

FlowFormer++: Masked Cost Volume Autoencoding for Pretraining Optical Flow Estimation

Xiaoyu Shi^{1,3*} Zhaoyang Huang^{1,3*} Dasong Li¹ Manyuan Zhang¹
Ka Chun Cheung³ Simon See³ Hongwei Qin⁴ Jifeng Dai⁵ Hongsheng Li^{1,2}

¹CUHK MMLab

²CPII under InnoHK

³NVIDIA AI Technology Center

⁴SenseTime Research

⁵Tsinghua University

{xiaoyushi@link, drinkingcoder@link, hsli@ee}.cuhk.edu.hk

Abstract

FlowFormer [24] introduces a transformer architecture into optical flow estimation and achieves state-of-the-art performance. The core component of FlowFormer is the transformer-based cost-volume encoder. Inspired by the recent success of masked autoencoding (MAE) pretraining in unleashing transformers' capacity of encoding visual representation, we propose Masked Cost Volume Autoencoding (MCVA) to enhance FlowFormer by pretraining the cost-volume encoder with a novel MAE scheme. Firstly, we introduce a block-sharing masking strategy to prevent masked information leakage, as the cost maps of neighboring source pixels are highly correlated. Secondly, we propose a novel pre-text reconstruction task, which encourages the cost-volume encoder to aggregate long-range information and ensures pretraining-finetuning consistency. We also show how to modify the FlowFormer architecture to accommodate masks during pretraining. Pretrained with MCVA, FlowFormer++ ranks 1st among published methods on both Sintel and KITTI-2015 benchmarks. Specifically, FlowFormer++ achieves 1.07 and 1.94 average endpoint error (AEPE) on the clean and final pass of Sintel benchmark, leading to 7.76% and 7.18% error reductions from FlowFormer. FlowFormer++ obtains 4.52 F1-all on the KITTI-2015 test set, improving FlowFormer by 0.16.

1. Introduction

Optical flow is a long-standing vision task, targeting at estimating per-pixel displacement between consecutive video frames. It can provide motion and correspondence information in many downstream video problems, including video object detection [65, 81, 82], action recogni-

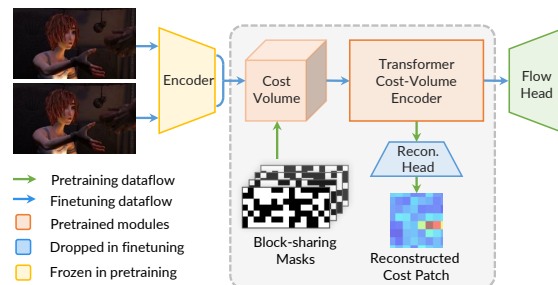


Figure 1. **Overview of FlowFormer++.** The core component of FlowFormer architecture is the transformer-based cost-volume encoder. We propose Masked Cost Volume Autoencoding to pretrain the cost-volume encoder. During pretraining, a portion of cost values are masked and the cost-volume encoder is required to reconstruct masked cost patches.

tion [48, 61, 80], and video restoration [5, 14, 32, 34, 51, 70].

Recently, FlowFormer [24] introduces a transformer architecture for optical flow estimation and achieves state-of-the-art performance. The core of its success lies on two aspects: the ImageNet-pretrained transformer-based image encoder and the transformer-based cost-volume encoder. Notably, adopting an ImageNet-pretrained visual backbone leads to considerable performance gain over the train-from-scratch counterpart, indicating that random weight initialization hinders the learning of correspondence estimation. This naturally begs the question: can we also pretrain the transformer-based cost-volume encoder and thus further unleash its power to achieve more accurate optical flow?

In this paper, we propose masked cost-volume autoencoding (MCVA), a self-supervised pretraining scheme to enhance the cost-volume encoding on top of the FlowFormer framework. We are inspired by the recent success of masked autoencoding, such as BERT [11] in NLP and MAE [20] in computer vision. The key idea of masked autoencoding is masking a portion of input data, and requiring networks to learn high-level representation for masked

*Xiaoyu Shi and Zhaoyang Huang assert equal contributions.

contents reconstruction. However, it is non-trivial to adapt the masked autoencoding strategy to learn a better cost volume encoder for optical flow estimation, because of the two following reasons. Firstly, the cost volume might contain redundancy and the cost maps (cost values between a source-image pixel to all target-image pixels) of neighboring source-image pixels are highly correlated. Randomly masking cost values, as done in other single-image pre-training methods [20], leads to information leakage and makes the model biased towards aggregating local information. Secondly, existing masked autoencoding methods target at reconstructing masked content randomly selected from fixed locations. This suffices to pretrain general-purpose single-image encoder in other fields. However, the cost-volume encoder of FlowFormer is deeply coupled with the follow-up recurrent decoder, which demands cost information of long range at flexible locations.

To tackle the aforementioned issues, we introduce two task-specific designs. Firstly, instead of randomly masking the cost volume, we partition source pixels into large varied-size blocks and let source pixels within the same block share a common mask pattern on their cost maps. This strategy, termed block-sharing masking, prevents the cost-volume encoder from reconstructing masked cost values by simply copying from neighboring source pixels' cost maps. Such design enforces the cost-volume encoder to abstract useful cues from cost maps belonging to far-away source pixels, which encourages long-range information aggregation. Secondly, to mimic the decoding process in finetuning and thus avoid pretraining-finetuning discrepancy, we propose a novel pre-text reconstruction task as shown in Fig. 3: small cost patches are randomly cropped from the cost maps to retrieve features from the cost-volume encoder, aiming to reconstruct larger cost map patches centered at the same locations. This is in line with the decoding process of FlowFormer in the finetuning stage. This pre-text task explicitly encourages the cost-volume encoder to capture long-range information for cost-volume encoding, which is critical for optical flow estimation.

In essence, the proposed masked cost-volume autoencoding (MCVA) has unique designs compared with conventional MAE methods, which encourages the cost-volume encoder 1) to construct high-level holistic representation of the cost volume, more effectively encoding long-range information, 2) to reason about occluded (*i.e.*, masked) information by aggregating faithful unmasked costs, and 3) to decode task-specific feature (*i.e.*, larger cost patches at required locations) to better align the pretraining process with that of the finetuning. These designs contribute to better handling of hard cases, such as noises, large-displacement motion and occlusion, for more accurate flow estimation.

To conclude, the contributions of this work are three-fold: 1) We propose the masked cost-volume autoencod-

ing scheme to better pretrain the cost-volume encoder of FlowFormer. 2) We propose task-specific masking strategy and reconstruction pre-text task to mitigate pretraining-finetuning discrepancy, fully taking advantage of the learned representations from pretraining. 3) With the proposed pretraining technique, our proposed FlowFormer++ obtains all-sided improvements over FlowFormer, setting new state-of-the-art performance on public benchmarks.

2. Related Work

Optical Flow. Image correspondence is a critical intermediate representation for downstream applications [6, 7, 23, 25, 35–37, 46, 53, 76]. Optical flow is defined as dense image correspondence between consecutive frames in a video. Compared with traditional optimization-based optical flow methods [2, 3, 22, 56] empirically formulating flow estimation, data-driven methods [13, 28] directly learn to estimate optical flow from labeled data. Since FlowNet [13, 28], learning optical flow with neural networks presents superior performance and is still fast progressing where network architecture design becomes the key to improving optical flow accuracy. A series of excellent works [1, 21, 30, 31, 49, 58, 59, 62, 67, 71, 75] are devoted to designing better network modules, which, indeed, introduced better inductive bias to the optical flow formulation. For example, encoding image feature with CNNs brings locality prior, and the all-pairs 4D cost volume [30, 62] outperforms the coarse-to-fine cost volumes [49, 58, 59] in modeling small fast-motion objects. Such a design also brings benefits to event-based optical flow [16, 38] and multi-frame optical flow [52]. However, the empirical network design may always ignore some unintended cases. Due to the success of transformers [10, 11, 39, 64] in image recognition [9, 12, 43], the optical flow community also tries transformers [24, 46, 55, 68] to further weaken the network-determined bias and learn feature relationships from data. By replacing the handcrafted modules, *i.e.*, the CNN image encoder, the cost pyramid, and the indexing-based costs retrieval, in RAFT [62] with transformers, FlowFormer [24] achieves state-of-the-art accuracy. However, transformers are known for requiring tremendous training data to capture feature relationships [12, 20] while collecting ground-truth flows for supervised optical flow learning is expensive. Inspired by the emerging pretraining-finetuning paradigm for vision transformers [15, 20], we explore to pretrain FlowFormer to capture the feature relationship for optical flow.

Masked Autoencoding (MAE). As a self-supervised learning technique, MAE, *e.g.*, BERT [11], achieved great success in NLP. Based on transformers, they mask a portion of the input tokens and require the models to predict the missing content from the reserved tokens. Pretraining with MAE encourages transformers to build effective long-range feature relationships. Recently, transform-

ers also stream into the computer vision area, such as image recognition [9, 12, 43], video inpainting [40, 41, 74], point cloud recognition [19, 77]. By breaking the limitations that convolution can only model local features, transformers present a significant performance gap compared to the previous counterparts. Pretraining with MAE is also introduced to these modalities, *e.g.*, image [8, 15, 20, 66], video [63], point cloud [47, 73]. These works show that MAE effectively releases the transformer power and do not require extra labeled data. FlowFormer [24] presents a transformer-based cost volume encoder and achieves state-of-the-art accuracy. In this paper, we propose the masked cost-volume autoencoding to pretrain the cost volume encoder on a video dataset, which further unleashes the power of the transformer-based cost-volume encoder.

3. Method

As presented in Fig. 2, we propose a masked cost-volume autoencoding (MCVA) scheme to pretrain the cost-volume encoder of FlowFormer framework for better performance. The key of general masked autoencoding methods is to mask a portion of data and encourage the network to reconstruct the masked tokens from visible ones. Due to the redundant nature of the cost volume and the original FlowFormer architecture being incompatible with masks, naively adopting this paradigm to pretrain the cost-volume encoder leads to inferior performance. Our proposed MCVA tackles the challenge and conducts masked autoencoding with three key components: a proper masking strategy on the cost volume, modifying FlowFormer architecture to accommodate masks, and a novel pre-text reconstruction task supervising the pretraining process.

In this section, we first revisit the FlowFormer architecture, and then elaborate the proposed three key designs. We first introduce the masking strategy, dubbed as block-sharing masking, and then show the masked cost-volume tokenization that makes the cost-volume encoder compatible with masks. Coupling these two designs prevents the masked autoencoding from being hindered by information leakage in pretraining. Finally, we present the pre-text cost reconstruction task, mimicking the decoding process in fine-tuning to pretrain the cost-volume encoder.

3.1. A Revisit of FlowFormer

Given a pair of source and target images, optical flow aims at recovering pixel-level correspondences for all source pixels. FlowFormer encodes the pair of images’ features with an ImageNet-pretrained Twins-SVT [9] as $\mathbb{R}^{H_I \times W_I \times 3} \rightarrow \mathbb{R}^{H \times W \times D}$, and creates a 4D cost volume of size $H \times W \times H \times W$ by computing all-pairs feature correlations. H_I, W_I and H, W respectively indicate the height and width of the images and the visual feature maps. The cost volume can also be viewed as a series of cost maps of

size $\mathbb{R}^{H \times W}$, each of which measures the similarity between one source pixel and all target pixels.

The 4D cost volume contains abundant but redundant information for optical flow estimation. FlowFormer projects it into a latent space of size $\mathbb{R}^{H \times W \times K \times D}$ with a cost tokenizer. In the latent space, each source pixel’s cost map is transformed into cost memory consisting of K tokens of dimension D , which is a more compact representation and is further processed by a transformer-based cost encoder, dubbed as alternate-group transformer (AGT). Finally, FlowFormer recurrently decodes the flow estimation from the cost memory with cross-attention.

FlowFormer is the first transformer architecture specifically designed for optical flow estimation, which enjoys the benefits of long-range information encoding via self-attention, but also encounters the similar problem to general vision transformers: it needs large-scale training data to model unbiased representations. The FlowFormer with the ImageNet-pretrained Twins-SVT backbone leads to boosted accuracy, while the same model with a train-from-scratch Twins-SVT or a shallow CNN achieve similar degraded performances, demonstrating the necessity of pretraining transformers for optical flow estimation. However, the ImageNet can only be used for pretraining the single-image encoder and the cost-volume encoder in FlowFormer is still trained from scratch and might not converge to the optimal point. To enable the pretraining of the cost-volume encoder to further enhance optical flow estimation, we propose the masked cost-volume autoencoding scheme.

3.2. Block-sharing Cost Volume Masking

A properly designed masking scheme is required to conduct autoencoding of the masked cost volume. For each source pixel \mathbf{x} , we need to create a binary mask $\mathbf{M}_{\mathbf{x}} \in \{0, 1\}^{H \times W}$ to its cost map $\mathbf{C}_{\mathbf{x}} \in \mathbb{R}^{H \times W}$, where 0 indicates masking (*i.e.*, removing) cost values from the masked locations. Naturally, neighboring source pixels’ cost maps are highly correlated. Randomly masking neighboring pixels’ cost maps might cause information leakage, *i.e.*, masked cost values might be easily reconstructed by copying the cost values from neighboring source pixels’ cost maps.

To prevent such an over-simplified learning process, we propose a block-sharing masking strategy. We partition source pixels into non-overlapping blocks in each iteration. All source pixels belonging to the same block share a common mask for masked region reconstruction. In this way, neighboring source pixels are unlikely to copy each other’s cost maps to over-simplify the autoencoding process. Besides, the size of block is designed to be large (height and width of blocks are of $32 \sim 120$ pixels) and randomly changes in each iteration, and thus encouraging the cost-volume encoder to aggregate information from long-range context and to filter noises of cost values. The details of the

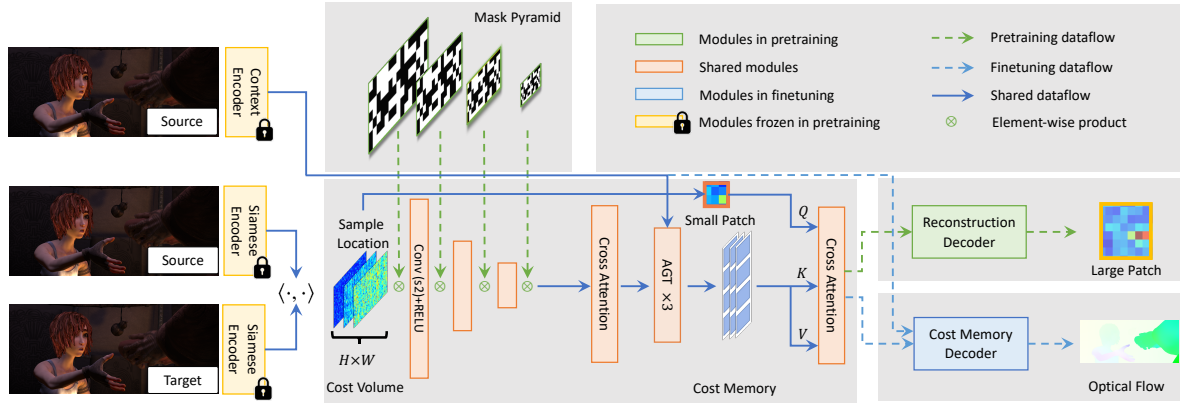


Figure 2. **Architecture of FlowFormer++.** During pretraining, FlowFormer++ freezes the image and context encoders, block-wisely masks the cost volume, and learns to reconstruct larger cost patches from small cost patches to pretrain cost-volume encoder. In fine-tuning, FlowFormer++ uses the full cost volume, removes the reconstruction decoder, and adds the cost memory decoder to learn optical flow, which naturally falls back to the FlowFormer architecture but inherits the pretrained parameters in the cost-volume encoder.

mask generation algorithm are provided in supplementary.

Specifically, for each source pixel’s cost map, we first generate the mask map $\mathbf{M}_x^3 \in \{0, 1\}^{\frac{H}{8} \times \frac{W}{8}}$ at $\frac{1}{8}$ resolution, and then up-sample it $2 \times$ for three times to obtain a pyramid of mask maps $\mathbf{M}_x^i \in \{0, 1\}^{\frac{H}{2^i} \times \frac{W}{2^i}}$, where $i \in \{0, 1, 2\}$, which are used for the down-sampling encoding process and will be discussed later in Sec. 3.3.

Another key design is that, in pretraining, we **freeze** the ImageNet-pretrained Twins-SVT backbone to build the cost volume from the pair of input images. Freezing the image encoder ensures the reconstruction targets (*i.e.*, raw cost values) to maintain static and avoids training collapse.

3.3. Masked Cost-volume Tokenization

Given the above generated mask for each source pixel \mathbf{x} ’s cost map, FlowFormer adopts a two-step cost-volume tokenization before the cost encoder. To prevent the masked costs from leaking into subsequent cost aggregation layers, the intermediate embeddings of the cost map need to be properly masked in the cost-volume tokenization process. We propose the masked cost-volume tokenization, which prevents mixing up masked and visible features. Firstly, FlowFormer patchifies the raw cost map $\mathbf{C}_x \in \mathbb{R}^{H \times W}$ of each source pixel \mathbf{x} (which is obtained by computing dot-product similarities between the source pixel \mathbf{x} and all target pixels) via 3 stacked stride-2 convolutions. We denote the feature maps after each of the 3 convolutions as \mathbf{F}_x^i , which have spatial sizes of $\frac{H}{2^i} \times \frac{W}{2^i}$ for $i \in \{0, 1, 2\}$. We propose to replace the vanilla convolutions used in the FlowFormer with masked convolutions [15, 18, 50]:

$$\mathbf{F}_x^{i+1} = \text{Conv}_{\text{stride}2}(\text{ReLU}(\mathbf{F}_x^i \odot \mathbf{M}_x^i)), \quad (1)$$

where \odot indicates element-wise multiplication, $i \in \{0, 1, 2\}$, and \mathbf{F}_x^0 is the raw cost map \mathbf{C}_x . The masked

convolutions with the three binary mask maps remove all cost features in the masked regions in pretraining. Secondly, FlowFormer further projects the patchified cost-map features into the latent space via cross-attention. We thus remove the tokens in \mathbf{F}_x^3 indicated by the mask map \mathbf{M}_x^3 and then only project the remaining tokens into the latent space via the same cross-attention. During finetuning, the mask maps are removed to utilize all cost features, which converts the masked convolution to the vanilla convolution but the pretrained parameters in the convolution kernels and cross-attention are maintained.

The masked cost-volume tokenization completes two tasks. Firstly, it ensures the subsequent cost-volume encoder only processes visible features in pretraining. Secondly, the network structure is consistent with the standard tokenization of FlowFormer and can directly be used for finetuning so that the pretrained parameters have the same semantic meanings. After the masked cost-volume tokenization, the cost aggregation layers (*i.e.*, AGT layers) take visible features as input which also don’t need to be modified in finetuning. The latent features interact with those of other source pixels in AGT layers and are transformed to the cost memory \mathbf{T}_x . We explain how to decode the cost memory to estimate flows in following section.

3.4. Reconstruction Target for Cost Memory Decoding

With the masked cost-volume tokenization, the cost encoder encodes the unmasked cost volume into the cost memory. The next step is decoding and reconstructing the masked regions from the cost memory.

In this section, we formulate the pre-text reconstruction targets, which supervises the decoding process as well as aforementioned embedding and aggregation layers. We start by revisiting the dynamic positional decoding scheme of

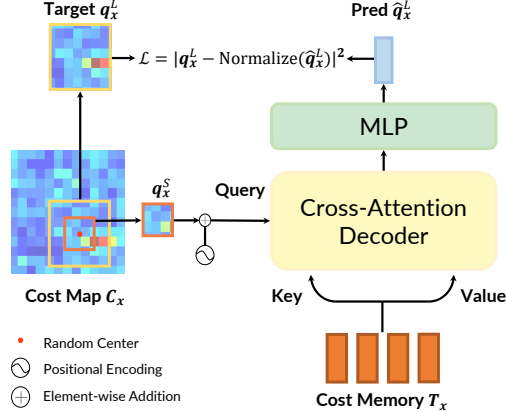


Figure 3. **Pre-text reconstruction for cost memory decoding.** For each source pixel \mathbf{x} and its corresponding cost map \mathbf{C}_x , a small cost patch \mathbf{q}_x^S is randomly cropped from the cost map to retrieve features from the cost memory \mathbf{T}_x , aiming to reconstruct larger cost patch \mathbf{q}_x^L centered at the same location. The superscript “S” and “L” of \mathbf{q}_x denote “small” and “large” respectively.

FlowFormer, and present our reconstruction targets which are highly consistent with the finetuning tasks. FlowFormer adopts recurrent flow prediction. In each iteration of the recurrent process, the flow of source pixel \mathbf{x} is decoded from cost memory \mathbf{T}_x , conditioned on current predicted flow, to update the flow prediction. Specifically, current predicted corresponding location in the target image \mathbf{p}_x is computed as $\mathbf{p}_x = \mathbf{x} + \mathbf{f}(\mathbf{x})$, where $\mathbf{f}(\mathbf{x})$ is current predicted flow. A local cost patch \mathbf{q}_x is then cropped from the 9×9 window centered at \mathbf{p}_x on the raw cost map \mathbf{C}_x . FlowFormer utilizes this local cost patch (with positional encoding) as the query feature to retrieve aggregated cost feature \mathbf{c}_x via cross-attention operation:

$$\begin{aligned} \mathbf{Q}_x &= \text{FFN}(\text{FFN}(\mathbf{q}_x) + \text{PE}(\mathbf{p})), \\ \mathbf{K}_x &= \text{FFN}(\mathbf{T}_x), \quad \mathbf{V}_x = \text{FFN}(\mathbf{T}_x), \\ \mathbf{c}_x &= \text{Attention}(\mathbf{Q}_x, \mathbf{K}_x, \mathbf{V}_x). \end{aligned} \quad (2)$$

Pre-text Reconstruction. Intuitively, \mathbf{c}_x should contain long-range cost information for better optical flow estimation and its center of the perceptual field should be similar to that of the local cost patch \mathbf{q}_x . We design a pre-text reconstruction task in line with these two characteristics to pretrain the cost-volume encoder as shown in Fig. 3: small cost-map patches are randomly cropped from the cost maps to retrieve cost features from the cost memory, targeting at reconstructing larger cost-map patches centered at the same locations.

Specifically, for each source pixel \mathbf{x} , we randomly sample a location \mathbf{o}_x , which is analogous to \mathbf{p}_x in finetuning. Taking this location as center, we crop a small cost-map patch $\mathbf{q}_x^S = \text{Crop}_{9 \times 9}(\mathbf{C}_x, \mathbf{o}_x)$ of shape 9×9 . Then we perform the decoding process shown in Equation 2 to obtain the cost feature \mathbf{c}_x , except that \mathbf{p}_x and \mathbf{q}_x are replaced by \mathbf{o}_x and \mathbf{q}_x^S in pretraining, respectively. To encourage the

extracted cost feature \mathbf{c}_x to carry long-range cost information conditioned on \mathbf{o}_x , we take larger cost-map patch as supervision. Specifically, we crop another larger cost-map patch $\mathbf{q}_x^L = \text{Crop}_{15 \times 15}(\mathbf{C}_x, \mathbf{o}_x)$ of shape 15×15 centered at the same location \mathbf{o}_x . We choose a light-weight MLP as prediction head. The MLP takes as input \mathbf{c}_x and its output $\hat{\mathbf{q}}_x^L = \text{MLP}(\mathbf{c}_x)$ is supervised by normalized \mathbf{q}_x^L . We take mean squared error (MSE) as loss function.

$$\mathcal{L} = \frac{1}{|\Omega|} \sum_{\mathbf{x} \in \Omega} \left| \mathbf{q}_x^L - \text{Normalize}(\hat{\mathbf{q}}_x^L) \right|^2, \quad (3)$$

where Ω is the set of source pixels.

Discussion. The key of pretraining is to maintain consistent with finetuning, in terms of both network architecture and prediction target. To this end, we keep the cross-attention decoding layer unchanged and construct inputs with the same semantic meaning (*e.g.*, replacing dynamically predicted \mathbf{p}_x with randomly sampled \mathbf{o}_x); we supervise the extracted feature \mathbf{c}_x with long-range cost values to encourage the cost-volume encoder to aggregate global information for better optical flow estimation. What’s more, our scheme only takes an extra light-weight MLP as prediction head, which is unused in finetuning. Compared with previous methods that use a stack of self-attention layers, it is much more computationally efficient.

4. Experiments

We evaluate our FlowFormer++ on the Sintel [4] and KITTI-2015 [17] benchmarks. We pretrain FlowFormer++ using the proposed Masked Cost-volume Autoencoding on YouTube-VOS [69] dataset. For the supervised finetuning, following previous works, we train FlowFormer++ on FlyingChairs [13] and FlyingThings [45], and then respectively finetune it on the Sintel and KITTI-2015 benchmarks. FlowFormer++ obtains all-sided improvements over FlowFormer, ranking 1st on both benchmarks.

Experimental Setup. We adopt the commonly-used average end-point-error (AEPE) as the evaluation metric. It measures the average l_2 distance between predictions and ground truth. For the KITTI-2015 dataset, we additionally use the F1-all (%) metric, which refers to the percentage of pixels whose flow error is larger than 3 pixels or over 5% of the length of ground truth flows. YouTube-VOS is a large-scale dataset containing video clips from YouTube website. The Sintel dataset is rendered from the same movie in two passes: the clean pass is rendered with easier smooth shading and specular reflections, while the final pass includes motion blur, camera depth-of-field blur and atmospheric effects. The motions in the Sintel dataset are relatively large and complicated. The KITTI-2015 dataset constitutes of real-world driving scenarios with sparse ground truth.

Training Data	Method	Sintel (train)		KITTI-15 (train)		Sintel (test)		KITTI-15 (test)
		Clean	Final	F1-epe	F1-all	Clean	Final	F1-all
A	Perceiver IO [29]	1.81	2.42	4.98	-	-	-	-
	PWC-Net [58]	2.17	2.91	5.76	-	-	-	-
	RAFT [62]	1.95	2.57	4.23	-	-	-	-
C+T	HD3 [72]	3.84	8.77	13.17	24.0	-	-	-
	LiteFlowNet [26]	2.48	4.04	10.39	28.5	-	-	-
	PWC-Net [58]	2.55	3.93	10.35	33.7	-	-	-
	LiteFlowNet2 [27]	2.24	3.78	8.97	25.9	-	-	-
	S-Flow [75]	1.30	2.59	4.60	15.9	-	-	-
	RAFT [62]	1.43	2.71	5.04	17.4	-	-	-
	FM-RAFT [31]	1.29	2.95	6.80	19.3	-	-	-
	GMA [30]	1.30	2.74	4.69	17.1	-	-	-
	GMFlow [68]	1.08	2.48	-	-	-	-	-
	GMFlowNet [79]	1.14	2.71	4.24	15.4	-	-	-
	CRAFT [55]	1.27	2.79	4.88	17.5	-	-	-
	SKFlow [60]	1.22	2.46	4.47	15.5	-	-	-
	FlowFormer [24]	0.94	2.33	4.09 [†]	14.72 [†]	-	-	-
	Ours	0.90	2.30	3.93[†]	14.13[†]	-	-	-
C+T+S+K+H	LiteFlowNet2 [27]	(1.30)	(1.62)	(1.47)	(4.8)	3.48	4.69	7.74
	PWC-Net+ [59]	(1.71)	(2.34)	(1.50)	(5.3)	3.45	4.60	7.72
	VCN [71]	(1.66)	(2.24)	(1.16)	(4.1)	2.81	4.40	6.30
	MaskFlowNet [78]	-	-	-	-	2.52	4.17	6.10
	S-Flow [75]	(0.69)	(1.10)	(0.69)	(1.60)	1.50	2.67	4.64
	RAFT [62]	(0.76)	(1.22)	(0.63)	(1.5)	1.94	3.18	5.10
	RAFT* [62]	(0.77)	(1.27)	-	-	1.61	2.86	5.10
	FM-RAFT [31]	(0.79)	(1.70)	(0.75)	(2.1)	1.72	3.60	6.17
	GMA [30]	-	-	-	-	1.40	2.88	5.15
	GMA* [30]	(0.62)	(1.06)	(0.57)	(1.2)	1.39	2.47	5.15
	GMFlow [68]	-	-	-	-	1.74	2.90	9.32
	GMFlowNet [79]	(0.59)	(0.91)	(0.64)	(1.51)	1.39	2.65	4.79
	CRAFT [55]	(0.60)	(1.06)	(0.57)	(1.20)	1.45	2.42	4.79
	SKFlow* [60]	(0.52)	(0.78)	(0.51)	(0.94)	1.28	2.23	4.84
	FlowFormer [24]	(0.48)	(0.74)	(0.53)	(1.11)	1.16	2.09	4.68 [†]
Ours	(0.40)	(0.60)	(0.57)	(1.16)	1.07	1.94	4.52[†]	

Table 1. **Experiments on Sintel [4] and KITTI [17] datasets.** ‘A’ denotes the autoflow dataset. ‘C + T’ denotes training only on the FlyingChairs and FlyingThings datasets. ‘+ S + K + H’ denotes finetuning on the combination of Sintel, KITTI, and HD1K training sets. * denotes that the methods use the warm-start strategy [62], which relies on previous image frames in a video, while other methods use two frames only. [†] denotes the result is obtained via the tile technique proposed in FlowFormer [24]. Our FlowFormer++ achieves the best generalization performance (C+T) and ranks 1st on both the Sintel and the KITTI-15 benchmarks (C+T+S+K+H).

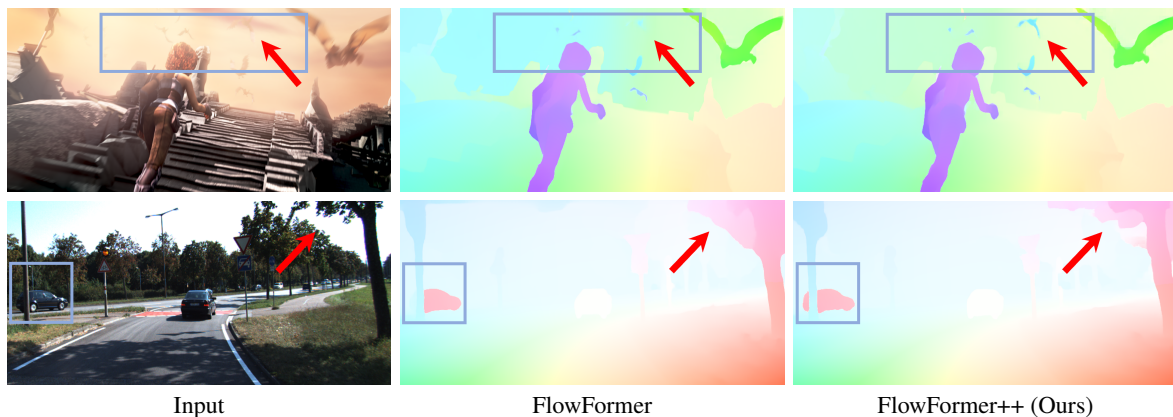


Figure 4. **Qualitative comparison on Sintel and KITTI test sets.** FlowFormer++ preserves clearer details (pointed by red arrows) and maintains better global consistency (indicated by blue boxes).

Implementation Details. We use the same architecture of FlowFormer for fair comparison. The image feature encoder and context feature encoder are chosen as the first two stages of ImageNet-pretrained Twins-SVT, which are frozen in pretraining for better performance. We pretrain our model on YouTube-VOS for 50k iterations with a batch size of 24. The highest learning rate is set as 5×10^{-4} . During finetuning, we follow the same training procedure of FlowFormer. We train our model on FlyingChairs for 120k iterations with a batch size of 8 and on FlyingThings with a batch size of 6 (denoted as ‘C+T’). Then, we train FlowFormer++ by combining data from Sintel, KITTI-2015 and HD1K [33] (denoted as ‘C+T+S+K+H’) for another 120k iterations with a batch size of 6. This model is submitted to Sintel online test benchmark for evaluation. To obtain the best performance on the KITTI-2015 benchmark, we further train FlowFormer++ on the KITTI-2015 dataset for 50k iterations with a batch size of 6. The highest learning rate is set as 2.5×10^{-4} for FlyingChairs and 1.25×10^{-4} on other training sets. In both pretraining and finetuning, we choose AdamW optimizer and one-cycle learning rate scheduler. We crop images and tile predictions from all patches to obtain full-resolution flow predictions following Perceiver IO [29] and FlowFormer.

4.1. Quantitative Experiments

As shown in Table 1, we evaluate FlowFormer++ on the Sintel and KITTI-2015 benchmarks. Following previous methods, we evaluate the generalization performance of models on the training sets of Sintel and KITTI-2015 (denoted as ‘C+T’). We also compare the dataset-specific accuracy of optical flow models after dataset-specific finetuning (denoted as ‘C+T+S+K+H’). Autoflow [57] is a synthetic dataset of complicated visual disturbance.

Generalization Performance. The ‘C+T’ setting in Table 1 reflects the generalization capacity of models. FlowFormer++ ranks 1st on both benchmarks among published methods. It achieves 0.90 and 2.30 on the clean and final pass of Sintel. Compared with FlowFormer, it achieves 4.26% error reduction on Sintel clean pass. On the KITTI-2015 training set, FlowFormer++ achieves 3.93 F1-epe and 14.13 F1-all, improving FlowFormer by 0.16 and 0.59, respectively. These results show that our proposed MCVA promotes the generalization capacity of FlowFormer.

Dataset-specific Performance. After training the FlowFormer++ in the ‘C+T+S+K+H’ setting, we evaluate its performance on the Sintel online benchmark. It achieves 1.07 and 2.09 on the clean and final passes, a 7.76% and 7.18% error reduction from previous best model FlowFormer.

We further finetune FlowFormer++ on the KITTI-2015 training set after the Sintel stage and evaluate its performance on the KITTI online benchmark. FlowFormer++ achieves 4.52 F1-all, improving FlowFormer by 0.16 while

also outperforming the previous best model S-Flow by 0.12.

To conclude, FlowFormer++ shows greater optical flow estimation capacity for both naturalistic non-rigid motions (Sintel) and real-world rigid scenarios (KITTI-2015). This validates that our proposed MCVA improves the FlowFormer architecture by enhancing the cost-volume encoder.

4.2. Qualitative Experiments

We visualize flow predictions by our FlowFormer++ and FlowFormer on Sintel and KITTI test sets in Fig. 4 to qualitatively show how FlowFormer++ outperforms FlowFormer. The red arrows highlight that FlowFormer++ preserves clearer details than FlowFormer: in the first row, FlowFormer misses the flying bird while FlowFormer++ produces clear results; in the second row, FlowFormer++ keeps the boundaries of leaves while FlowFormer generates blurry prediction. FlowFormer++ also shows greater global aggregation capacity indicated by blue boxes. In the first row, FlowFormer produces obviously inconsistent prediction on the large-area sky, while FlowFormer++ yields consistent prediction. In the second row, the black car is partially occluded by the foreground tree, which challenges the optical flow model to aggregate information in long range. FlowFormer++ generates consistent prediction for the two separated parts of the car, while FlowFormer mixes the left part of the car with background and thus produces inconsistent optical flow prediction.

4.3. Ablation Study

We conduct a set of ablation studies to show the impact of designs in the Masked Cost-volume Autoencoding (MCVA). All models in the experiments are first pretrained and then finetuned on ‘C+T’. We report the test results on Sintel and KITTI training sets.

Masking Strategy. Masking strategy is one important design of our MCVA. As shown in Table 2, pretraining FlowFormer with random masking already improves the performance on three of the four metrics. But the proposed block-sharing masking strategy brings even larger gain, which demonstrates the effectiveness of this design. Besides, we observe higher pretraining loss with block-sharing masking than that with random masking, validating that the block-sharing masking makes the pretraining task harder.

Masking Ratio. Masking ratio influences the difficulty of the pre-text reconstruction task. We empirically find that the mask ratio of 50% yields the best overall performance.

Pre-text Reconstruction Design. The conventional MAE methods aim to reconstruct input data at fixed locations and use the positional encodings as query features to absorb information for reconstruction (the first row of Table 4). To keep consistent with the dynamic positional query of FlowFormer architecture, we propose to reconstruct contents at random locations (the second row of Table 4) and addition-

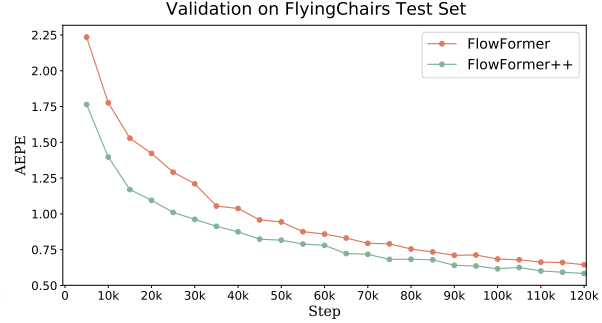
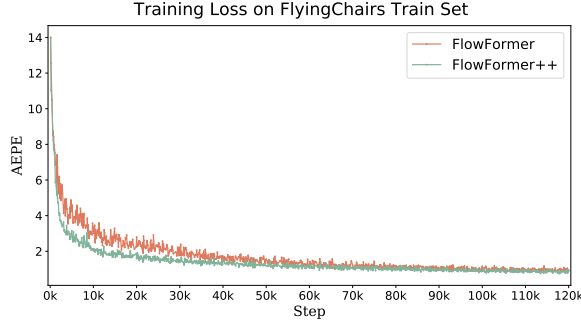


Figure 5. **Comparison on FlyingChairs.** FlowFormer++ converges faster and achieves lower validation error.

Case	Sintel (train)		KITTI-15 (train)	
	Clean	Final	F1-epe	F1-all
FlowFormer	0.94	2.33	4.09	14.72
+ Random Masking	0.93	2.35	4.03	14.30
+ Block-sharing Masking	0.90	2.30	3.93	14.13

Table 2. **Masking strategy.** Pretraining with block-sharing masking brings greater performance gain than random masking.

Masking Ratio	Sintel (train)		KITTI-15 (train)	
	Clean	Final	F1-epe	F1-all
20%	0.97	2.41	3.91	14.27
50%	0.90	2.30	3.93	14.13
80%	0.94	2.35	4.00	14.07

Table 3. **Masking ratio.** The masking ratio influences the difficulty of the reconstruction task. Masking 50% cost values yields best overall results.

Location	Query Feature	Sintel (train)		KITTI-15 (train)	
		Clean	Final	F1-epe	F1-all
Fixed	PE	0.99	2.40	4.35	15.33
Random	PE	0.95	2.42	3.99	14.47
Random	PE + Cropped patch	0.90	2.30	3.93	14.13

Table 4. **Pre-text reconstruction design.** Our pre-text reconstruction task leads to better performance over conventional MAE task (first row) and its improved version with random reconstruction locations (second row).

ally use local patches as query features (the third row of Table 4). The results validate the necessity of ensuring semantic consistency between pretraining and finetuning.

Freezing Image and Context Encoders. The FlowFormer architecture has an image encoder to encode visual appearance features for constructing the cost volume, and a context encoder to encode context features for flow prediction. As shown in Table 5, freezing the image encoder is necessary, otherwise the model diverges. We hypothesize that the frozen image encoder ensures the reconstruction targets (*i.e.*, raw cost values) to keep static. Freezing the context encoder leads to better overall performance.

Comparisons with Unsupervised Methods. We also use conventional unsupervised methods to pretrain FlowFormer with photometric loss and smooth loss following [42, 44, 54] and then finetune it in the ‘C+T’ setting as FlowFormer++. As shown in Table 6, our MCVA outperforms the unsupervised counterpart for pretraining FlowFormer.

FlowFormer++ v.s. FlowFormer on FlyingChairs. We show the training and validating loss of the training process on FlyingChairs [13] in Fig. 5. FlowFormer++ presents faster convergence during training and better validation loss

Image Encoder	Context Encoder	Sintel (train)		KITTI-15 (train)	
		Clean	Final	F1-epe	F1-all
✗	✗	-	-	-	-
✗	✓	-	-	-	-
✓	✗	0.92	2.34	3.90	14.23
✓	✓	0.90	2.30	3.93	14.13

Table 5. **Freezing image and context encoders in pretraining.** Freezing the image encoder ensures the reconstruction targets (*i.e.*, raw cost values) to maintain static, otherwise the model diverges. Freezing the context encoder leads to better overall performance.

Methods	Sintel (train)		KITTI-15 (train)	
	Clean	Final	F1-epe	F1-all
Unsupervised Baseline	0.99	2.54	4.38	15.22
MCVA (ours)	0.90	2.30	3.93	14.13

Table 6. **Comparisons with unsupervised methods.** We use the conventional unsupervised algorithm [42, 44, 54] (using photometric loss and smooth loss) to pretrain FlowFormer for comparison. Our MCVA outperforms the unsupervised counterpart.

at the end, which reveals that FlowFormer++ learns effective feature relationships during pretraining and benefits the supervised finetuning.

5. Conclusion and Limitation

In this paper, we propose Masked Cost Volume Autoencoding (MCVA) to enhance the cost-volume encoder of FlowFormer by pretraining. We show that the naive adaptation of MAE scheme to cost volume does not work due to the redundant nature of cost volumes and the incurred pretraining-finetuning discrepancy. We tackle these issues with a specially designed block-sharing masking strategy and the novel pre-text reconstruction task. These designs ensure semantic integrity between pretraining and finetuning and encourage the cost-volume to aggregate information in a long range. Experiments demonstrate clear generalization and dataset-specific performance improvements. The major limitation is that our proposed pretraining strategy is mainly designed for FlowFormer.

Acknowledgements. This project is funded in part by National Key R&D Program of China Project 2022ZD0161100, by the Centre for Perceptual and Interactive Intelligence (CPII) Ltd under the Innovation and Technology Commission (ITC)’s InnoHK, by General Research Fund of Hong Kong RGC Project 14204021. Hongsheng Li is a PI of CPII under the InnoHK.

References

- [1] Shaojie Bai, Zhengyang Geng, Yash Savani, and J Zico Kolter. Deep equilibrium optical flow estimation. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 620–630, 2022. 2
- [2] 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
- [3] 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
- [4] 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. 5, 6
- [5] Kelvin CK Chan, Xintao Wang, Ke Yu, Chao Dong, and Chen Change Loy. Basicvsr: The search for essential components in video super-resolution and beyond. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 4947–4956, 2021. 1
- [6] Zhu Chao, Dong Hang, Pan Jinshan, Liang Boyang, Huang Yuhao, Fu Lean, and Wang Fei. Deep recurrent neural network with multi-scale bi-directional propagation for video deblurring. In *AAAI*, 2022. 2
- [7] Chen Chen, Qifeng Chen, Minh N. Do, and Vladlen Koltun. Seeing motion in the dark. In *Proceedings of the IEEE/CVF International Conference on Computer Vision (ICCV)*, October 2019. 2
- [8] Mark Chen, Alec Radford, Rewon Child, Jeffrey Wu, Heewoo Jun, David Luan, and Ilya Sutskever. Generative pre-training from pixels. In *International conference on machine learning*, pages 1691–1703. PMLR, 2020. 3
- [9] Xiangxiang Chu, Zhi Tian, Yuqing Wang, Bo Zhang, Haibing Ren, Xiaolin Wei, Huaxia Xia, and Chunhua Shen. Twins: Revisiting spatial attention design in vision transformers. *arXiv preprint arXiv:2104.13840*, 2021. 2, 3
- [10] Zihang Dai, Zhilin Yang, Yiming Yang, Jaime Carbonell, Quoc V Le, and Ruslan Salakhutdinov. Transformer-xl: Attentive language models beyond a fixed-length context. *arXiv preprint arXiv:1901.02860*, 2019. 2
- [11] Jacob Devlin, Ming-Wei Chang, Kenton Lee, and Kristina Toutanova. Bert: Pre-training of deep bidirectional transformers for language understanding. *arXiv preprint arXiv:1810.04805*, 2018. 1, 2
- [12] Alexey Dosovitskiy, Lucas Beyer, Alexander Kolesnikov, Dirk Weissenborn, Xiaohua Zhai, Thomas Unterthiner, Mostafa Dehghani, Matthias Minderer, Georg Heigold, Sylvain Gelly, et al. An image is worth 16x16 words: Transformers for image recognition at scale. *arXiv preprint arXiv:2010.11929*, 2020. 2, 3
- [13] 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. 2, 5, 8
- [14] 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
- [15] Peng Gao, Teli Ma, Hongsheng Li, Jifeng Dai, and Yu Qiao. Convmae: Masked convolution meets masked autoencoders. *arXiv preprint arXiv:2205.03892*, 2022. 2, 3, 4
- [16] Mathias Gehrig, Mario Millhäusler, Daniel Gehrig, and Davide Scaramuzza. E-raft: Dense optical flow from event cameras. In *2021 International Conference on 3D Vision (3DV)*, pages 197–206. IEEE, 2021. 2
- [17] 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. 5, 6
- [18] Benjamin Graham and Laurens van der Maaten. Submanifold sparse convolutional networks. *arXiv preprint arXiv:1706.01307*, 2017. 4
- [19] Meng-Hao Guo, Jun-Xiong Cai, Zheng-Ning Liu, Tai-Jiang Mu, Ralph R Martin, and Shi-Min Hu. Pct: Point cloud transformer. *Computational Visual Media*, 7(2):187–199, 2021. 3
- [20] Kaiming He, Xinlei Chen, Saining Xie, Yanghao Li, Piotr Dollár, and Ross Girshick. Masked autoencoders are scalable vision learners. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 16000–16009, 2022. 1, 2, 3
- [21] Markus Hofinger, Samuel Rota Bulò, Lorenzo Porzi, Arno Knapitsch, Thomas Pock, and Peter Kontschieder. Improving optical flow on a pyramid level. In *European Conference on Computer Vision*, pages 770–786. Springer, 2020. 2
- [22] Berthold KP Horn and Brian G Schunck. Determining optical flow. *Artificial intelligence*, 17(1-3):185–203, 1981. 2
- [23] Zhaoyang Huang, Xiaokun Pan, Weihong Pan, Weikang Bian, Yan Xu, Ka Chun Cheung, Guofeng Zhang, and Hongsheng Li. Neuralmarker: A framework for learning general marker correspondence. *ACM Transactions on Graphics (TOG)*, 41(6):1–10, 2022. 2
- [24] 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, 3, 6
- [25] Zhaoyang Huang, Han Zhou, Yijin Li, Bangbang Yang, Yan Xu, Xiaowei Zhou, Hujun Bao, Guofeng Zhang, and Hongsheng Li. Vs-net: Voting with segmentation for visual localization. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 6101–6111, 2021. 2
- [26] Tak-Wai Hui, Xiaoou Tang, and Chen Change Loy. Liteflownet: 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. 6
- [27] Tak-Wai Hui, Xiaoou Tang, and Chen Change Loy. A lightweight optical flow cnn—revisiting data fidelity and reg-

- ularization. *IEEE transactions on pattern analysis and machine intelligence*, 43(8):2555–2569, 2020. 6
- [28] 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. 2
- [29] 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. 6, 7
- [30] Shihao Jiang, Dylan Campbell, Yao Lu, Hongdong Li, and Richard Hartley. Learning to estimate hidden motions with global motion aggregation. *arXiv preprint arXiv:2104.02409*, 2021. 2, 6
- [31] Shihao Jiang, Yao Lu, Hongdong Li, and Richard Hartley. Learning optical flow from a few matches. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 16592–16600, 2021. 2, 6
- [32] Dahun Kim, Sanghyun Woo, Joon-Young Lee, and In So Kweon. Deep video inpainting. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 5792–5801, 2019. 1
- [33] Daniel Kondermann, Rahul Nair, Katrin Honauer, Karsten Krispin, Jonas Andrusis, Alexander Brock, Burkhard Gusefeld, 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 Conference on Computer Vision and Pattern Recognition Workshops*, pages 19–28, 2016. 7
- [34] Wei-Sheng Lai, Jia-Bin Huang, Narendra Ahuja, and Ming-Hsuan Yang. Deep laplacian pyramid networks for fast and accurate super-resolution. In *Proceedings of the IEEE conference on computer vision and pattern recognition*, pages 624–632, 2017. 1
- [35] Dasong Li, Xiaoyu Shi, Yi Zhang, Xiaogang Wang, Hongwei Qin, and Hongsheng Li. No attention is needed: Grouped spatial-temporal shift for simple and efficient video restorers, 2022. 2
- [36] Dasong Li, Yi Zhang, Ka Chun Cheung, Xiaogang Wang, Hongwei Qin, and Hongsheng Li. Learning degradation representations for image deblurring. In *ECCV*, 2022. 2
- [37] Dasong Li, Yi Zhang, Ka Lung Law, Xiaogang Wang, Hongwei Qin, and Hongsheng Li. Efficient burst raw denoising with variance stabilization and multi-frequency denoising network, 2022. 2
- [38] Yijin Li, Zhaoyang Huang, Shuo Chen, Xiaoyu Shi, Hongsheng Li, Hujun Bao, Zhaopeng Cui, and Guofeng Zhang. Blinkflow: A dataset to push the limits of event-based optical flow estimation. *arXiv preprint arXiv:2303.07716*, 2023. 2
- [39] Yijin Li, Xinyang Liu, Wenqi Dong, Han Zhou, Hujun Bao, Guofeng Zhang, Yinda Zhang, and Zhaopeng Cui. Deltar: Depth estimation from a light-weight tof sensor and rgb image. In *Computer Vision—ECCV 2022: 17th European Conference, Tel Aviv, Israel, October 23–27, 2022, Proceedings, Part I*, pages 619–636. Springer, 2022. 2
- [40] Rui Liu, Hanming Deng, Yangyi Huang, Xiaoyu Shi, Lewei Lu, Wenxiu Sun, Xiaogang Wang, Jifeng Dai, and Hongsheng Li. Decoupled spatial-temporal transformer for video inpainting. *arXiv preprint arXiv:2104.06637*, 2021. 3
- [41] Rui Liu, Hanming Deng, Yangyi Huang, Xiaoyu Shi, Lewei Lu, Wenxiu Sun, Xiaogang Wang, Jifeng Dai, and Hongsheng Li. Fuseformer: Fusing fine-grained information in transformers for video inpainting. In *Proceedings of the IEEE/CVF International Conference on Computer Vision*, pages 14040–14049, 2021. 3
- [42] Shuaicheng Liu, Kunming Luo, Nianjin Ye, Chuan Wang, Jue Wang, and Bing Zeng. Oiflow: Occlusion-inpainting optical flow estimation by unsupervised learning. *IEEE Transactions on Image Processing*, 30:6420–6433, 2021. 8
- [43] Ze Liu, Yutong Lin, Yue Cao, Han Hu, Yixuan Wei, Zheng Zhang, Stephen Lin, and Baining Guo. Swin transformer: Hierarchical vision transformer using shifted windows. *arXiv preprint arXiv:2103.14030*, 2021. 2, 3
- [44] Kunming Luo, Chuan Wang, Shuaicheng Liu, Haoqiang Fan, Jue Wang, and Jian Sun. Upflow: Upsampling pyramid for unsupervised optical flow learning. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*, pages 1045–1054, June 2021. 8
- [45] 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. 5
- [46] Junjie Ni, Yijin Li, Zhaoyang Huang, Hongsheng Li, Hujun Bao, Zhaopeng Cui, and Guofeng Zhang. Pats: Patch area transportation with subdivision for local feature matching. *arXiv preprint arXiv:2303.07700*, 2023. 2
- [47] Yatian Pang, Wenxiao Wang, Francis EH Tay, Wei Liu, Yonghong Tian, and Li Yuan. Masked autoencoders for point cloud self-supervised learning. *arXiv preprint arXiv:2203.06604*, 2022. 3
- [48] AJ Piergiovanni and Michael S Ryoo. Representation flow for action recognition. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 9945–9953, 2019. 1
- [49] 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. 2
- [50] Mengye Ren, Andrei Pokrovsky, Bin Yang, and Raquel Urtasun. Sbnnet: Sparse blocks network for fast inference. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, pages 8711–8720, 2018. 4
- [51] Mehdi SM Sajjadi, Raviteja Vemulapalli, and Matthew Brown. Frame-recurrent video super-resolution. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, pages 6626–6634, 2018. 1
- [52] Xiaoyu Shi, Zhaoyang Huang, Weikang Bian, Dasong Li, Manyuan Zhang, Ka Chun Cheung, Simon See, Hongwei

- Qin, Jifeng Dai, and Hongsheng Li. Videoflow: Exploiting temporal cues for multi-frame optical flow estimation. *arXiv preprint arXiv:2303.08340*, 2023. [2](#)
- [53] Hyeonseok Son, Junyong Lee, Jonghyeop Lee, Sunghyun Cho, and Seungyong Lee. Recurrent video deblurring with blur-invariant motion estimation and pixel volumes. *ACM Transactions on Graphics (TOG)*, 40(5), 2021. [2](#)
- [54] 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. [8](#)
- [55] 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. [2](#), [6](#)
- [56] Deqing Sun, Stefan Roth, and Michael J Black. A quantitative analysis of current practices in optical flow estimation and the principles behind them. *International Journal of Computer Vision*, 106(2):115–137, 2014. [2](#)
- [57] 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. [7](#)
- [58] 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. [2](#), [6](#)
- [59] 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. [2](#), [6](#)
- [60] Shangkun Sun, Yuanqi Chen, Yu Zhu, Guodong Guo, and Ge Li. Skflow: Learning optical flow with super kernels. *arXiv preprint arXiv:2205.14623*, 2022. [6](#)
- [61] 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](#)
- [62] 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. [2](#), [6](#)
- [63] Zhan Tong, Yibing Song, Jue Wang, and Limin Wang. Videomae: Masked autoencoders are data-efficient learners for self-supervised video pre-training. *arXiv preprint arXiv:2203.12602*, 2022. [3](#)
- [64] Ashish Vaswani, Noam Shazeer, Niki Parmar, Jakob Uszkoreit, Llion Jones, Aidan N Gomez, Łukasz Kaiser, and Illia Polosukhin. Attention is all you need. In *Advances in neural information processing systems*, pages 5998–6008, 2017. [2](#)
- [65] Shiyao Wang, Yucong Zhou, Junjie Yan, and Zhidong Deng. Fully motion-aware network for video object detection. In *Proceedings of the European Conference on Computer Vision (ECCV)*, September 2018. [1](#)
- [66] Zhenda Xie, Zheng Zhang, Yue Cao, Yutong Lin, Jianmin Bao, Zhuliang Yao, Qi Dai, and Han Hu. Simmim: A simple framework for masked image modeling. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 9653–9663, 2022. [3](#)
- [67] Haofei Xu, Jiaolong Yang, Jianfei Cai, Juyong Zhang, and Xin Tong. High-resolution optical flow from 1d attention and correlation. In *Proceedings of the IEEE/CVF International Conference on Computer Vision*, pages 10498–10507, 2021. [2](#)
- [68] Haofei Xu, Jing Zhang, Jianfei Cai, Hamid Reza Tofighi, 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. [2](#), [6](#)
- [69] Ning Xu, Linjie Yang, Yuchen Fan, Dingcheng Yue, Yuchen Liang, Jianchao Yang, and Thomas Huang. Youtube-vos: A large-scale video object segmentation benchmark. *arXiv preprint arXiv:1809.03327*, 2018. [5](#)
- [70] Rui Xu, Xiaoxiao Li, Bolei Zhou, and Chen Change Loy. Deep flow-guided video inpainting. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 3723–3732, 2019. [1](#)
- [71] Gengshan Yang and Deva Ramanan. Volumetric correspondence networks for optical flow. *Advances in neural information processing systems*, 32:794–805, 2019. [2](#), [6](#)
- [72] Zhichao Yin, Trevor Darrell, and Fisher Yu. Hierarchical discrete distribution decomposition for match density estimation. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 6044–6053, 2019. [6](#)
- [73] Xumin Yu, Lulu Tang, Yongming Rao, Tiejun Huang, Jie Zhou, and Jiwen Lu. Point-bert: Pre-training 3d point cloud transformers with masked point modeling. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 19313–19322, 2022. [3](#)
- [74] Yanhong Zeng, Jianlong Fu, and Hongyang Chao. Learning joint spatial-temporal transformations for video inpainting. In *European Conference on Computer Vision*, pages 528–543. Springer, 2020. [3](#)
- [75] 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. [2](#), [6](#)
- [76] Yi Zhang, Dasong Li, Xiaoyu Shi, Dailan He, Kangning Song, Xiaogang Wang, Hongwei Qin, and Hongsheng Li. Kbnnet: Kernel basis network for image restoration, 2023. [2](#)
- [77] Hengshuang Zhao, Li Jiang, Jiaya Jia, Philip HS Torr, and Vladlen Koltun. Point transformer. In *Proceedings of the IEEE/CVF International Conference on Computer Vision*, pages 16259–16268, 2021. [3](#)
- [78] 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. 6

- [79] 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. 6
- [80] Yuxuan Zhao, Ka Lok Man, Jeremy Smith, Kamran Siddique, and Sheng-Uei Guan. Improved two-stream model for human action recognition. *EURASIP Journal on Image and Video Processing*, 2020(1):1–9, 2020. 1
- [81] Xizhou Zhu, Jifeng Dai, Lu Yuan, and Yichen Wei. Towards high performance video object detection. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, pages 7210–7218, 2018. 1
- [82] Xizhou Zhu, Yujie Wang, Jifeng Dai, Lu Yuan, and Yichen Wei. Flow-guided feature aggregation for video object detection. In *Proceedings of the IEEE international conference on computer vision*, pages 408–417, 2017. 1