

Rethinking Optical Flow from Geometric Matching Consistent Perspective

Qiaole Dong*, Chenjie Cao*, Yanwei Fu†
 School of Data Science, Fudan University

{qldong18, 20110980001, yanweifu}@fudan.edu.cn

Abstract

Optical flow estimation is a challenging problem remaining unsolved. Recent deep learning based optical flow models have achieved considerable success. However, these models often train networks from the scratch on standard optical flow data, which restricts their ability to robustly and geometrically match image features. In this paper, we propose a rethinking to previous optical flow estimation. We particularly leverage Geometric Image Matching (GIM) as a pre-training task for the optical flow estimation (MatchFlow) with better feature representations, as GIM shares some common challenges as optical flow estimation, and with massive labeled real-world data. Thus, matching static scenes helps to learn more fundamental feature correlations of objects and scenes with consistent displacements. Specifically, the proposed MatchFlow model employs a QuadTree attention-based network pre-trained on MegaDepth to extract coarse features for further flow regression. Extensive experiments show that our model has great cross-dataset generalization. Our method achieves 11.5% and 10.1% error reduction from GMA on Sintel clean pass and KITTI test set. At the time of anonymous submission, our MatchFlow(G) enjoys state-of-the-art performance on Sintel clean and final pass compared to published approaches with comparable computation and memory footprint. Codes and models will be released in <https://github.com/DQiaole/MatchFlow>.

1. Introduction

This paper studies optical flow estimation, which is the problem of estimating the per-pixel displacement vector between two frames. It is very useful to various real-world applications, such as video frame interpolation [24], video inpainting [17], and action recognition [48]. The recent

direct-regression based methods [16, 22, 25, 34, 44, 50] have achieved great success by using powerful deep models, especially the recent transformers [20, 54]. Among them, RAFT [50] employs a convolutional GRU for iterative refinements, which queries local correlation features from a multi-scale 4D correlation volume. And GMA [25] further proposes a global motion aggregation module based on the self-similarity of image context, which greatly improves the performance within the occluded regions without degrading the performance in non-occluded regions. Typically, these models often train networks from the scratch on standard optical flow data, with the matching module (correlation volume) to help align the features of different images/frames. Generally, these current optical flow estimation algorithms still can not robustly handle several intractable cases, *e.g.*, small and fast-moving objects, occlusions, and textureless regions, as these estimators have very restricted ability of robustly learning the local image feature correspondence of different frames.

In this paper, we aim to provide a rethinking to the importance of Geometric Image Matching (GIM) to the optical flow estimation. In particular, despite GIM being designed to deal with the geometrically matching of static scenes, it indeed shares some common challenges with optical flow estimation, such as large displacement and appearance change [52]. Thus, we advocate that the deep models for optical flow estimation should be trained from matching static scene pairs with consistent displacements. This can potentially help these models to learn the local low-level features and color correlations at the early stages of networks, before extracting the priors for 3D multi-object motion. Furthermore, compared to optical flow data, it is much easier and simple to collect the real-world GIM data [14, 27], labeled by camera poses and depth computed from ground-truth or pioneering multi-view stereo manners [40]. Such extensive real-world data largely improves the generalization of optical flow.

On the other hand, we can also emphasize a rethinking of the general training pipeline of optical flow. Specifically, since the creative work of FlowNet2 [22], optical flow models [25, 42, 42] are often trained following the schedule

*Equal contributions.

†Corresponding author. The author is also with Shanghai Key Lab of Intelligent Information Processing, and Fudan ISTBI-ZJNU Algorithm Centre for Brain-inspired Intelligence, Zhejiang Normal University, Jinhua, China.




Figure 1. Qualitative results on KITTI test set. Red dashed boxes mark the regions of substantial improvements. Please zoom in for details.

of *Curriculum Learning* [2], *i.e.*, from FlyingChair [16] to FlyThings3D [31] and finally to Sintel [7] or KITTI [18]. Nevertheless, the motion contained in FlyingChair is still far from the simplest scenario. Empirically, the static scene under viewpoint changes of GIM [46, 49], can also be taken as one special type of optical flow estimation, which is even much simpler than FlyingChair. Therefore, it is reasonable to take GIM amenable for being the very first stage of the curriculum learning pipeline. Essentially, as mentioned above, the GIM data can be easily collected at large-scale, and thus will greatly benefit the learning of deep models.

Formally, this paper well elaborates on the key idea of taking GIM as the prefixed task for optical flow. We draw the inspiration from recent GIM works, and present a novel MatchFlow that can effectively generalize the pre-trained image matching module to estimate optical flow. The key component of MatchFlow is a new module of Feature Matching Extractor, composed of Resnet-16 and 8 interleaving self/cross-Quadtree attention blocks, trained on GIM task and used to get the 4D correlation volume. After Feature Matching Extractor, our MatchFlow still takes the common practice of using GRUs module to handle optical flow estimation, with an optional GMA modelling the context information. In this paper, we denote the full model based on GMA as **MatchFlow(G)**, and the model without GMA module (a.k.a. RAFT) as **MatchFlow(R)**.

Following the standard optical flow training procedure [25, 50], we conduct extensive experiments on FlyingChair [16], FlyingThings3D [31], Sintel [7], and KITTI [18]. Experiments results show that MatchFlow enjoys good performance and great cross-dataset generalization. Formally, RAFT-based MatchFlow(R) obtains an F1-all error 13.6% on KITTI training set after being trained on synthetic datasets. In addition, GMA-based MatchFlow(G) achieves 11.5% and 10.1% error reduction from GMA on the Sintel clean pass and KITTI test set. The qualitative comparison on KITTI also shows the superior performance of MatchFlow as in Fig. 1. Ablation studies verify that GIM pre-training indeed helps to learn better feature representation for optical flow estimation.

We highlight our contributions as follows. (1) We reformulate the optical flow pipeline, and propose the idea

of employing GIM as the prelude task for optical flow. This offers a rethinking to the learning based optical flow estimation. (2) We further present a novel matching-based optical flow estimation model – MatchFlow, which has the new module of Feature Matching Extractor, learned by the GIM pre-training task for optical flow estimation. Accordingly, the pipeline of curriculum learning has also been updated to effectively train our MatchFlow. (3) We introduce the massive real-world matching data to train our model. And thus our model can extract robust features to handle with the consistent motion of scenes, and common challenges faced by both tasks. (4) We conduct extensive experiments and ablation studies to show that both the matching based pre-training and interleaving self/cross-attention modules are critical for the final optical flow performance. The proposed model shows great cross-dataset generalization and better performance over several optical flow competitors on several standard benchmarks.

2. Related Work

Optical Flow Estimation. Traditionally, optical flow estimation [3–6, 19, 36, 55] is treated as an energy minimization problem. Nowadays, deep models [16, 21, 22, 44] formulate the optical flow estimation as a regression problem with an end-to-end trainable Convolutional Neural Network. More recently, deep learning based iterative refinement for optical flow [21, 25, 42, 44, 50, 57] has resulted in a big breakthrough. RAFT [50] constructs a 4D multi-scale correlation volume and utilizes a convolutional GRU block as an update operator. And GMA [25] further proposes a global motion aggregation module to tackle the occlusion problem. Besides, recent methods [20, 29, 30, 47, 58] mainly focus on the recurrent decoder of RAFT. These methods are orthogonal to ours as we mainly focus on GIM pretraining of feature encoder and provide a rethinking for optical flow from the geometric matching consistent scene perspective.

Geometric Image Matching. Geometric image matching [28, 35, 38] tends to find correspondence among images with different views. Different from optical flow estimation, image matching usually assumes that the scene is static and the geometric displacements are due to the change

of viewpoints (camera pose). Among recent detector-free matching methods [11, 12, 37, 46, 49], Tang *et al.* [49] propose QuadTree attention to capture both fine-level details and long-range dependencies, which outperforms the Linear attention [26] used in [46]. In this work, we also employ a QuadTree attention-based network as our feature extractor to learn feature correlation between two-view images. Benefited from the matching pre-training, we can get much better flow estimation within non-occluded regions. Furthermore, global motion aggregation module [25] can help propagate the accurate flow within non-occluded regions to more challenging occluded ones, which can boost the performance of optical flow estimation a lot.

Besides, there is also a line of works [1, 51, 52, 54, 57] which try to reduce the gap between optical flow estimation and GIM. GMFlow [54] formulates the optical flow estimation as a dense global matching problem, while GMFlowNet [57] introduces the matching loss and matching initialization into the optical flow. On the other hand, Depthstill [1] constructs matching pairs based on estimated depth for direct optical flow training and achieves superior generalization to unseen real data. We should clarify that our work is different from them because MatchFlow focuses on the GIM based pre-training for better feature representations from the perspective of curriculum learning [2]. Moreover, MatchFlow outperforms these competitors with superior generalization on Sintel.

Curriculum Learning. Large datasets are one of the driving forces of deep learning. However, due to its difficult to provide reliable ground truth optical flow labels, there is limited target data. Thus optical flow models typically rely on the pre-training on the synthetic FlyingChair [16] and FlyingThings3D [31] dataset with the curriculum learning [2]. In contrast, we give a rethinking to this pipeline and propose using GIM as a pre-training task for optical flow.

3. Methodology

Problem Setup. Given a pair of consecutive images, I_1, I_2 , optical flow estimation is the task of estimating per-pixel dense displacement field (f^1, f^2) which maps the location (u, v) in I_1 to I_2 as $(u + f^1(u), v + f^2(v))$.

Pipeline Overview. The training pipeline is shown in Fig. 2, which includes two stages. For stage one, we first train the Feature Matching Extractor (FME) \mathcal{F} on massive real-world GIM data to learn static scene matching caused by viewpoint changes (Sec. 3.1). Thus given a pair of images I_1, I_2 , the coarse 4D correlation volume can be achieved with $C = \mathcal{F}(I_1, I_2)$. For stage two, we follow [22, 50], and jointly finetune the FME \mathcal{F} with the specialized Optical Flow Estimator (OFE) \mathcal{O} to get the final displacement field as $(f^1, f^2) = \mathcal{O}(\mathcal{F}(I_1, I_2))$ (Sec. 3.2). Stage two is gradually applied on FlyingChair, FlyingThings3D, Sintel, and KITTI with increasing difficulty. Such a

progressive pipeline makes MatchFlow address intractable object motion within complex scenes.

3.1. Geometric Image Matching Pre-training

First, we elaborate on the idea of training FME on GIM task in this section. Previous methods [22, 25, 42, 50, 57] are trained on optical flow data directly, without learning the feature matching through static scenes. FlowNet2 [22] finds that not only the quality of data is important, but also the data presenting order during the optical flow training is critical, which is closely related to the curriculum learning [2]. It is trained on FlyingChair first, then trained on FlyingThings3D, and finally fine-tuned on target data Sintel. Among them, FlyingChair is the simplest dataset, composed of 22k image pairs of chairs superimposed on randomly selected background images. Furthermore, FlyingChair only contains planar motions. FlyingThings3D contains about 88k image pairs generated by rendering random scenes of 3D models from ShapeNet [39] in front of static 3D backgrounds, which manifests real 3D motion and lighting effects with more challenges.

However, as both FlyingChair and FlyingThings3D are synthetic datasets, there is a lack of authenticity. Moreover, motions among different objects or between foregrounds and backgrounds in FlyingChair are not consistent. Thus FlyingChair is still not the ‘simplest motion’ for curriculum learning. Therefore, we propose a simpler scenario to better assist the early representation learning of the network: optical flow estimation of consistent matching scenes, which can be regarded as GIM. Besides, GIM enjoys much larger real-world datasets with various displacement and appearance changes, *e.g.*, MegaDepth [27] and Scannet [14] with reliable enough matching labels.

On the other hand, limited by the hardware [14], algorithm precision, and data quality [27], GIM data is usually labeled sparsely with only non-occluded regions. So we only train FME on GIM data. Specifically, we extract the feature F_1 and F_2 from the reference frame and matching frame respectively on 1/8 of input resolution. And stacked QuadTree attention blocks [49] are incorporated into FME with self and cross-attention to enhance the feature learning between the image pair efficiently. Such an interleaving self and cross-attention manner is critical for the final optical flow performance as verified in Tab. 2. We then construct the 4D correlation volume C as follows:

$$C(i, j) = \langle F_1(i), F_2(j) \rangle \in \mathbb{R}^{H \times W \times H \times W}, \quad (1)$$

where i, j indicate the index of feature map F_1 and F_2 ; H, W indicate 1/8 height and width of the input image.

Pre-training Loss. We apply the dual-softmax [37, 53] to get the probability before the optimization. Specifically, the matching probability \mathcal{P}_c is calculated as:

$$\mathcal{P}_c(i, j) = \text{softmax}(C(i, \cdot)/\tau)_j \cdot \text{softmax}(C(\cdot, j)/\tau)_i, \quad (2)$$

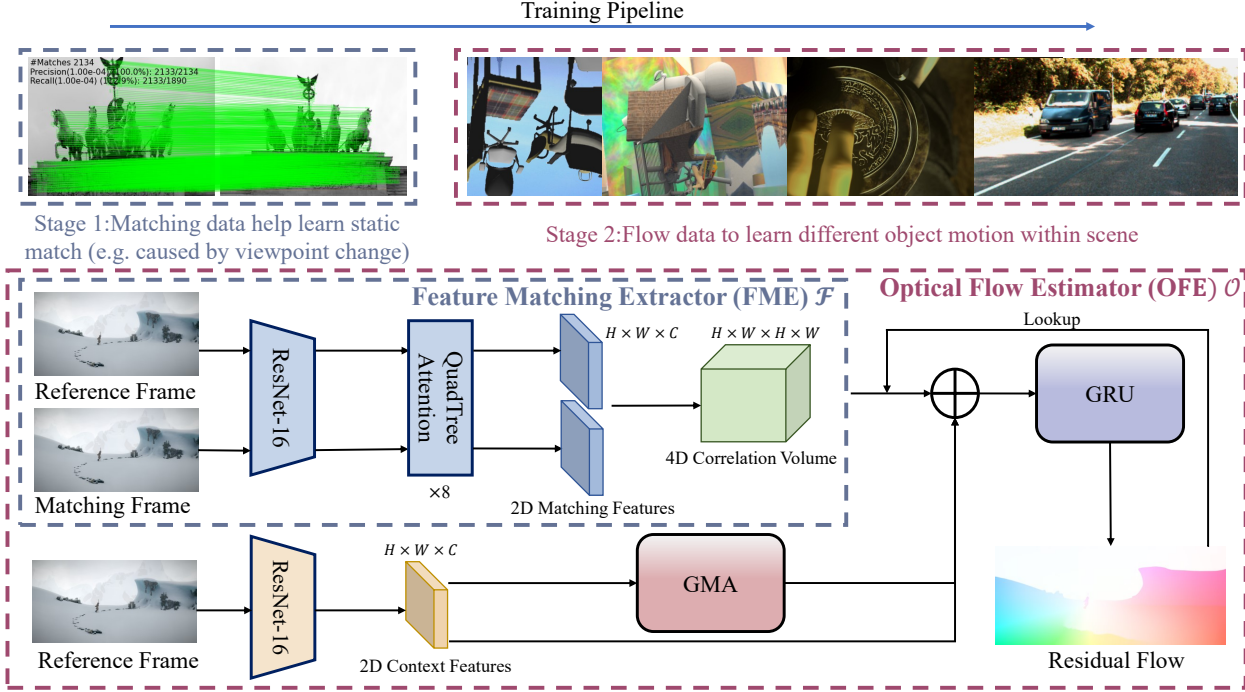


Figure 2. Overview of our MatchFlow. The simplified training pipeline is shown at the top, while details of each stage are specifically listed below. H, W indicate $1/8$ height and width of the input image respectively. Here the GMA means global motion aggregation module [25].

where temperature $\tau = 0.1$. We get ground-truth corresponding matches \mathcal{M}_c^{gt} through depth and camera poses from GIM data in $1/8$ -resolution. Our loss function \mathcal{L}_c is:

$$\mathcal{L}_c = -\frac{1}{|\mathcal{M}_c^{gt}|} \sum_{(i,j) \in \mathcal{M}_c^{gt}} \log \mathcal{P}_c(i, j). \quad (3)$$

We also use additional l_2 loss in $1/2$ -resolution for fine-grained supervision [46].

3.2. Optical Flow Refinement

After the pre-training of FME on matching dataset, we then finetune our FME with OFE on FlyingChair, FlyingThings3D, Sintel, and KITTI following previous data schedule [22, 25, 42, 50, 57]. We employ the successful RAFT [50] and GMA [25] as our OFE as shown in Fig 2. Given I_1 and I_2 , we first use FME to extract 2D matching features F_1, F_2 and calculate the inner product to get a multi-scale 4D correlation volume. QuadTree attention blocks are utilized in FME to learn more discriminative features within non-occluded regions. Furthermore, the global motion aggregation module [25] is used to propagate accurate motion features within non-occluded regions learned from matching data to occluded ones based on the self-similarity of 2D context features, which can successfully resolve ambiguities caused by occlusions. Finally, GRU can learn to lookup from the correlation volume and iteratively decode

the motion feature and context feature into a series of residual flows. Sum of the residual flows is the final output flow. **Optical Flow Loss.** Given the partially summed residual flows $\{f_1, \dots, f_N\}$, we supervise our network use l_1 loss following previous work [50]:

$$\mathcal{L} = \sum_{i=1}^N \gamma^{N-i} \|f_{gt} - f_i\|_1, \quad (4)$$

where f_{gt} is the ground truth flow, $\gamma = 0.8$ on FlyingChair and FlyingThings3D, and 0.85 on Sintel and KITTI.

4. Experiments

Implementation Details. We utilize MegaDepth [27] as our matching dataset and train the feature encoder with randomly sampled 36,800 image pairs per epoch for a total of 30 epochs [46]. We then follow the standard training pipeline [25, 50] of first training our model on FlyingChair (C) for 120k iterations (batch size 8) and FlyingThings3D (T) for another 120k iterations (batch size 6). Following training on C+T, we finetune the model using data from clean and final passes of Sintel (S), KITTI (K), and HD1K (H). Our method is implemented with Pytorch [33] on 2 RTX 3090 GPUs. And we use the one-cycle learning rate [41], setting the highest learning rate for FlyingChairs to 2.5×10^{-4} and the remaining to 1.25×10^{-4} . We denote the full model as **MatchFlow(G)**, and the model without GMA module as **MatchFlow(R)**.

4.1. Sintel

Following [50], we first evaluate the generalization error on the Sintel training set after training on C+T. As shown in Tab. 1, MatchFlow(G) achieves a 20.8% reduction in Average End-Point-Error (AEPE) of Sintel clean pass from 1.30 to 1.03 based on GMA. For AEPE of the final pass, MatchFlow(G) achieves a 10.6% reduction from 2.74 to 2.45 based on GMA. MatchFlow(R) also improves the RAFT by 20.3% and 3.7%. These results indicate the great generalization of our method. For the Sintel test set, MatchFlow(G) and MatchFlow(R) improve the GMA and RAFT by 11.5% and 17.4% on the clean pass. And MatchFlow(G) ranks first on Sintel clean pass among all published works. On Sintel final pass, MatchFlow(G) obtains an AEPE of 2.37 and outperforms GMA by 0.1 pixels. We also provide qualitative results on Sintel from the official website in Fig. 3. The flow estimation from MatchFlow(G) exhibits clearer details (first, second, and fourth rows) and captures small and fast-moving birds in the third row.

4.2. KITTI

For the KITTI training set, MatchFlow(G) reduces the generalization error EPE and FI-all of GMA by 13.0% and 8.8% as shown in Tab. 1. MatchFlow(R) even achieves the best generalization performance with FI-all 13.6% with slightly worse EPE compared to MatchFlow(G). FI-all here refers to the percentage of optical flow vectors whose EPE is larger than 3 pixels or 5% of the ground truth magnitude. Note that the FI-all gap between MatchFlow(G) and MatchFlow(R) is somehow reasonable. As KPA-Flow [29] found that GMA’s attention map contains many globally misleading clues. These globally misleading clues can lead to more large errors. For the KITTI test set, MatchFlow(R) and MatchFlow(G) obtain FI-all of 4.72% and 4.63% respectively, which reduce the error of the base model by 7.5% and 10.1%. Fig. 1 provides two qualitative examples from the KITTI test set. Compared with GMA, our MatchFlow(G) can successfully recognize the small signage in the first image and reduces the flow leakage around car boundaries in the second image. These results show the effectiveness and superiority of our GIM pre-training.

4.3. Ablation Studies

We conduct a series of ablation studies on MatchFlow with training on C+T and evaluating on the test set of FlyingThings3D and the training set of Sintel and KITTI. The results are shown in Tab. 2. Since KITTI has a special aspect ratio, we perform more KITTI ablations in Tab. 3.

Geometric Image Matching Pre-training. We first assess the effect of GIM pre-training with the same network structure. As seen in Tab. 2, solely adding more parameters for FME only brings marginal improvement, and performs much worse than GMA on the Sintel final pass. With

GIM pre-training, we can get much better performance on all datasets, except for the slightly worse FI-all of KITTI, which validates the benefit of pre-training.

Number of Attention Blocks. We attempt to reduce the attention blocks in Tab. 2. When the number of attention blocks was reduced, the performance decreased sharply on Sintel while slowly on KITTI. When there are no attention blocks, our FME is a pure CNN-based network. Although FME is pre-trained with GIM, it performs even worse on Sintel than random initialized GMA. This may be caused by the network ‘overfitting’ on the matching task. A recent work [32] suggests that attention layers can smooth out irregularities and high-frequency noises with a better generalization. We thus use 8 QuadTree attention blocks by default. Furthermore, to compare GMA and our model with aligned settings, we also provide the results of GMA-L in Tab. 2. GMA-L is a large version of GMA by doubling the channels. And the results of GMA-L are taken from [20]. However, its performance is much worse than our model with more parameters. To sum up, the new design of GIM pre-training and interleaved attention-based FME outperforms the one with trivially extended model scales.

Attention Type. We also try Linear attention [26] and Global-Local Attention (GLA) [11] with comparable parameters in this subsection. Except for KITTI’s FI-all of GLA, Linear attention and GLA are inferior to QuadTree attention with much worse results on FlyingThings3D and cross-dataset generalization on Sintel.

Global Matching. Finally, we compare the technique proposed by GMFlowNet [57], where match initialization (Match Init.) refers to the usage of robust matching point based on 4D correlation volume to initialize the flow. Moreover, Match Loss indicates further optimizing the flow network with matching loss [37, 53] when trained on optical flow datasets. Tab. 2 shows that after training on the GIM datasets, either adding match initialization or match loss on the flow dataset can not provide better performance.

More Ablations on KITTI. As KITTI owns a much smaller aspect ratio, it is difficult for our attention block to generalize to such images as verified in the first row of Tab. 3. So we use the tile technique proposed by Perceiver IO [23]. Formally, we split the KITTI images into 2 smaller overlapping sub-images and weighted average the results on overlapping regions. The tile technique can reduce the FI-all from 18.2% to 16.8%. Furthermore, recent multi-scale training methods [8, 10, 15] can be generalized to much higher resolution after being trained on multi-scale images with few large-scale ones. We thus try to train our model on several different resolution images to generalize better on KITTI. However, it only achieves marginal improvement with much more computational resources as shown in Tab. 3. And we do not use it in our final model. Finally, we find that simply zero-padding test images discourages

Table 1. Quantitative comparison on standard benchmark. ‘A’ indicates the AutoFlow [43] dataset. ‘C+T’: Succeeding training on FlyingChairs (C) and FlyingThings3D (T), we evaluate the capacity of generalization on Sintel (S) and KITTI (K) training sets. ‘C+T+S+K+H’: Training samples from T, S, K, and HD1K (H) are included in our training set for further finetuning. Results on training set are shown in parentheses. The top and second-place results are bolded and underlined, respectively. † indicates tile technique [23]. And * indicates evaluating with RAFT’s “warm-start” strategy.

Training	Method	Sintel (train)		KITTI-15 (train)		Sintel (test)		KITTI-15 (test)
		Clean	Final	Fl-epe	Fl-all	Clean	Final	Fl-all
A	Perceiver IO [23]	1.81 [†]	2.42 [†]	4.98 [†]	-	-	-	-
	RAFT-A [43]	1.95	2.57	4.23	-	-	-	-
C+T	PWC-Net [44]	2.55	3.93	10.35	33.7	-	-	-
	FlowNet2 [22]	2.02	3.54	10.08	30.0	3.96	6.02	-
	RAFT [50]	1.43	2.71	5.04	17.4	-	-	-
	Separable Flow [56]	1.30	2.59	4.60	15.9	-	-	-
	GMA [25]	1.30	2.74	4.69	17.1	-	-	-
	AGFlow [30]	1.31	2.69	4.82	17.0	-	-	-
	KPA-Flow [29]	1.28	2.68	4.46	15.9	-	-	-
	DIP [58]	1.30	2.82	4.29	<u>13.7</u>	-	-	-
	GMFlowNet [57]	1.14	2.71	4.24	15.4	-	-	-
	GMFlow [54]	1.08	2.48	7.77	23.40	-	-	-
	CRAFT [42]	1.27	2.79	4.88	17.5	-	-	-
	FlowFormer [20]	1.01	2.40	4.09 [†]	14.7 [†]	-	-	-
	SKFlow [47]	1.22	2.46	4.27	15.5	-	-	-
	MatchFlow(R) (Ours)	1.14	2.61	4.19 [†]	13.6[†]	-	-	-
	MatchFlow(G) (Ours)	<u>1.03</u>	<u>2.45</u>	4.08[†]	15.6 [†]	-	-	-
C+T+S+K+H	PWC-Net+ [45]	(1.71)	(2.34)	(1.50)	(5.3)	3.45	4.60	7.72
	RAFT [50]	(0.76)	(1.22)	(0.63)	(1.5)	1.61*	2.86*	5.10
	RAFT-A [43]	-	-	-	-	2.01	3.14	4.78
	Separable Flow [56]	(0.69)	(1.10)	(0.69)	(1.60)	1.50	2.67	4.64
	GMA [25]	(0.62)	(1.06)	(0.57)	(1.2)	1.39*	2.47*	5.15
	AGFlow [30]	(0.65)	(1.07)	(0.58)	(1.2)	1.43*	2.47*	4.89
	KPA-Flow [29]	(0.60)	(1.02)	(0.52)	(1.1)	1.35*	2.36*	<u>4.60</u>
	DIP [58]	-	-	-	-	1.44*	2.83*	4.21
	GMFlowNet [57]	(0.59)	(0.91)	(0.64)	(1.51)	1.39	2.65	4.79
	GMFlow [54]	-	-	-	-	1.74	2.90	9.32
	CRAFT [42]	(0.60)	(1.06)	(0.58)	(1.34)	1.45*	2.42*	4.79
	FlowFormer [20]	(0.48)	(0.74)	(0.53)	(1.11)	<u>1.20</u>	2.12	4.68 [†]
	SKFlow [47]	(0.52)	(0.78)	(0.51)	(0.94)	1.28*	<u>2.23*</u>	4.84
	MatchFlow(R) (Ours)	(0.51)	(0.81)	(0.59)	(1.3)	1.33*	2.64*	4.72 [†]
	MatchFlow(G) (Ours)	(0.49)	(0.78)	(0.55)	(1.1)	1.16*	2.37*	4.63 [†]

the final performance of attention-based models in KITTI. Because we employ the QuadTree attention, the image resolution should be multiplied by 32 instead of 8 compared with previous methods. The traditional zero-padding brings much thicker black borders and degrades the performance a lot in extreme cases [42]. So we use bilinear interpolation to adjust the resolution with the tile technique for better performance as shown in Tab. 3.

4.4. Where are Gains Coming from?

Occlusion Analysis. To determine where our gains come from, we use the occlusion map provided by Sintel on the training set to divide the images into occluded (‘Occ’) and non-occluded (‘Noc’) areas. Following GMA [25], we further divide the occluded area into in-frame (‘Occ-in’) and out-of-frame (‘Occ-out’) occlusions. Using the model

trained on FlyingThings3D, we evaluate it on three rendering passes of Sintel. As shown in Tab. 4, the relative improvement of our method compared to the baseline model, *i.e.*, GMA, is mainly attributed to the error reduction within the Noc region. We further visualize the correlation volume on 100 Sintel final pass images with the model trained on FlyingThings3D in Fig. 4. The first row is for GMA, and the second row is MatchFlow(G). The much higher central value of our method validates that GIM pre-training can benefit optical flow a lot, especially within the Noc region.

Real-world Dataset based Training. We further supply a comparison of generalization error to Depthstill [1], self-supervised pre-trained DINO feature [9], and class-supervised Twins-SVT feature [13]. We follow Depthstill [1] to generate 20K virtual images and flow maps on MegaDepth (dubbed **dMegaDepth**) to finetune the model.

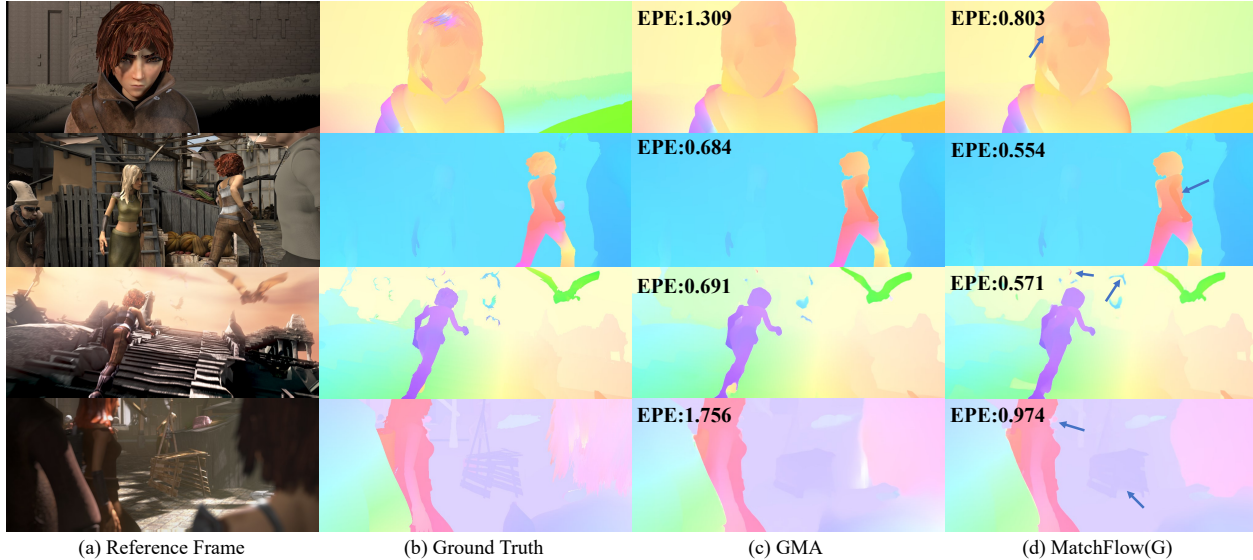


Figure 3. Qualitative comparison on Sintel test set. First two rows are from clean pass, and the last two from final pass. Notable areas are pointed out by arrows. Please zoom in for details.

Table 2. Ablation studies. Settings used in our final model are underlined. * indicates a large version of GMA, results from [20].

Experiment	Method	Things (test)		Sintel (train)		KITTI-15 (train)		Parameters
		Clean	Final	Clean	Final	Fl-epe	Fl-all	
Baseline	GMA [25]	3.14	2.80	1.30	2.74	4.69	17.1	5.9M
	GMA-L* [20]	-	-	1.33	2.56	4.40	15.9	17.0M
Local Matching Pre-training	No	2.82	2.56	1.27	2.84	4.12	14.4	15.4M
	Yes	2.12	2.07	1.03	2.45	4.08	15.6	15.4M
Number of Attention Blocks	0	2.97	2.81	1.36	2.98	4.58	16.3	5.9M
	4	2.65	2.40	1.18	2.79	4.55	15.6	9.1M
	<u>8</u>	2.12	2.07	1.03	2.45	4.08	15.6	15.4M
Attention Type	Linear Atten.	2.82	2.56	1.24	3.00	4.32	15.6	14.3M
	<u>QuadTree Atten.</u>	2.12	2.07	1.03	2.45	4.08	15.6	15.4M
	GLA	2.66	2.55	1.24	2.82	4.29	15.0	18.7M
Global Matching	Match Init.	2.23	2.29	1.44	2.93	5.30	17.4	15.4M
	Match Loss	2.22	2.29	1.00	2.50	4.78	16.1	15.4M
	<u>No</u>	2.12	2.07	1.03	2.45	4.08	15.6	15.4M

Table 3. More Ablations on KITTI using MatchFlow(G) trained on C+T. We use Resize+Tile in our final model.

Resolution		Training	Test	KITTI-15 (train)	
Pad	Resize	Multiscale	Tile	Fl-epe	Fl-all
✓				5.59	18.2
✓			✓	4.67	16.8
✓		✓		5.43	18.0
✓		✓	✓	4.61	16.5
	✓	✓		4.35	15.7
	✓	✓	✓	4.04	15.2
	✓			4.54	16.3
	✓		✓	<u>4.08</u>	<u>15.6</u>

Our method still performs best among them in more complex Sintel. And feature encoders trained by DINO [9] and classification also perform worse than ours on both Sintel

and KITTI. In contrast, the motion types of objects in KITTI are much simpler than Sintel, such as mostly forward translations or steerings. Depthstill [1] greatly benefited from the last dense matching-based flow training on KITTI, while greatly degrading the results on Sintel.

4.5. Parameters, Timing, and GPU Memory

We show the computational overhead of our models in Tab. 6. The parameter numbers of MatchFlow(R) and MatchFlow(G) are 14.8M and 15.4M respectively, which mainly attribute to 8 QuadTree attention blocks. We also test the GPU training memory on FlyingThings3D with 2 48G RTX A6000 GPUs. Since FlowFormer [20] is trained on FlyingThings3D with image size 432×960 to get better cross-dataset performance on Sintel, we can only support batch size 2 per GPU with peak memory consumption

Table 4. Optical flow error within occluded (‘Occ’) and non-occluded (‘Noc’) regions for three Sintel datasets. Each dataset’s top outcomes and biggest relative improvement are shown in bold. Models are trained using the FlyingThings3D training set.

Sintel Pass	Type	GMA (AEPE)	MatchFlow(G) (AEPE)	Rel. Impr. (%)
Clean	Noc	0.58	0.44	24.1
	Occ	10.58	8.51	19.6
	Occ-in	7.68	6.42	16.4
	Occ-out	12.52	9.63	23.1
	All	1.30	1.03	20.8
Final	Noc	1.72	1.46	15.1
	Occ	17.33	15.00	13.4
	Occ-in	14.96	12.31	17.7
	Occ-out	16.44	15.32	6.8
	All	2.74	2.45	10.6
Albedo	Noc	0.48	0.37	22.9
	Occ	9.54	7.98	16.4
	Occ-in	7.51	6.37	15.2
	Occ-out	10.22	8.39	17.9
	All	1.15	0.92	20.0




Figure 4. Visualizations of normalized correlation volumes within different region on 100 Sintel final pass images. The first row is for GMA, and the second row is MatchFlow(G). $(\Delta u, \Delta v)$ are the offsets away from ground truth flow. The center value (corresponding to ground truth flow) is $2.0\times$ and $1.6\times$ higher for our method vs. GMA within Noc and Occ regions respectively.

43.7G. And the training image size of GMA and our models are 400×720 and 416×736 (as our resolution should be multiplied by 32 for QuadTree attention) respectively. Testing with batch size 3 per GPU, our models consume around 22.1 and 23.6G memory. And our training time is around 58 hours. We test the inference time on Sintel with a single RTX A6000 GPU. GMA takes 74ms on average while MatchFlow(G) takes 126ms per image pair. The GRU iteration number is set to 12 for all models. The overall com-

Table 5. Generalization error of real-world dataset based training.

Method	Dataset	Sintel		KITTI-15	
		Clean	Final	Fl-epe	Fl-all
Baseline	C+T	1.27	2.84	4.12	14.4
Depthstill	C+T+dMegaDepth	2.10	3.47	3.41	10.9
DINO	ImageNet+C+T	1.52	3.08	5.50	19.2
Twins-SVT	ImageNet+C+T	1.15	2.73	4.98	16.8
Ours	MegaDepth+C+T	1.03	2.45	4.08	15.6

Table 6. Parameters, inference and training time, and memory.

Model	Para.	Infer.	Train.	Memory
GMA [25]	5.9M	74 ms	30 h	11.7 G
FlowFormer [20]	18.2 M	149 ms	107 h	43.7 G
MatchFlow(R)	14.8 M	110 ms	53 h	22.1 G
MatchFlow(G)	15.4M	126 ms	58 h	23.6 G




Figure 5. Failure cases on Sintel final pass.

putational overhead of MatchFlow is modest compared to FlowFormer, while our performance is still competitive.

4.6. Failure Cases and Limitations

Though we achieve impressive results on KITTI and Sintel, we fail to generalize well to some extreme situations. Fig. 5 shows two failed cases. In the first case, our model ignores the right part of the weapon caused by motion blur. And the motion of shadow is estimated mistakenly in the second case. We conjecture that this is caused by inadequate training on such blur data, which discourages the performance of attention blocks. We believe that training on more diverse final-type data can improve the results.

5. Conclusion

This paper proposed to use geometric image matching as the pre-training task for learning more robust matching features for optical flow estimation. Sharing some common challenges with optical flow estimation while with much more real-world data, geometric image matching pre-training can help learn fundamental feature correlations of scenes. To improve the feature matching extractor, we further proposed adding stacked QuadTree attention layers for self and cross-attention learning. Extensive experiments showed great cross-dataset generalization and superior performance of our models. Future works may focus on reducing the attention layers in feature matching extractor while preserve competitive performance.

Acknowledgements. This work was supported by Shanghai Municipal Science and Technology Major Project (2018SHZDZX01).

References

- [1] Filippo Aleotti, Matteo Poggi, and Stefano Mattoccia. Learning optical flow from still images. In *2021 IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*, pages 15196–15206, 2021. 3, 6, 7
- [2] Yoshua Bengio, Jérôme Louradour, Ronan Collobert, and Jason Weston. Curriculum learning. In *Proceedings of the 26th annual international conference on machine learning*, pages 41–48, 2009. 2, 3
- [3] Michael J Black and Padmanabhan Anandan. A framework for the robust estimation of optical flow. In *1993 (4th) International Conference on Computer Vision*, pages 231–236. IEEE, 1993. 2
- [4] Michael J Black and Paul Anandan. The robust estimation of multiple motions: Parametric and piecewise-smooth flow fields. *Computer vision and image understanding*, 63(1):75–104, 1996. 2
- [5] Thomas Brox, Andrés Bruhn, Nils Papenberg, and Joachim Weickert. High accuracy optical flow estimation based on a theory for warping. In *European conference on computer vision*, pages 25–36. Springer, 2004. 2
- [6] Andrés Bruhn, Joachim Weickert, and Christoph Schnörr. Lucas/kanade meets horn/schunck: Combining local and global optic flow methods. *International journal of computer vision*, 61(3):211–231, 2005. 2
- [7] Daniel J Butler, Jonas Wulff, Garrett B Stanley, and Michael J Black. A naturalistic open source movie for optical flow evaluation. In *European conference on computer vision*, pages 611–625. Springer, 2012. 2
- [8] Chenjie Cao, Xinlin Ren, and Yanwei Fu. Mvsformer: Learning robust image representations via transformers and temperature-based depth for multi-view stereo. *arXiv preprint arXiv:2208.02541*, 2022. 5
- [9] Mathilde Caron, Hugo Touvron, Ishan Misra, Hervé Jégou, Julien Mairal, Piotr Bojanowski, and Armand Joulin. Emerging properties in self-supervised vision transformers. In *2021 IEEE/CVF International Conference on Computer Vision (ICCV)*, pages 9630–9640, 2021. 6, 7
- [10] Lucy Chai, Michaël Gharbi, Eli Shechtman, Phillip Isola, and Richard Zhang. Any-resolution training for high-resolution image synthesis. In Shai Avidan, Gabriel Brostow, Moustapha Cissé, Giovanni Maria Farinella, and Tal Hassner, editors, *Computer Vision – ECCV 2022*, pages 170–188, Cham, 2022. Springer Nature Switzerland. 5
- [11] Hongkai Chen, Zixin Luo, Lei Zhou, Yurun Tian, Mingmin Zhen, Tian Fang, David Mckinnon, Yanghai Tsing, and Long Quan. Aspanformer: Detector-free image matching with adaptive span transformer. *arXiv preprint arXiv:2208.14201*, 2022. 3, 5
- [12] Christopher B Choy, JunYoung Gwak, Silvio Savarese, and Manmohan Chandraker. Universal correspondence network. *Advances in neural information processing systems*, 29, 2016. 3
- [13] Xiangxiang Chu, Zhi Tian, Yuqing Wang, Bo Zhang, Haibing Ren, Xiaolin Wei, Huaxia Xia, and Chunhua Shen. Twins: Revisiting the design of spatial attention in vision transformers. In M. Ranzato, A. Beygelzimer, Y. Dauphin, P.S. Liang, and J. Wortman Vaughan, editors, *Advances in Neural Information Processing Systems*, volume 34, pages 9355–9366. Curran Associates, Inc., 2021. 6
- [14] Angela Dai, Angel X. Chang, Manolis Savva, Maciej Halber, Thomas Funkhouser, and Matthias Nießner. Scannet: Richly-annotated 3d reconstructions of indoor scenes. In *Proc. Computer Vision and Pattern Recognition (CVPR)*, IEEE, 2017. 1, 3
- [15] Qiaole Dong, Chenjie Cao, and Yanwei Fu. Incremental transformer structure enhanced image inpainting with masking positional encoding. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 11358–11368, 2022. 5
- [16] 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 international conference on computer vision*, pages 2758–2766, 2015. 1, 2, 3
- [17] Chen Gao, Ayush Saraf, Jia-Bin Huang, and Johannes Kopf. Flow-edge guided video completion. In *European Conference on Computer Vision*, pages 713–729. Springer, 2020. 1
- [18] Andreas Geiger, Philip Lenz, and Raquel Urtasun. Are we ready for autonomous driving? the kitti vision benchmark suite. In *2012 IEEE conference on computer vision and pattern recognition*, pages 3354–3361. IEEE, 2012. 2
- [19] Berthold KP Horn and Brian G Schunck. Determining optical flow. *Artificial intelligence*, 17(1-3):185–203, 1981. 2
- [20] 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. *arXiv preprint arXiv:2203.16194*, 2022. 1, 2, 5, 6, 7, 8
- [21] Tak-Wai Hui, Xiaou Tang, and Chen Change Loy. Lite-flownet: A lightweight convolutional neural network for optical flow estimation. In *Proceedings of the IEEE conference on computer vision and pattern recognition*, pages 8981–8989, 2018. 2
- [22] 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 conference on computer vision and pattern recognition*, pages 2462–2470, 2017. 1, 2, 3, 4, 6
- [23] Andrew Jaegle, Sebastian Borgeaud, Jean-Baptiste Alayrac, Carl Doersch, Catalin Ionescu, David Ding, Skanda Koppala, Daniel Zoran, Andrew Brock, Evan Shelhamer, et al. Perceiver io: A general architecture for structured inputs & outputs. *arXiv preprint arXiv:2107.14795*, 2021. 5, 6
- [24] Huaizu Jiang, Deqing Sun, Varun Jampani, Ming-Hsuan Yang, Erik Learned-Miller, and Jan Kautz. Super slo-mo: High quality estimation of multiple intermediate frames for video interpolation. In *Proceedings of the IEEE conference on computer vision and pattern recognition*, pages 9000–9008, 2018. 1
- [25] Shihao Jiang, Dylan Campbell, Yao Lu, Hongdong Li, and Richard Hartley. Learning to estimate hidden motions with global motion aggregation. In *Proceedings of the IEEE/CVF*

- International Conference on Computer Vision*, pages 9772–9781, 2021. [1](#), [2](#), [3](#), [4](#), [6](#), [7](#), [8](#)
- [26] Angelos Katharopoulos, Apoorv Vyas, Nikolaos Pappas, and François Fleuret. Transformers are rnns: Fast autoregressive transformers with linear attention. In *International Conference on Machine Learning*, pages 5156–5165. PMLR, 2020. [3](#), [5](#)
- [27] Zhengqi Li and Noah Snavely. Megadepth: Learning single-view depth prediction from internet photos. In *Proceedings of the IEEE conference on computer vision and pattern recognition*, pages 2041–2050, 2018. [1](#), [3](#), [4](#)
- [28] David G Lowe. Distinctive image features from scale-invariant keypoints. *International journal of computer vision*, 60(2):91–110, 2004. [2](#)
- [29] Ao Luo, Fan Yang, Xin Li, and Shuaicheng Liu. Learning optical flow with kernel patch attention. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 8906–8915, 2022. [2](#), [5](#), [6](#)
- [30] Ao Luo, Fan Yang, Kunming Luo, Xin Li, Haoqiang Fan, and Shuaicheng Liu. Learning optical flow with adaptive graph reasoning. In *Proceedings of the AAAI Conference on Artificial Intelligence (AAAI)*, 2022. [2](#), [6](#)
- [31] Nikolaus Mayer, Eddy Ilg, Philip Hausser, Philipp Fischer, Daniel Cremers, Alexey Dosovitskiy, and Thomas Brox. A large dataset to train convolutional networks for disparity, optical flow, and scene flow estimation. In *Proceedings of the IEEE conference on computer vision and pattern recognition*, pages 4040–4048, 2016. [2](#), [3](#)
- [32] Namuk Park and Songkuk Kim. How do vision transformers work? *arXiv preprint arXiv:2202.06709*, 2022. [5](#)
- [33] Adam Paszke, Sam Gross, Francisco Massa, Adam Lerer, James Bradbury, Gregory Chanan, Trevor Killeen, Zeming Lin, Natalia Gimelshein, Luca Antiga, et al. Pytorch: An imperative style, high-performance deep learning library. *Advances in neural information processing systems*, 32, 2019. [4](#)
- [34] Anurag Ranjan and Michael J Black. Optical flow estimation using a spatial pyramid network. In *Proceedings of the IEEE conference on computer vision and pattern recognition*, pages 4161–4170, 2017. [1](#)
- [35] Jerome Revaud, Philippe Weinzaepfel, César De Souza, Noe Pion, Gabriela Csurka, Johann Cabon, and Martin Humenberger. R2d2: repeatable and reliable detector and descriptor. *arXiv preprint arXiv:1906.06195*, 2019. [2](#)
- [36] Jerome Revaud, Philippe Weinzaepfel, Zaid Harchaoui, and Cordelia Schmid. Epicflow: Edge-preserving interpolation of correspondences for optical flow. In *Proceedings of the IEEE conference on computer vision and pattern recognition*, pages 1164–1172, 2015. [2](#)
- [37] Ignacio Rocco, Mircea Cimpoi, Relja Arandjelović, Akihiko Torii, Tomas Pajdla, and Josef Sivic. Neighbourhood consensus networks. *Advances in neural information processing systems*, 31, 2018. [3](#), [5](#)
- [38] Paul-Edouard Sarlin, Daniel DeTone, Tomasz Malisiewicz, and Andrew Rabinovich. Superglue: Learning feature matching with graph neural networks. In *Proceedings of the IEEE/CVF conference on computer vision and pattern recognition*, pages 4938–4947, 2020. [2](#)
- [39] Manolis Savva, Angel X Chang, and Pat Hanrahan. Semantically-enriched 3d models for common-sense knowledge. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition Workshops*, pages 24–31, 2015. [3](#)
- [40] Johannes L Schönberger, Enliang Zheng, Jan-Michael Frahm, and Marc Pollefeys. Pixelwise view selection for unstructured multi-view stereo. In *European Conference on Computer Vision*, pages 501–518. Springer, 2016. [1](#)
- [41] Leslie N Smith and Nicholay Topin. Super-convergence: Very fast training of neural networks using large learning rates. In *Artificial intelligence and machine learning for multi-domain operations applications*, volume 11006, pages 369–386. SPIE, 2019. [4](#)
- [42] Xiuchao Sui, Shaohua Li, Xue Geng, Yan Wu, Xinxing Xu, Yong Liu, Rick Goh, and Hongyuan Zhu. Craft: Cross-attentional flow transformer for robust optical flow. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 17602–17611, 2022. [1](#), [2](#), [3](#), [4](#), [6](#)
- [43] 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. [6](#)
- [44] 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 conference on computer vision and pattern recognition*, pages 8934–8943, 2018. [1](#), [2](#), [6](#)
- [45] Deqing Sun, Xiaodong Yang, Ming-Yu Liu, and Jan Kautz. Models matter, so does training: An empirical study of cnns for optical flow estimation. *IEEE transactions on pattern analysis and machine intelligence*, 42(6):1408–1423, 2019. [6](#)
- [46] Jiaming Sun, Zehong Shen, Yuang Wang, Hujun Bao, and Xiaowei Zhou. Loftr: Detector-free local feature matching with transformers. In *Proceedings of the IEEE/CVF conference on computer vision and pattern recognition*, pages 8922–8931, 2021. [2](#), [3](#), [4](#)
- [47] Shangkun Sun, Yuanqi Chen, Yu Zhu, Guodong Guo, and Ge Li. Skflow: Learning optical flow with super kernels. *arXiv preprint arXiv:2205.14623*, 2022. [2](#), [6](#)
- [48] Shuyang Sun, Zhanghui Kuang, Lu Sheng, Wanli Ouyang, and Wei Zhang. Optical flow guided feature: A fast and robust motion representation for video action recognition. In *Proceedings of the IEEE conference on computer vision and pattern recognition*, pages 1390–1399, 2018. [1](#)
- [49] Shitao Tang, Jiahui Zhang, Siyu Zhu, and Ping Tan. Quadtree attention for vision transformers. *arXiv preprint arXiv:2201.02767*, 2022. [2](#), [3](#)
- [50] Zachary Teed and Jia Deng. Raft: Recurrent all-pairs field transforms for optical flow. In *European conference on computer vision*, pages 402–419. Springer, 2020. [1](#), [2](#), [3](#), [4](#), [5](#), [6](#)
- [51] Prune Truong, Martin Danelljan, Luc V Gool, and Radu Timofte. Gocor: Bringing globally optimized correspon-

- dence volumes into your neural network. *Advances in Neural Information Processing Systems*, 33:14278–14290, 2020. [3](#)
- [52] Prune Truong, Martin Danelljan, and Radu Timofte. Glunet: Global-local universal network for dense flow and correspondences. In *Proceedings of the IEEE/CVF conference on computer vision and pattern recognition*, pages 6258–6268, 2020. [1](#), [3](#)
- [53] Michał Tyszkiewicz, Pascal Fua, and Eduard Trulls. Disk: Learning local features with policy gradient. *Advances in Neural Information Processing Systems*, 33:14254–14265, 2020. [3](#), [5](#)
- [54] Haofei Xu, Jing Zhang, Jianfei Cai, Hamid Rezaatfighi, and Dacheng Tao. Gmflow: Learning optical flow via global matching. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 8121–8130, 2022. [1](#), [3](#), [6](#)
- [55] Christopher Zach, Thomas Pock, and Horst Bischof. A duality based approach for realtime tv-l1 optical flow. In *Joint pattern recognition symposium*, pages 214–223. Springer, 2007. [2](#)
- [56] Feihu Zhang, Oliver J Woodford, Victor Adrian Prisacariu, and Philip HS Torr. Separable flow: Learning motion cost volumes for optical flow estimation. In *Proceedings of the IEEE/CVF International Conference on Computer Vision*, pages 10807–10817, 2021. [6](#)
- [57] Shiyu Zhao, Long Zhao, Zhixing Zhang, Enyu Zhou, and Dimitris Metaxas. Global matching with overlapping attention for optical flow estimation. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 17592–17601, 2022. [2](#), [3](#), [4](#), [5](#), [6](#)
- [58] Zihua Zheng, Ni Nie, Zhi Ling, Pengfei Xiong, Jiangyu Liu, Hao Wang, and Jiankun Li. Dip: Deep inverse patch-match for high-resolution optical flow. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 8925–8934, 2022. [2](#), [6](#)