conference on computer vision. pp. 2758–2766 (2015) [1](#), [2](#), [3](#), [7](#)
8. Finn, C., Abbeel, P., Levine, S.: Model-agnostic meta-learning for fast adaptation of deep networks. In: International Conference on Machine Learning. pp. 1126–1135. PMLR (2017) [6](#)
9. Ganin, Y., Lempitsky, V.: Unsupervised domain adaptation by backpropagation. In: International conference on machine learning. pp. 1180–1189. PMLR (2015) [3](#)
10. Grill, J.B., Strub, F., Altché, F., Tallec, C., Richemond, P.H., Buchatskaya, E., Doersch, C., Pires, B.A., Guo, Z.D., Azar, M.G., et al.: Bootstrap your own latent: A new approach to self-supervised learning. arXiv preprint arXiv:2006.07733 (2020) [3](#), [5](#), [6](#)
11. He, K., Zhang, X., Ren, S., Sun, J.: Deep residual learning for image recognition. In: Proceedings of the IEEE conference on computer vision and pattern recognition. pp. 770–778 (2016) [6](#)
12. Hinton, G., Vinyals, O., Dean, J.: Distilling the knowledge in a neural network. arXiv preprint arXiv:1503.02531 (2015) [3](#)
13. Hur, J., Roth, S.: Iterative residual refinement for joint optical flow and occlusion estimation. In: Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition. pp. 5754–5763 (2019) [3](#)
14. Ilg, E., Mayer, N., Saikia, T., Keuper, M., Dosovitskiy, A., Brox, T.: FlowNet 2.0: Evolution of optical flow estimation with deep networks. In: Proceedings of the IEEE conference on computer vision and pattern recognition. pp. 2462–2470 (2017) [1](#), [2](#), [3](#), [7](#)
15. Im, W., Kim, T.K., Yoon, S.E.: Unsupervised learning of optical flow with deep feature similarity. In: The European Conference on Computer Vision (ECCV) (2020) [3](#)
16. Jaderberg, M., Simonyan, K., Zisserman, A., et al.: Spatial transformer networks. In: Advances in neural information processing systems. pp. 2017–2025 (2015) [4](#)
17. Jiang, S., Campbell, D., Lu, Y., Li, H., Hartley, R.: Learning to estimate hidden motions with global motion aggregation. In: Proceedings of the IEEE/CVF International Conference on Computer Vision (ICCV). pp. 9772–9781 (October 2021) [13](#), [14](#)

18. Jonschkowski, R., Stone, A., Barron, J.T., Gordon, A., Konolige, K., Angelova, A.: What matters in unsupervised optical flow. arXiv preprint arXiv:2006.04902 (2020) [3](#), [7](#)
19. Kondermann, D., Nair, R., Honauer, K., Krispin, K., Andrusis, J., Brock, A., Gusefeld, B., Rahimimoghaddam, M., Hofmann, S., Brenner, C., et al.: The hci benchmark suite: Stereo and flow ground truth with uncertainties for urban autonomous driving. In: Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition Workshops. pp. 19–28 (2016) [13](#)
20. Lai, W.S., Huang, J.B., Yang, M.H.: Semi-supervised learning for optical flow with generative adversarial networks. In: Advances in neural information processing systems. pp. 354–364 (2017) [1](#), [3](#), [4](#), [9](#)
21. Liu, P., Lyu, M., King, I., Xu, J.: Selfflow: Self-supervised learning of optical flow. In: Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition. pp. 4571–4580 (2019) [3](#)
22. Liu, P., Lyu, M.R., King, I., Xu, J.: Learning by distillation: A self-supervised learning framework for optical flow estimation. IEEE Transactions on Pattern Analysis and Machine Intelligence (2021) [3](#)
23. Long, M., Zhu, H., Wang, J., Jordan, M.I.: Unsupervised domain adaptation with residual transfer networks. arXiv preprint arXiv:1602.04433 (2016) [6](#)
24. Mayer, N., Ilg, E., Fischer, P., Hazirbas, C., Cremers, D., Dosovitskiy, A., Brox, T.: What makes good synthetic training data for learning disparity and optical flow estimation? International Journal of Computer Vision **126**(9), 942–960 (2018) [1](#), [3](#)
25. Meister, S., Hur, J., Roth, S.: Unflow: Unsupervised learning of optical flow with a bidirectional census loss. In: Thirty-Second AAAI Conference on Artificial Intelligence (2018) [1](#), [3](#), [4](#)
26. Menze, M., Geiger, A.: Object scene flow for autonomous vehicles. In: Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition. pp. 3061–3070 (2015) [7](#)
27. Novák, T., Šochman, J., Matas, J.: A new semi-supervised method improving optical flow on distant domains. In: Computer Vision Winter Workshop (2020) [3](#), [4](#)
28. Perazzi, F., Pont-Tuset, J., McWilliams, B., Van Gool, L., Gross, M., Sorkine-Hornung, A.: A benchmark dataset and evaluation methodology for video object segmentation. In: Computer Vision and Pattern Recognition (2016) [11](#), [12](#)
29. Stanton, S., Izmailov, P., Kirichenko, P., Alemi, A.A., Wilson, A.G.: Does knowledge distillation really work? arXiv preprint arXiv:2106.05945 (2021) [3](#), [11](#)
30. Stone, A., Maurer, D., Ayvaci, A., Angelova, A., Jonschkowski, R.: Smurf: Self-teaching multi-frame unsupervised raft with full-image warping. In: Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition. pp. 3887–3896 (2021) [1](#), [3](#), [8](#), [10](#), [14](#)
31. Sun, D., Vlastic, D., Herrmann, C., Jampani, V., Krainin, M., Chang, H., Zabih, R., Freeman, W.T., Liu, C.: Autoflow: Learning a better training set for optical flow. In: Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition. pp. 10093–10102 (2021) [1](#), [2](#), [3](#), [9](#)
32. Sun, D., Yang, X., Liu, M.Y., Kautz, J.: Pwc-net: Cnns for optical flow using pyramid, warping, and cost volume. In: Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition. pp. 8934–8943 (2018) [1](#), [2](#), [3](#), [4](#)
33. Tarvainen, A., Valpola, H.: Mean teachers are better role models: Weight-averaged consistency targets improve semi-supervised deep learning results. arXiv preprint arXiv:1703.01780 (2017) [6](#), [8](#)



34. Teed, Z., Deng, J.: Raft: Recurrent all-pairs field transforms for optical flow. In: European Conference on Computer Vision. pp. 402–419. Springer (2020) [1](#), [2](#), [3](#), [6](#), [7](#), [13](#), [14](#)
35. Vapnik, V., Izmailov, R., et al.: Learning using privileged information: similarity control and knowledge transfer. *J. Mach. Learn. Res.* **16**(1), 2023–2049 (2015) [3](#)
36. Yan, W., Sharma, A., Tan, R.T.: Optical flow in dense foggy scenes using semi-supervised learning. In: Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition. pp. 13259–13268 (2020) [1](#)
37. Yang, G., Ramanan, D.: Volumetric correspondence networks for optical flow. In: Advances in neural information processing systems. pp. 794–805 (2019) [3](#)
38. Zhang, F., Woodford, O.J., Prisacariu, V.A., Torr, P.H.: Separable flow: Learning motion cost volumes for optical flow estimation. In: Proceedings of the IEEE/CVF International Conference on Computer Vision. pp. 10807–10817 (2021) [11](#), [13](#), [14](#)
39. Zoph, B., Ghiasi, G., Lin, T.Y., Cui, Y., Liu, H., Cubuk, E.D., Le, Q.V.: Rethinking pre-training and self-training. arXiv preprint arXiv:2006.06882 (2020) [6](#), [8](#)