

OCAI: Improving Optical Flow Estimation by Occlusion and Consistency Aware Interpolation

Jisoo Jeong Hong Cai Risheek Garrepalli Jamie Menjay Lin Munawar Hayat Fatih Porikli
Qualcomm AI Research[†]

{jisoojeon, hongcai, rgarrepa, jmlin, hayat, fporikli}@qti.qualcomm.com

Abstract

The scarcity of ground-truth labels poses one major challenge in developing optical flow estimation models that are both generalizable and robust. While current methods rely on data augmentation, they have yet to fully exploit the rich information available in labeled video sequences. We propose OCAI, a method that supports robust frame interpolation by generating intermediate video frames alongside optical flows in between. Utilizing a forward warping approach, OCAI employs occlusion awareness to resolve ambiguities in pixel values and fills in missing values by leveraging the forward-backward consistency of optical flows. Additionally, we introduce a teacher-student style semi-supervised learning method on top of the interpolated frames. Using a pair of unlabeled frames and the teacher model’s predicted optical flow, we generate interpolated frames and flows to train a student model. The teacher’s weights are maintained using Exponential Moving Averaging of the student. Our evaluations demonstrate perceptually superior interpolation quality and enhanced optical flow accuracy on established benchmarks such as Sintel and KITTI.

1. Introduction

Optical flow estimation and Video Frame Interpolation (VFI) share a complementary relationship. Accurate optical flow contributes significantly to various downstream tasks such as video compression [26, 44], video denoising and blur removal [3, 46, 49], action recognition [5, 22], and VFI stands as one of these applications. Pixel-level correspondence by optical flow enables estimating pixel-level movement and generating intermediate frames (or inter-frames). While utilizing flow-based methods is a common practice in VFI [12, 20, 23, 27], leveraging inter-frames to train optical flow models is relatively less explored.

While the scarcity of ground-truth data has long been a critical challenge in learning-based optical flow estimation [8, 11, 13, 34, 38, 41], there has been little attention

[†] Qualcomm AI Research is an initiative of Qualcomm Technologies, Inc.

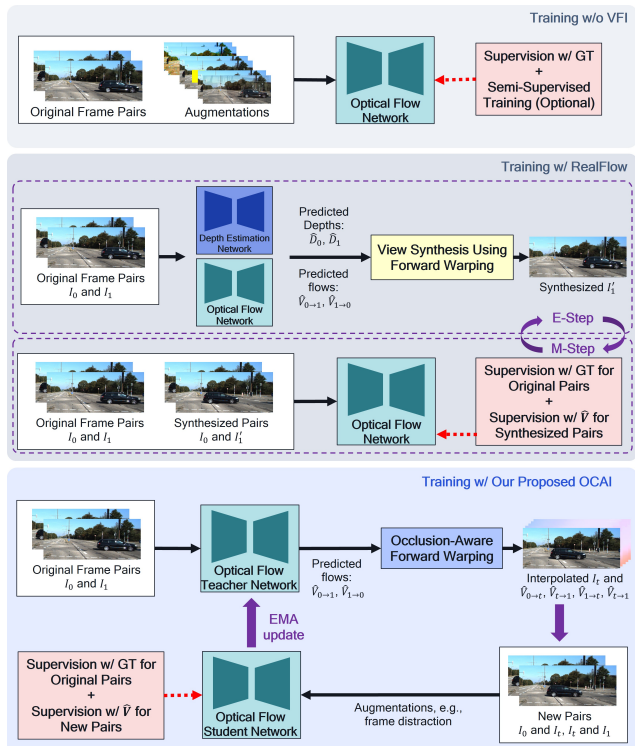


Figure 1. **Top:** Many existing data augmentation approaches focus on modifying the existing frames [18, 24, 41]. **Middle:** While RealFlow [10] employs optical flow model to generate new frames and iteratively update the model with generated frames, it requires compute expensive steps of depth estimation and EM optimization. **Bottom:** Our proposed OCAI allows flexible, robust video interpolation at any intermediate time step, and leverages interpolated frames and flows to efficiently train the model in a semi-supervised setting, significantly improving optical flow estimation.

into leveraging Video Frame Interpolation (VFI) to augment the training of optical flow networks. State-of-the-art VFI models predominantly employ deep networks trained to interpolate the exact middle frame between two consecutive time steps within a video sequence, which restricts their capability to generate frames at other intermediate time instances thus hampering their ability to produce opti-

cal flows between existing and intermediate frames reliably. Besides, these models lack generalizability across new domains without necessitating finetuning or retraining. These limitations hinder the potential use of existing VFI models for augmenting the training data for optical flow models.

To address data scarcity in optical flow training: data augmentation [18, 39], data generation [8, 10, 39], and semi-supervised learning [10, 14, 17, 18] have been explored. Most of the current data augmentation methods prioritize modifying the existing frames, e.g., [18], as illustrated in Fig. 1 (top). Notably, RealFlow[10] stands out by synthesizing a new second frame via forward warping. Given a pair of frames and the model’s prediction, the original first and second frames form a new training pair, used to update model weights. Model prediction, frame synthesis, and model update are iterated in an Expectation-Maximization (EM) framework, as depicted in Fig. 1 (middle). While it is possible to use forward warping to interpolate frames, RealFlow focuses only on making the model’s prediction consistent with frame synthesis. Moreover, the EM steps significantly increase the training workload.

In this paper, we introduce OCAI, a novel approach for training optical flow networks within a semi-supervised framework using readily available unlabeled pairs. Our method leverages video frames and flow interpolation to achieve this goal. OCAI implements an occlusion-aware forward warping technique, facilitating the interpolation of both inter-frames and intermediate flows, and can effectively tackle pixel value ambiguities during the warping process without needing any depth information and can do confidence-aware estimation of *missing values* utilizing forward-backward consistency. Our algorithm can perform interpolation at any intermediate time step, thus offering a broad diversity of data and motion ranges essential for training optical flow models. Our innovative teacher-student style semi-supervised learning scheme utilizes these resulting inter-frames to train optical flow networks effectively.

In summary, our main contributions are as follows:

- We propose a novel approach, OCAI, that tackles the data scarcity challenge in training optical flow models, by exploiting useful, hidden information in existing videos. Specifically, we interpolating frames and flows to generate supplementary data, and use them in a semi-supervised learning framework.
- To do this, we propose a new, effective video interpolation method that derives occlusion information to address ambiguous pixels and fill in holes by exploiting optical flow consistency. Our algorithm flexibly generates high-quality intermediate frames and reliable intermediate flows along with corresponding confidence maps at any intermediate time step.
- We devise a new teacher-student semi-supervised learning strategy leveraging VFI to train an optical flow

network, incorporating exponential moving averaging (EMA) to enhance training stability.

- We demonstrate that OCAI achieves higher quality in video interpolation than existing SOTA methods on standard datasets including Sintel and KITTI. By incorporating interpolated video information, our semi-supervised learning scheme significantly improves the optical flow estimation performance, e.g., 0.5+ F1-all reduction on KITTI test set when comparing to latest SOTA.

2. Related Work and Preliminaries

Consider two consecutive video frames, I_0 and I_1 . We denote the optical flow from I_0 to I_1 as $V_{0 \rightarrow 1}$ and the inter-frame as I_t , where $t \in (0, 1)$. We use ω_b and ω_f to denote backward and forward warping operations, respectively. For instance, backward warped image \hat{I}_0^b can be obtained by $\omega_b(I_1, V_{0 \rightarrow 1})$ and forward warped image \hat{I}_1^f can be computed by $\omega_f(I_0, V_{0 \rightarrow 1})$.

2.1. Video Frame Interpolation

Video Frame Interpolation (VFI) algorithms can be divided into three categories: flow-based [1, 2, 20, 23], kernel-based [6, 32] and hallucination-based [7, 35] approaches. Among these, flow-based methods have shown SOTA performance [20, 23]. These methods predict optical flows, and apply forward or backward warping using existing frames and the flows. The warped images are fused with residual information computed by the deep learning network.

Backward-warping-based VFI: The VFI network computes two optical flows, $V_{t \rightarrow 0}$ and $V_{t \rightarrow 1}$, a weighting mask (M), and the residual (R) using the two frames, I_0 and I_1 , where t is an intermediate time step. It applies backward warping to I_0 and I_1 , combines them using the mask, and finally, adds the residual information, as follows:

$$\hat{I}_t = M \cdot \omega_b(I_0, V_{t \rightarrow 0}) + (1 - M) \cdot \omega_b(I_1, V_{t \rightarrow 1}) + R, \quad (1)$$

where \hat{I}_t is the interpolated frame. This approach requires predicting $V_{t \rightarrow 0}$ and $V_{t \rightarrow 1}$ without I_t , which is challenging when there are large displacements.

Forward-warping-based VFI: Softmax Splating [31] is introduced for forward warping. It predicts two optical flows, $V_{0 \rightarrow 1}$ and $V_{1 \rightarrow 0}$, between I_0 and I_1 , and assumes that $V_{0 \rightarrow t} = t \cdot V_{0 \rightarrow 1}$ and $V_{1 \rightarrow t} = (1 - t) \cdot V_{1 \rightarrow 0}$. It generates I_t with the following steps:

$$\begin{aligned} \text{Let } u &= p - (q + V_{0 \rightarrow t}), \\ b(u) &= \max(0, 1 - |u_x|) \cdot \max(0, 1 - |u_y|). \end{aligned} \quad (2)$$

$$\text{Then, } \hat{I}_t(p) = \frac{\sum_q \exp(Z_0(q)) \cdot I_0(q) \cdot b(u)}{\sum_q \exp(Z_0(q)) \cdot b(u)}, \quad (3)$$

where p is target grid index, q is source grid index, $V_{0 \rightarrow t}$ is the optical flow between I_0 and I_t , u_x and u_y are x and y

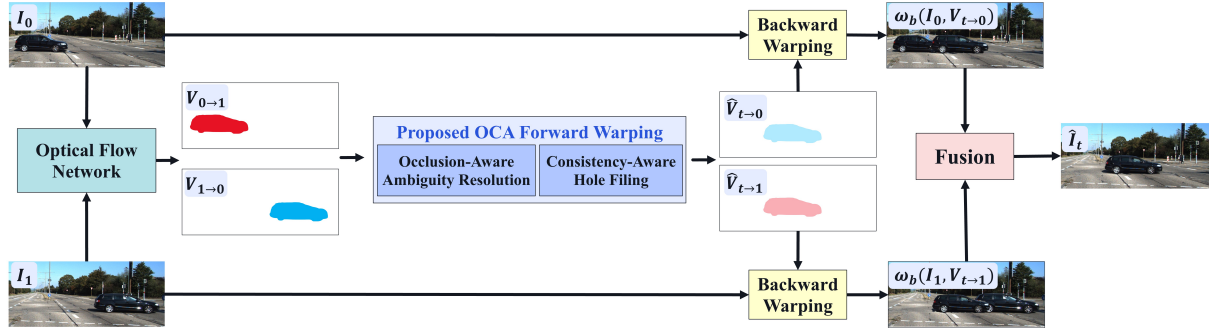


Figure 2. Our proposed video interpolation algorithm using occlusion and consistency-aware forward warping.

components of u , $b(u)$ indicates the mapping between pixels in I_0 and I_t , i.e., $b(u) = 1$ if p and q map to each other.

Softmax Splatting uses a learnable weighting Z and Realistic Image Pair Rendering (RIPR) in RealFlow [10] uses the inverse depth map of I_0 from a monocular depth estimation model. In forward warping, there can be pixel ambiguities, i.e., two source pixels can be mapped to the same target pixel location. In RIPR of RealFlow, as the weighting is based on depth, source pixels closer to the camera are chosen to resolve ambiguities. However, this requires additional computation, e.g., RIPR uses a depth network with $>300M$ parameters, and is susceptible to depth estimation errors. In addition, RIPR introduced the Bi-directional Hole Filling which fills in the holes in $w_f(I_0, V_{0 \rightarrow t})$ image using $w_f(I_1, V_{1 \rightarrow t})$ values. In OCAI, we introduce an occlusion-aware forward-warping algorithm that can generate accurate inter-frame without depth estimation.

2.2. Semi-Supervised Optical Flow Model Training

RAFT-OCTC [17] introduces transformation consistency [21, 40] to regressing optical flows. FlowSupervisor [14] proposes a new teacher network that can be trained stably. It uses the same encoder weights for both teacher and student, and computes the loss over all the pixels without using confidence masking. DistractFlow [18] introduces semantic augmentation which is inspired by Interpolation Regularization [16, 42, 48] and FixMatch [36]. In particular, DistractFlow proposes to use a confidence map derived from forward-backward consistency [29], which improves the stability of semi-supervised training. The confidence map is computed as follows:

$$C_{0,1} = \exp\left(-\frac{|\hat{V}_{0 \rightarrow 1}(x) + \hat{V}_{1 \rightarrow 0}(x + \hat{V}_{0 \rightarrow 1}(x))|^2}{\gamma_1(|\hat{V}_{0 \rightarrow 1}|^2 + |\hat{V}_{1 \rightarrow 0}(x + \hat{V}_{0 \rightarrow 1})|^2) + \gamma_2}\right), \quad (4)$$

where $\gamma_1 = 0.01$ and $\gamma_2 = 0.5$ from [29].

RealFlow [10] uses Expectation-Maximization (EM) to train an optical flow model in a semi-supervised setting. It first trains the model using supervision from existing ground-truth data. Then, it synthesizes new data based on the predicted optical flows and forward warping. After that, it trains the network with the new data. RealFlow repeats

these steps several times in the training process, which is computationally expensive.

Many semi-supervised learning algorithms [25, 40] for other tasks, such as classification and object detection, employ Exponential Moving Average (EMA) to robustly and stably update the teacher network, using a temporal ensemble of the student network. In this paper, we leverage video interpolation and several semi-supervised training techniques to enhance model performance.

3. Proposed Approach

OCAI improves the accuracy of optical flow models by generating diverse, high-quality intermediate frames and flows, and trains the network with new pairs in a semi-supervised learning framework. In Section 3.1, we present our occlusion and consistency aware forward warping algorithm for video interpolation. Next, in Section 3.2, we propose a semi-supervised learning strategy to leverage video interpolation to better train optical flow networks.

3.1. Occlusion and Consistency Aware Interpolation Using Forward Warping (OCAI)

Our goal is to generate an accurate estimation of the intermediate frame I_t by using I_0 , I_1 , $V_{0 \rightarrow 1}$, and $V_{1 \rightarrow 0}$. In order to do this, we first need to estimate $V_{t \rightarrow 0}$ and $V_{t \rightarrow 1}$, based on which we can perform backward warping from I_0 and I_1 and fuse the warped images, $\omega_b(I_0, \hat{V}_{t \rightarrow 0})$ and $\omega_b(I_1, \hat{V}_{t \rightarrow 1})$, to generate the estimated inter-frame, \hat{I}_t . In addition to generating I_t , we compute two intermediate optical flows, $V_{t \rightarrow 0}$ and $V_{t \rightarrow 1}$. This has several advantages. First, $V_{t \rightarrow 0}$ and $V_{t \rightarrow 1}$ can be used in optical flow training. Specifically, these optical flows enable computing a confidence map via forward-backward consistency, which is important to self-/semi-supervised training. Second, using the confidence map allows us to more accurately fuse warped versions of I_0 and I_1 to produce \hat{I}_t . Previous approach [10] generates \hat{I}_t from I_0 and fills hole using content generated from I_1 ; however, there are discrepancies since pixels have moved between I_0 and I_1 . Finally, our method can generate better background. For instance, RealFlow has missing




Figure 3. Visual illustration of deriving occlusion-aware weighting mask M_0 . The top part shows backward warping with a moving object and static background. Ghosting effect happens since the background is occluded in I_1 . The bottom part illustrates how we derive the foreground mask on I_0 , by performing backward warping and removing the ghosting effect. The foreground mask is used in the forward warping when synthesizing $\hat{V}_{t \rightarrow 1}$.

values near image boundaries whereas we do not have this issue (see last two rows in Fig. 5).

Next, we present our OCAI in detail.

Generating Inter-Frame via Forward Warping

First, we decompose $V_{0 \rightarrow 1}$ into $V_{0 \rightarrow t}$ and $V_{t \rightarrow 1}$:

$$V_{0 \rightarrow 1}(x) = V_{0 \rightarrow t}(x) + V_{t \rightarrow 1}(x + V_{0 \rightarrow t}(x)). \quad (5)$$

Assuming linear projected motion, $V_{0 \rightarrow t}$ can be properly approximated by $t \cdot V_{0 \rightarrow 1}$ and we have

$$V_{0 \rightarrow 1}(x) = t \cdot V_{0 \rightarrow 1}(x) + V_{t \rightarrow 1}(x + t \cdot V_{0 \rightarrow 1}(x)), \quad (6)$$

$$V_{t \rightarrow 1}(x + t \cdot V_{0 \rightarrow 1}(x)) = (1 - t) \cdot V_{0 \rightarrow 1}(x). \quad (7)$$

We see that by performing forward warping on $(1 - t) \cdot V_{0 \rightarrow 1}(x)$ with optical flow $t \cdot V_{0 \rightarrow 1}(x)$, we can obtain $V_{t \rightarrow 1}(x)$, that is

$$V_{t \rightarrow 1}(x) = \omega_f((1 - t) \cdot V_{0 \rightarrow 1}(x), t \cdot V_{0 \rightarrow 1}(x)). \quad (8)$$

Similarly, we can compute $V_{t \rightarrow 0}(x)$ based on decomposing $V_{1 \rightarrow 0}$ and forward warping.

By using $V_{t \rightarrow 0}$ and $V_{t \rightarrow 1}$, we can generate two versions of the intermediate frame I_t : $w_b(I_0, V_{t \rightarrow 0})$ and $w_b(I_0, V_{t \rightarrow 1})$, by backward warping. We can then use confidence maps, $C_{t,0}$ and $C_{t,1}$, to fuse them to estimate the inter-frame as follows:

$$\hat{I}_t = \frac{C_{t,0}}{C_{t,0} + C_{t,1}} w_b(I_0, V_{t \rightarrow 0}) + \frac{C_{t,1}}{C_{t,0} + C_{t,1}} w_b(I_0, V_{t \rightarrow 1}), \quad (9)$$

where the confidence maps are calculated based on Eq. 4. For instance, $C_{t,0}$ is computed using $V_{t \rightarrow 0}$ and $V_{0 \rightarrow t}$.

In order to correctly perform forward warping, we need to resolve two issues: **pixel value ambiguity** and **missing pixel values**. Ambiguity is due to two pixels in the source frame moving to the same location in the target frame, in which case we need to understand which is closer to the camera and thus, should be chosen. Missing values is because a pixel location in the target frame can correspond to an object that is occluded in the source frame (and the occluding object moves away in target frame), where there are no pixels representing this occluded object. This can also be caused by an object moving closer to the camera and there are not enough pixels in the source frame to represent the object in the target frame.

We propose occlusion-aware weighting to resolve pixel value ambiguity and choose the pixel that corresponds to what is closer to the camera. We further introduce a hole-filling method based on the forward-backward consistency of optical flow. Fig. 2 provides an overview of our proposed forward warping approach.

Occlusion-Aware Weighting to Resolve Ambiguity

We resolve pixel value ambiguity via occlusion understanding. Specifically, we assume that *when a pixel is not occluded but creates occlusion for other pixels, it corresponds to an object closer to the camera*.

More specifically, we derive an occlusion-aware weighting mask to be used in warping. First, we obtain occlusion map $O_{0,1}$ via forward-backward consistency [29], which indicates occlusion region on I_0 ; see how to compute forward-backward consistency in Eq. 1 of [29]. Next, we apply backward warping to $O_{0,1}$ using $V_{0 \rightarrow 1}$ and by further applying the non-occlusion mask $1 - O_{0,1}$, we can infer the pixels in I_0 that produce occlusion, i.e., foreground pixels, and accordingly generate the mask to select these pixels. More specifically, the occlusion-aware weighting mask is computed as follows:

$$M_0 = \alpha \cdot (1 - O_{0,1}) \cdot \omega_b(O_{0,1}, V_{0 \rightarrow 1}), \quad (10)$$

where α is a coefficient controlling the weighting.

Fig. 3 provides a visual example. Assume that the disk moves to the right from I_0 to I_1 and the background is static. As indicated by $O_{0,1}$, the disk is not occluded and creates an occlusion blocking the background, i.e., it is the foreground. Due to this occlusion, when performing backward warping on I_1 or $O_{0,1}$ using $V_{0 \rightarrow 1}$, we not only move the disk back to its original location at $t = 0$, but also create a duplicate at its current location since $V_{0 \rightarrow 1}$ is zero for the background (known as the ghosting effect [51]). By applying the non-occlusion mask, we can remove the ghosting effect and recover the disk's original location; in other words, we obtain the weighting mask for selecting foreground pixels.

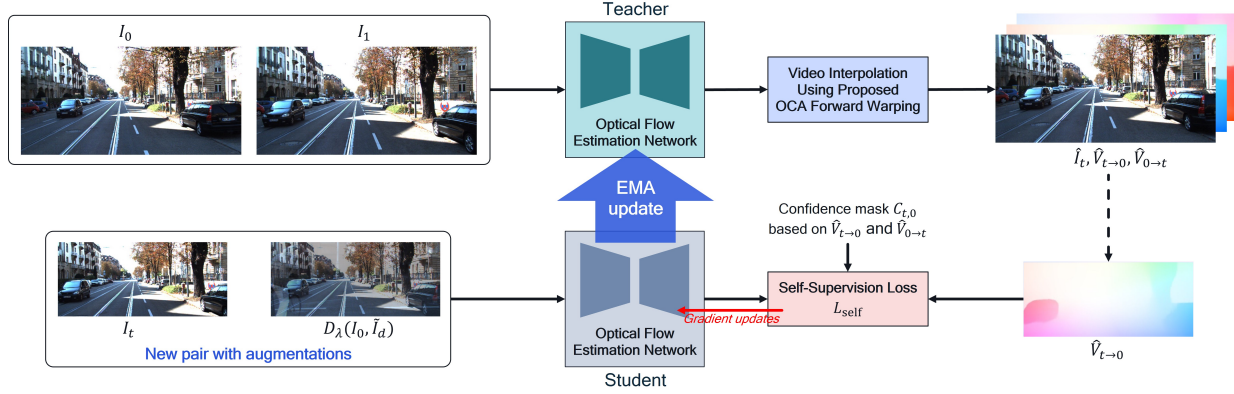


Figure 4. Self-supervision using interpolated video frames and flows in a teacher-student learning setting. Note that the student model is trained with both the self-supervision shown in the figure and the supervision from available ground-truth data.

We replace the depth-based masking in Eq. 3 with our occlusion-aware weighting mask M_0 and apply forward warping in Eq. 8, which gives

$$\hat{V}_{t \rightarrow 1}(p) = \frac{\sum_q \exp(M_0(q)) \cdot (1-t) \cdot V_{0 \rightarrow 1} \cdot b(u)}{\sum_q \exp(M_0(q)) \cdot b(u)}, \quad (11)$$

where q, p , and $b(u)$ are from Eq. 2. $V_{t \rightarrow 0}$ can be computed in a similar way.

Hole Filling Using Optical Flow Consistency

Assuming linear motion, the directions in the optical flows $V_{t \rightarrow 0}$ and $V_{t \rightarrow 1}$ should be exactly opposite, and the magnitude ratio of the two flows should be $t : (1-t)$. Based on this assumption, we fill in the missing values in the synthesized optical flow maps as follows:

$$V_{t \rightarrow 0}(p) = -\frac{t}{1-t} V_{t \rightarrow 1}(p), \quad \text{if } \sum_q b(u) = 0, \quad (12)$$

where $\sum_q b(u) = 0$ indicates that no source pixels are assigned to the target location, thus creating a hole.

Now that we can resolve the ambiguities and missing values in forward warping, we can use our computed $\hat{V}_{t \rightarrow 0}$ and $\hat{V}_{t \rightarrow 1}$, along with the confidence maps, to perform backward warping and fusion to produce the inter-frame, based on Eq. 9.

3.2. Teacher-Student Semi-Supervised Learning

The generated intermediate video frames and optical flows create an opportunity to significantly augment the training of optical flow models. We propose a new semi-supervised training strategy to leverage the new image pairs and flows.

More specifically, we adopt a teacher-student training approach. The teacher network consumes an original pair of video frames and predicts the forward and backward optical flows. By using our proposed video interpolation algorithm, we can generate an inter-frame I_t as well as the corresponding inter-flows $V_{0 \rightarrow t}$ and $V_{t \rightarrow 0}$, for any intermediate time step. We randomly sample $t \in [0, 1]$ to expand motion

diversity in training data. We use (I_0, I_t) to form a new training pair. The generated intermediate optical flows not only supply the supervision signal for the new pair, but also allows us to compute a confidence mask to make training more stable, i.e., we only train the network with the reliable optical flows. We train the student model with these new pairs derived from interpolation as well as with the available pseudo ground-truth data. Inspired by Mean Teacher [40], we employ the Exponential Moving Average (EMA) update method to update the teacher model with a temporal ensemble of the student model.

In addition, as inspired by Smurf [37], RealFlow [10], and DistractFlow [18], we further impose semantic distraction to I_0 to generate an augmented version $D_\lambda(I_0, \tilde{I}_d)$ (see DistractFlow for details of the frame distraction operation), and crop I_t and $D_\lambda(I_0, \tilde{I}_d)$. Our self-supervision loss is given as follows for each new pair:

$$\mathcal{L}_{\text{self}} = [C_{t,0} \geq \tau] \|f_{\text{student}}(I_t, D_\lambda(I_0, \tilde{I}_d)) - \hat{V}_{t \rightarrow 0}\|_1, \quad (13)$$

where $f_{\text{student}}(I_t, D_\lambda(I_0, \tilde{I}_d))$ denotes the optical flow prediction on frame pair I_t and $D_\lambda(I_0, \tilde{I}_d)$ by the student network, $C_{t,0}$ is the confidence map derived from $\hat{V}_{t \rightarrow 0}$ and $\hat{V}_{0 \rightarrow t}$ based on Eq. 4, and τ is a threshold.

Combining this with the supervision from available ground-truth data, the total training loss for the student network is given by

$$\mathcal{L}_{\text{total}} = \mathcal{L}_{\text{sup}} + w \cdot \mathcal{L}_{\text{self}}, \quad (14)$$

where w is a weighting coefficient.

4. Experiment

We evaluate our proposed OCAI method for video interpolation and for semi-supervised optical flow training on benchmark datasets, and compare OCAI with baselines and latest state-of-the-art (SOTA) methods.

¹Note that PSNR and SSIM do not always correlate well with perceived visual quality, as shown in previous studies [33, 50]. For instance, in Fig. 5, VFIFormer has higher SSIM scores but worse visual quality.

Table 1. Video Frame Interpolation (VFI) results on Sintel (Clean), KITTI datasets. There are five state of the art backward warping based VFI algorithms (First to fifth rows), two forward warping based VFI algorithms (sixth to seventh) and ours (bottom). Algorithms of first to sixth rows trained their model on Vimeo dataset using pre-trained optical flow model, and two algorithms of seventh and eighth use RAFT trained on FlyingChairs and FlyingThings3D. A/V denotes AlexNet/VGG used in LPIPS. **Bold/Underline**: Best and second best results.

Method	Sintel (12FPS → 24 FPS)		Sintel (6FPS → 12 FPS)		KITTI (5FPS → 10 FPS)		Param (M)
	PSNR / SSIM ↑	LPIPS (A) / (V) ↓	PSNR / SSIM ↑	LPIPS (A) / (V) ↓	PSNR / SSIM ↑	LPIPS (A) / (V) ↓	
IFRNet-B [20] (CVPR 2022)	30.06 / 0.901	0.093 / 0.128	25.87 / 0.837	0.154 / 0.195	21.64 / 0.760	0.140 / 0.217	5.0
VFIFormer [27] (CVPR 2022)	30.37 / 0.909	0.088 / 0.115	<u>26.22</u> / 0.849	0.154 / 0.184	22.53 / 0.787	0.149 / 0.227	24.1
RIFE [12] (ECCV 2022)	30.04 / 0.899	0.103 / 0.135	25.79 / 0.835	0.169 / 0.206	21.50 / 0.752	0.158 / 0.235	9.8
EMA-VFI [47] (CVPR 2023)	30.51 / 0.906	0.101 / 0.128	26.29 / 0.844	0.165 / 0.196	22.00 / 0.767	0.177 / 0.253	66.0
AMT-L [23] (CVPR 2023)	30.30 / 0.905	0.087 / 0.119	26.19 / 0.843	0.149 / 0.186	21.97 / 0.773	0.151 / 0.224	12.9
SoftSplat [31] (CVPR 2020)	<u>30.42</u> / <u>0.907</u>	0.090 / 0.118	26.23 / <u>0.846</u>	0.153 / 0.186	21.95 / <u>0.776</u>	0.169 / 0.244	12.2
RIPR [10] (ECCV 2022)	28.26 / 0.894	<u>0.084</u> / 0.120	24.91 / 0.826	<u>0.135</u> / 0.181	21.18 / 0.733	0.112 / 0.195	349.5
OCAI (ours)	29.47 / 0.904	0.078 / 0.117	25.88 / 0.838	0.129 / 0.178	<u>22.08</u> / 0.758	0.112 / 0.190	5.3

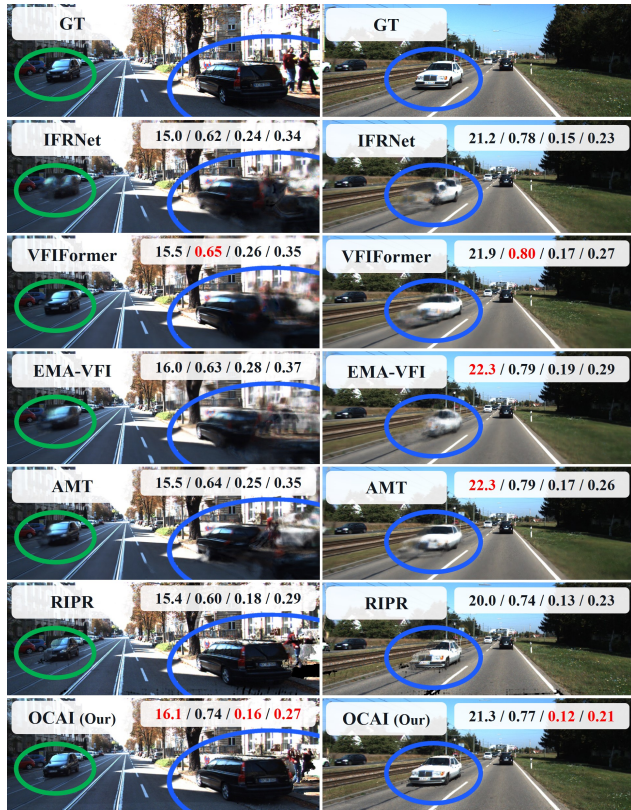


Figure 5. Video Frame Interpolation (VFI) results on KITTI. First row is the ground truth. Second to fifth rows are outputs of SOTA VFI models [20, 23, 27, 47]. Sixth row is the output of using RealFlow [10] for VFI. Bottom row shows our OCAI results. For each interpolated frame, we show the PSNR, SSIM, and LPIPS (using AlexNet and VGG) scores. Best scores are shown in red.¹

4.1. Experimental Setup

Video Frame Interpolation (VFI): We compare with latest SOTA VFI algorithms [12, 20, 23, 27, 47]. We use their official codes and weights trained on Vimeo90k [45]. In addition, we compare with RIPR of RealFlow [10] as our forward warping baseline. We use their official code to generate inter-frames and RAFT [41] trained on FlyingChair (C) [8] and FlyingThings3D (T) [28] as the optical flow

model. In OCAI, we use the same optical flow network, i.e., RAFT trained on C+T, for fair comparison. For evaluation, we use Sintel (S) [4] and KITTI (K) [9, 30], which are standard optical flow datasets.² More specifically, each test sample consists of three consecutive frames, with the first and third used as existing frames, and the second as the interpolation target. We use common image similarity metrics to evaluate VFI quality, including PSNR, SSIM [43], and LPIPS (using AlexNet and VGG) [50]. More details can be found in the supplementary file.

Semi-Supervised Optical Flow (SSOF): We use RAFT as the network architecture, following previous semi-supervised optical flow model training settings [10, 14, 18]. When evaluating on Sintel (train) and SlowFlow [15], we first pretrain the network on C+T and then, use FlyingThings3D (T) as the labeled dataset and Sintel (S) as unlabeled dataset. For KITTI (train) evaluation, we use FlyingThings3D as the labeled dataset and KITTI (multiview) test as the unlabeled dataset, with initialization from C+T pre-trained weight. For Sintel and KITTI (test) evaluations, we use the same labeled datasets (i.e., C+T+S+K+HD1K [19]) following the original RAFT supervised training setting, and use Sintel training ((I_t, I_{t+2}) pairs), Monkaa [28], and KITTI (multiview) training dataset as unlabeled data. Note that Sintel and KITTI test sets are not used as unlabeled data for training in these test evaluations.

4.2. Video Frame Interpolation

Table 1 and Fig. 5 show the performance of SOTA VFI methods, RIPR of RealFlow, and our method on Sintel (clean) and KITTI. We achieve the best LPIPS scores and PSNR/SSIM scores on-par with existing SOTA solutions; note that PSNR and SSIM do not always correctly reflect visual quality [33]. Since SOTA methods use backward warping and predict intermediate two optical flows ($V_{t \rightarrow 0}$ and $V_{t \rightarrow 1}$) without the inter-frame I_t , they do not work well when there are large displacements, which results in blurriness (see cars in Fig 5). In contrast, forward warping methods (RIPR of RealFlow and ours) can predict accurate $V_{0 \rightarrow 1}$

²The frames per second (FPS) numbers of Sintel and KITTI are 24 and 12, respectively.

Table 2. Optical flow results on SlowFlow, Sintel (train), and KITTI (train) datasets. We train the model on FlyingChairs (C) and FlyingThings3D (T) as labeled data, and Sintel and KITTI multiview (S/K) as unlabeled data. BD in SlowFlow represents Blur Duration. **Bold/Underline**: Best and second best results.

Method	Labeled data	Unlabeled data	SlowFlow (100px)		Sintel (train)		KITTI (train)	
			(3BD/epe)	(5BD/epe)	(Clean-epe)	(Final-epe)	(FI-epe)	(FI-all)
RAFT-Supervised	C + T		7.98	6.72	1.43	2.71	5.04	17.4
RAFT-A [39] (CVPR 2021)	AutoFlow [39]		-	-	1.95	2.57	4.23	-
RAFT-OCTC [17] (CVPR 2022)		T (subsampled)	-	-	1.31	2.67	4.72	16.3
Fixed Teacher [14] (ECCV 2022)			-	-	1.32	2.58	4.91	15.9
FlowSupervisor [14] (ECCV 2022)	C + T	S/K	-	-	1.30	2.46	3.35	11.1
RealFlow [10] (ECCV 2022)			-	-	1.34	2.38	<u>2.16</u>	<u>8.5</u>
DistractFlow [18] (CVPR 2023)			3.60	<u>5.15</u>	<u>1.25</u>	<u>2.35</u>	3.01	11.7
OCAI (ours)			2.97	5.04	1.20	2.32	2.07	7.6

Table 3. Optical flow results on Sintel and KITTI test. * indicates “warm-start” results that use previous flow prediction.

Method	Sintel (test) (Final-epe)	KITTI (test) (FI-all)
RAFT-Supervised	3.18/2.86*	5.10
RAFT-A [39] (CVPR 2021)	3.14	4.78
RAFT-OCTC [17] (CVPR 2022)	3.09	4.72
FlowSupervisor [14] (ECCV 2022)	2.79*	4.85
RealFlow [10] (ECCV 2022)	-	4.63
DistractFlow [18] (CVPR 2023)	2.71*	4.71
OCAI (ours)	2.63*	4.13

and $V_{1 \rightarrow 0}$, and handles fast moving objects better. Furthermore, our OCAI method produces better VFI quality than RealFlow, with sharper image details and fewer holes. More interpolation results using different t values (e.g., 0.2, 0.4, 0.6, 0.8) can be found in the supplementary material.

OCAI only requires the RAFT optical flow model with 5.3M parameters and shows better VFI performance as compared to existing SOTA methods using similar or more network parameters. Backward-warping-based approaches require additional training on Vimeo data with a pre-trained optical flow model, while RIPR of RealFlow and our OCAI only needs the pre-trained optical flow model without any additional training. Moreover, we achieve better VFI scores than RIPR on both Sintel and KITTI. These results demonstrate that OCAI generates accurate inter-frames without needing depth estimation in forward warping.

4.3. Optical Flow

Table 2 shows the optical flow estimation evaluation results on SlowFlow, Sintel, and KITTI. We see that our proposed OCAI achieves the best performance on all datasets. Notably, it has significantly more accurate optical flow estimation as compared to latest SOTA such as RealFlow and DistractFlow. Specifically, on KITTI, OCAI brings nearly 1-point reduction in FI-epe when comparing to DistractFlow and 1-point smaller FI-all than RealFlow.

Table 3 shows the evaluation result on KITTI test dataset. We achieve SOTA performance in semi-supervised optical flow on KITTI, bringing a significant improvement as compared to existing SOTA semi-supervised optical flow algorithms.

Table 4. Ablation Study of video frame interpolation on KITTI. M refers to using our occlusion-aware weighting mask (Eq. 10) in Z (Eq. 3). BHF is Bi-directional Hole Filling proposed in RealFlow, CBF refers to Confidence-Based Fusion (Eq. 9).

Method	Z (Eq. 3)	Warping	Hole Filling	Fusion	KITTI	
					SSIM \uparrow	LPIPS (V) \downarrow
RIPR [10]	Depth	Image	Image	BHF	0.733	0.195
	M	Image	Image	BHF	0.734	0.198
OCAI	M	Flow	-	BHF	0.721	0.213
	M	Flow	Flow	BHF	0.739	0.195
	M	Flow	Flow	CBF	0.758	0.190




Figure 6. Video Frame Interpolation (VFI) results on Sintel. RIPR (D) and RIPR (M) images are generated using RIPR with depth or occlusion weighting mask (M). OCAI is generated using M , with flow hole-filling and confidence fusion.

5. Ablation Studies

5.1. Video Frame Interpolation

Depth weighting vs. occlusion-aware weighting. Table 4 shows the effectiveness of our occlusion- and consistency-Aware forward warping for VFI. The first row shows forward image warping using depth weighting. In the second row, we replace depth weighting with our occlusion-aware weighting mask, which shows comparable performance without using depth. Fig. 6 provides sample qualitative results. RIPR (from RealFlow) using depth and us-




Figure 7. Forward warped Optical Flow on Sintel (Same images from 6). In top, there are two optical flows ($V_{t \rightarrow 1}$, $V_{t \rightarrow 0}$) from SOTA VFI algorithm. Two optical flows from images are shown in second row, and forward warped flows with holes are shown in third row ($V_{t \rightarrow 1}^*$, $V_{t \rightarrow 0}^*$). There are hole filled optical flow in the bottom. (Since visualization code applies normalization, color intensity looks similar, but the magnitude was reduced and shift and hole filling are visible.)

ing occlusion-aware weighting both have holes (See RIPR (D) and RIPR (M)). Since RIPR performs hole-filling using another warped image, a hole cannot be filled when both images have holes in the same corresponding regions (e.g., wings and head in Fig. 6).

Image warping vs. flow warping. In the third row of Table 4, we show the effect of applying forward warping to flow instead of image (see optical flows in the third row of Fig. 7). After performing hole-filling based on optical flow consistency, our intermediate flows are significantly improved (see inpainted flows in the fourth row of Fig. 7). By generating the confidence maps, we can then combine two warped images to more accurately generate the inter-frame (see OCAI output in Fig. 6).

5.2. Semi-Supervised Optical Flow

Using baseline VFI for semi-supervised learning. Table 5 shows an ablation study of the semi-supervised training. RealFlow trains the model using EM algorithm, which requires significantly more iterations. In our training, when only using EMA, our model already has a lower FI-all score as compared to RealFlow with 1 EM iteration. By addi-

Table 5. Ablation study of our semi-supervised optical flow training. We denote EMA, confidence-base loss masking, and image distraction as E, C, and D, respectively. We denote KITTI test dataset as K-Test, and RealFlow as RF.

Method	E	C	D	Dataset	Training iterations	KITTI 15	
						(FI-epe)	(FI-all)
RAFT [41]						5.04	17.4
RealFlow [10]				RF-K-Test	200k (1 EM)	2.79	10.7
				RF-K-Test	800k (4 EM)	2.16	8.5
DistractFlow			✓	S/K-Test	150k	3.01	11.7
OCAI (ours)	✓			K-Test	100k	Diverged	
	✓			K-Test	100k	2.85	10.1
	✓	✓		K-Test	100k	2.81	9.7
	✓	✓	✓	K-Test	100k	2.27	8.4
	✓	✓	✓	OCAI-K-Test	100k	2.07	7.6

Table 6. Effectiveness of new semi-supervised trained model for Video Frame interpolation.

VFI dataset	Flow Train Dataset		PSNR / SSIM \uparrow	LPIPS (A) / (V) \downarrow
	Label	Unlabel		
Sintel (12 FPS \rightarrow 24 FPS)		-	29.47 / 0.904	0.078 / 0.117
		OCAI-S-Test	29.51 / 0.905	0.756 / 0.115
Sintel (6 FPS \rightarrow 12 FPS)		-	25.88 / 0.838	0.129 / 0.178
		OCAI-S-Test	25.87 / 0.840	0.127 / 0.176
KITTI (5 FPS \rightarrow 10 FPS)		-	22.08 / 0.758	0.112 / 0.190
		OCAI-K-Test	22.29 / 0.759	0.108 / 0.186

tionally masking the loss with confidence map and imposing image distractions, the model further improves. Note that EMA is crucial for training stability; the training fails to converge when teacher and student models are directly weight-shared. Finally, when using intermediate frames and flows generated by OCAI in training, our model achieves significantly lower FI-epe and FI-all.

VFI using optical flow model trained from semi-supervised scheme. After we train the optical flow in semi-supervised manner, we evaluate the VFI performance of semi-supervised training in Table 6. Our semi-supervised training weight shows the improvements on all evaluation metrics on all dataset except for Sintel PSNR in 6 FPS to 12 FPS setting. While our OCAI trains the optical flow in a semi-supervised manner, improved optical flow can generate more accurate inter-frame. Our OCAI can boost the performances of not only optical flow but also VFI.

6. Conclusion

In this paper, we proposed a novel scheme that significantly augments the training of optical flow models. This effectively alleviates the lack of ground-truth optical flow labels in existing datasets. More specifically, we first proposed an occlusion-aware video frame interpolation method, which can robustly generate interframes despite large motions, as well as the intermediate optical flows. This allows us to significantly expand existing optical flow training data for free. We further proposed a semi-supervised training approach by leveraging the video frame interpolation. Through extensive experiments on standard optical flow benchmarks like Sintel and KITTI, we demonstrate the efficacy of our proposed approach and that it sets the new state of the art.

References

- [1] Wenbo Bao, Wei-Sheng Lai, Chao Ma, Xiaoyun Zhang, Zhiyong Gao, and Ming-Hsuan Yang. Depth-aware video frame interpolation. In *Proceedings of the IEEE/CVF conference on computer vision and pattern recognition*, pages 3703–3712, 2019. **2**
- [2] Wenbo Bao, Wei-Sheng Lai, Xiaoyun Zhang, Zhiyong Gao, and Ming-Hsuan Yang. Memc-net: Motion estimation and motion compensation driven neural network for video interpolation and enhancement. *IEEE transactions on pattern analysis and machine intelligence*, 43(3):933–948, 2019. **2**
- [3] Antoni Buades, Jose-Luis Lisani, and Marko Miladinović. Patch-based video denoising with optical flow estimation. *IEEE Transactions on Image Processing*, 25(6):2573–2586, 2016. **1**
- [4] Daniel J Butler, Jonas Wulff, Garrett B Stanley, and Michael J Black. A naturalistic open source movie for optical flow evaluation. In *Proceedings of the European Conference on Computer Vision*, pages 611–625. Springer, 2012. **6**
- [5] Zixi Cai, Helmut Neher, Kanav Vats, David A Clausi, and John Zelek. Temporal hockey action recognition via pose and optical flows. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition Workshops*, pages 0–0, 2019. **1**
- [6] Xianhang Cheng and Zhenzhong Chen. Video frame interpolation via deformable separable convolution. In *Proceedings of the AAAI Conference on Artificial Intelligence*, pages 10607–10614, 2020. **2**
- [7] Myungsub Choi, Heewon Kim, Bohyung Han, Ning Xu, and Kyoung Mu Lee. Channel attention is all you need for video frame interpolation. In *Proceedings of the AAAI Conference on Artificial Intelligence*, pages 10663–10671, 2020. **2**
- [8] Alexey Dosovitskiy, Philipp Fischer, Eddy Ilg, Philip Hausser, Caner Hazirbas, Vladimir Golkov, Patrick Van Der Smagt, Daniel Cremers, and Thomas Brox. FlowNet: Learning optical flow with convolutional networks. In *Proceedings of the IEEE/CVF International Conference on Computer Vision*, pages 2758–2766, 2015. **1, 2, 6**
- [9] Andreas Geiger, Philip Lenz, Christoph Stiller, and Raquel Urtasun. Vision meets robotics: The kitti dataset. *The International Journal of Robotics Research*, 32(11):1231–1237, 2013. **6**
- [10] Yunhui Han, Kunming Luo, Ao Luo, Jiangyu Liu, Haoqiang Fan, Guiming Luo, and Shuaicheng Liu. Realflow: Embased realistic optical flow dataset generation from videos. *arXiv preprint arXiv:2207.11075*, 2022. **1, 2, 3, 5, 6, 7, 8**
- [11] Zhaoyang Huang, Xiaoyu Shi, Chao Zhang, Qiang Wang, Ka Chun Cheung, Hongwei Qin, Jifeng Dai, and Hongsheng Li. Flowformer: A transformer architecture for optical flow. In *Proceedings of the European Conference on Computer Vision*, 2022. **1**
- [12] Zhewei Huang, Tianyuan Zhang, Wen Heng, Boxin Shi, and Shuchang Zhou. Real-time intermediate flow estimation for video frame interpolation. In *European Conference on Computer Vision*, pages 624–642. Springer, 2022. **1, 6**
- [13] Eddy Ilg, Nikolaus Mayer, Tonmoy Saikia, Margret Keuper, Alexey Dosovitskiy, and Thomas Brox. FlowNet 2.0: Evolution of optical flow estimation with deep networks. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 2462–2470, 2017. **1**
- [14] Woobin Im, Sebin Lee, and Sung-Eui Yoon. Semi-supervised learning of optical flow by flow supervisor. In *Proceedings of the European Conference on Computer Vision*, 2022. **2, 3, 6, 7**
- [15] Joel Janai, Fatma Güney, Jonas Wulff, Michael Black, and Andreas Geiger. Slow flow: Exploiting high-speed cameras for accurate and diverse optical flow reference data. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, 2017. **6**
- [16] Jisoo Jeong, Vikas Verma, Minsung Hyun, Juho Kannala, and Nojun Kwak. Interpolation-based semi-supervised learning for object detection. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 11602–11611, 2021. **3**
- [17] Jisoo Jeong, Jamie Menjay Lin, Fatih Porikli, and Nojun Kwak. Imposing consistency for optical flow estimation. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 3181–3191, 2022. **2, 3, 7**
- [18] Jisoo Jeong, Hong Cai, Risheek Garrepalli, and Fatih Porikli. Distractflow: Improving optical flow estimation via realistic distractions and pseudo-labeling. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 13691–13700, 2023. **1, 2, 3, 5, 6, 7**
- [19] Daniel Kondermann, Rahul Nair, Katrin Honauer, Karsten Krispin, Jonas Andrulis, Alexander Brock, Burkhard Gussefeld, Mohsen Rahimimoghaddam, Sabine Hofmann, Claus Brenner, et al. The hci benchmark suite: Stereo and flow ground truth with uncertainties for urban autonomous driving. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition Workshops*, pages 19–28, 2016. **6**
- [20] Lingtong Kong, Boyuan Jiang, Donghao Luo, Wenqing Chu, Xiaoming Huang, Ying Tai, Chengjie Wang, and Jie Yang. Ifrnet: Intermediate feature refine network for efficient frame interpolation. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 1969–1978, 2022. **1, 2, 6**
- [21] Samuli Laine and Timo Aila. Temporal ensembling for semi-supervised learning. *arXiv preprint arXiv:1610.02242*, 2016. **3**
- [22] Myunggi Lee, Seungeui Lee, Sungjoon Son, Gyutae Park, and Nojun Kwak. Motion feature network: Fixed motion filter for action recognition. In *Proceedings of the European Conference on Computer Vision*, pages 387–403, 2018. **1**
- [23] Zhen Li, Zuo-Liang Zhu, Ling-Hao Han, Qibin Hou, Chun-Le Guo, and Ming-Ming Cheng. Amt: All-pairs multi-field transforms for efficient frame interpolation. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 9801–9810, 2023. **1, 2, 6**
- [24] Pengpeng Liu, Michael Lyu, Irwin King, and Jia Xu. Self-low: Self-supervised learning of optical flow. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 4571–4580, 2019. **1**

- [25] Yen-Cheng Liu, Chih-Yao Ma, Zijian He, Chia-Wen Kuo, Kan Chen, Peizhao Zhang, Bichen Wu, Zsolt Kira, and Peter Vajda. Unbiased teacher for semi-supervised object detection. *arXiv preprint arXiv:2102.09480*, 2021. 3
- [26] Guo Lu, Wanli Ouyang, Dong Xu, Xiaoyun Zhang, Chunlei Cai, and Zhiyong Gao. Dvc: An end-to-end deep video compression framework. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 11006–11015, 2019. 1
- [27] Liying Lu, Ruizheng Wu, Huaijia Lin, Jiangbo Lu, and Jiaya Jia. Video frame interpolation with transformer. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 3532–3542, 2022. 1, 6
- [28] Nikolaus Mayer, Eddy Ilg, Philip Hausser, Philipp Fischer, Daniel Cremers, Alexey Dosovitskiy, and Thomas Brox. A large dataset to train convolutional networks for disparity, optical flow, and scene flow estimation. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 4040–4048, 2016. 6
- [29] Simon Meister, Junhwa Hur, and Stefan Roth. Unflow: Unsupervised learning of optical flow with a bidirectional census loss. In *Proceedings of the AAAI conference on artificial intelligence*, 2018. 3, 4
- [30] Moritz Menze and Andreas Geiger. Object scene flow for autonomous vehicles. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 3061–3070, 2015. 6
- [31] Simon Niklaus and Feng Liu. Softmax splatting for video frame interpolation. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 5437–5446, 2020. 2, 6
- [32] Tomer Peleg, Pablo Szekely, Doron Sabo, and Omry Sendik. Im-net for high resolution video frame interpolation. In *Proceedings of the IEEE/CVF conference on computer vision and pattern Recognition*, pages 2398–2407, 2019. 2
- [33] Ekta Prashnani, Hong Cai, Yasamin Mostofi, and Pradeep Sen. Pieapp: Perceptual image-error assessment through pairwise preference. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, pages 1808–1817, 2018. 5, 6
- [34] Anurag Ranjan and Michael J Black. Optical flow estimation using a spatial pyramid network. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 4161–4170, 2017. 1
- [35] Zhihao Shi, Xiangyu Xu, Xiaohong Liu, Jun Chen, and Ming-Hsuan Yang. Video frame interpolation transformer. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 17482–17491, 2022. 2
- [36] Kihyuk Sohn, David Berthelot, Nicholas Carlini, Zizhao Zhang, Han Zhang, Colin A Raffel, Ekin Dogus Cubuk, Alexey Kurakin, and Chun-Liang Li. Fixmatch: Simplifying semi-supervised learning with consistency and confidence. *Advances in neural information processing systems*, 33:596–608, 2020. 3
- [37] Austin Stone, Daniel Maurer, Alper Ayvaci, Anelia Angelova, and Rico Jonschkowski. Smurf: Self-teaching multi-frame unsupervised raft with full-image warping. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 3887–3896, 2021. 5
- [38] Deqing Sun, Xiaodong Yang, Ming-Yu Liu, and Jan Kautz. Pwc-net: Cnns for optical flow using pyramid, warping, and cost volume. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 8934–8943, 2018. 1
- [39] Deqing Sun, Daniel Vlasic, Charles Herrmann, Varun Jampani, Michael Krainin, Huiwen Chang, Ramin Zabih, William T Freeman, and Ce Liu. Autoflow: Learning a better training set for optical flow. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 10093–10102, 2021. 2, 7
- [40] Antti Tarvainen and Harri Valpola. Mean teachers are better role models: Weight-averaged consistency targets improve semi-supervised deep learning results. In *Advances in neural information processing systems*, pages 1195–1204, 2017. 3, 5
- [41] Zachary Teed and Jia Deng. Raft: Recurrent all-pairs field transforms for optical flow. In *Proceedings of the European Conference on Computer Vision*, pages 402–419. Springer, 2020. 1, 6, 8
- [42] Vikas Verma, Alex Lamb, Juho Kannala, Yoshua Bengio, and David Lopez-Paz. Interpolation consistency training for semi-supervised learning. *arXiv preprint arXiv:1903.03825*, 2019. 3
- [43] Zhou Wang, Alan C Bovik, Hamid R Sheikh, and Eero P Simoncelli. Image quality assessment: from error visibility to structural similarity. *IEEE transactions on image processing*, 13(4):600–612, 2004. 6
- [44] Chao-Yuan Wu, Nayan Singhal, and Philipp Krahenbuhl. Video compression through image interpolation. In *Proceedings of the European Conference on Computer Vision*, pages 416–431, 2018. 1
- [45] Tianfan Xue, Baian Chen, Jiajun Wu, Donglai Wei, and William T Freeman. Video enhancement with task-oriented flow. *International Journal of Computer Vision*, 127:1106–1125, 2019. 6
- [46] Yuan Yuan, Wei Su, and Dandan Ma. Efficient dynamic scene deblurring using spatially variant deconvolution network with optical flow guided training. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 3555–3564, 2020. 1
- [47] Guozhen Zhang, Yuhuan Zhu, Haonan Wang, Youxin Chen, Gangshan Wu, and Limin Wang. Extracting motion and appearance via inter-frame attention for efficient video frame interpolation. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 5682–5692, 2023. 6
- [48] Hongyi Zhang, Moustapha Cisse, Yann N. Dauphin, and David Lopez-Paz. mixup: Beyond empirical risk minimization. In *International Conference on Learning Representations*, 2018. 3
- [49] Jiawei Zhang, Jinshan Pan, Daoye Wang, Shangchen Zhou, Xing Wei, Furong Zhao, Jianbo Liu, and Jimmy Ren. Deep dynamic scene deblurring from optical flow. *IEEE Transac-*

tions on Circuits and Systems for Video Technology, 32(12): 8250–8260, 2021. 1

- [50] Richard Zhang, Phillip Isola, Alexei A Efros, Eli Shechtman, and Oliver Wang. The unreasonable effectiveness of deep features as a perceptual metric. In *Proceedings of the IEEE conference on computer vision and pattern recognition*, pages 586–595, 2018. 5, 6
- [51] Shengyu Zhao, Yilun Sheng, Yue Dong, Eric I Chang, Yan Xu, et al. Maskflownet: Asymmetric feature matching with learnable occlusion mask. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 6278–6287, 2020. 4