

SAMFlow: Eliminating Any Fragmentation in Optical Flow with Segment Anything Model

Shili Zhou, Ruian He, Weimin Tan*, Bo Yan*

School of Computer Science, Shanghai Key Laboratory of Intelligent Information Processing, Fudan University
slzhou19@fudan.edu.cn, rahe16@fudan.edu.cn, wmtan@fudan.edu.cn, byan@fudan.edu.cn

Abstract

Optical Flow Estimation aims to find the 2D dense motion field between two frames. Due to the limitation of model structures and training datasets, existing methods often rely too much on local clues and ignore the integrity of objects, resulting in fragmented motion estimation. Through theoretical analysis, we find the pre-trained large vision models are helpful in optical flow estimation, and we notice that the recently famous Segment Anything Model (SAM) demonstrates a strong ability to segment complete objects, which is suitable for solving the fragmentation problem. We thus propose a solution to embed the frozen SAM image encoder into FlowFormer to enhance object perception. To address the challenge of in-depth utilizing SAM in non-segmentation tasks like optical flow estimation, we propose an Optical Flow Task-Specific Adaption scheme, including a Context Fusion Module to fuse the SAM encoder with the optical flow context encoder, and a Context Adaption Module to adapt the SAM features for optical flow task with Learned Task-Specific Embedding. Our proposed SAMFlow model reaches **0.86/2.10** clean/final EPE and **3.55/12.32** EPE/F1-all on Sintel and KITTI-15 training set, surpassing FlowFormer by **8.5%/9.9%** and **13.2%/16.3%**. Furthermore, our model achieves state-of-the-art performance on the Sintel and KITTI-15 benchmarks, **ranking #1** among all two-frame methods on Sintel clean pass.

Introduction

Optical flow is a fundamental task in computer vision, which estimates the pixel-level correspondences between frames. As an important paradigm to exploit video temporal continuity, it has applications in many video-related downstream tasks, such as frame interpolation (Huang et al. 2022b), video inpainting (Gao et al. 2020) and action recognition (Sun et al. 2018b). With the advent of advanced neural network architectures, many powerful optical flow estimation models have been proposed (Dosovitskiy et al. 2015; Sun et al. 2018a; Teed and Deng 2020; Jiang et al. 2021a; Huang et al. 2022a).

*Corresponding author: Weimin Tan and Bo Yan. This work is supported by NSFC (GrantNo.: U2001209 and 62372117) and Natural Science Foundation of Shanghai (21ZR1406600). Copyright © 2024, Association for the Advancement of Artificial Intelligence (www.aaai.org). All rights reserved.

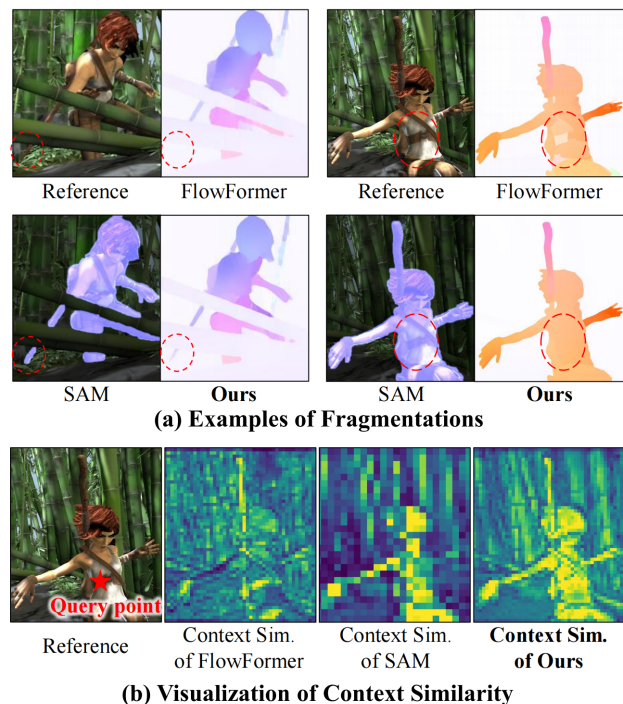


Figure 1: (a) Examples of fragmentation in optical flow estimation. We observe that SAM is able to segment the whole object. Thus, we propose our SAMFlow to eliminate fragmentation. (b) Visualization of the context similarity with the query point.

Although leaps and bounds have been made, existing optical flow estimation methods are still limited by two factors: 1) **Scarcity of well-labeled datasets**. Since it is difficult to obtain pixel-level motion annotations in the real world, optical flow datasets are usually constructed using artificial synthesis schemes. For example, some works (Dosovitskiy et al. 2015; Sun et al. 2021) try to construct datasets from images and generate motions with simple 2D transformations. At the same time, other works (Mayer et al. 2016; Gaidon et al. 2016) generate datasets of virtual scenes with the 3D rendering engine. Compared with natural scenes, these synthetic datasets have limited diversity and realism, resulting

in insufficient training of existing optical flow models. 2) **Lack of high-level understanding.** Human perception of motion is closely linked to the understanding of objects. Instead, optical flow models only focus on the local low-level clues, leading to incorrect “**fragmentation**” results. Here, fragmentation refers to erroneous fragmented optical flow predictions for the same object. Figure 1 (a) give examples of fragmentation in optical flow caused by occlusion and complex lighting/textures. Some previous studies (Sun et al. 2022; Jiang et al. 2021a) also try to solve this problem by using larger receptive fields or global motion aggregation. However, these simple structural improvements cannot fundamentally eliminate fragmentation.

The recent pre-trained large vision models that have received considerable attention are highly suitable for addressing the aforementioned two challenges. (i) (Shi et al. 2023) and (Dong, Cao, and Fu 2023) have shown that pre-training with data and supervision beyond optical flow can strengthen optical flow estimation, which implies that pre-trained large vision models can leverage a wide range of unlabeled image and video data to circumvent the problem of insufficient optical flow datasets. (ii) The visual representation learned through pre-training contains the high-level understanding we need. Therefore, fusion with large vision models may further enhance optical flow estimation. Among the large vision models, Segment Anything Model (SAM) (Kirillov et al. 2023) is one of the most suitable for optical flow estimation. As shown in Figure 1 (a), SAM can segment entire objects under occlusion and other confusing environments, which is exactly the solution to fragmentation in optical flow. Here, we propose using SAM’s features as the SAM image encoder occupies most of the SAM parameters and knowledge.

However, it is challenging to effectively harness SAM features for non-segmentation tasks such as optical flow estimation due to the absence of task-specific knowledge. As illustrated in Figure 1(b), while SAM’s feature yields a superior similarity map compared to FlowFormer, it loses numerous details, posing an obstacle to optical flow estimation. Therefore, we propose an Optical Flow Task-Specific Adaptation scheme to address the challenge. First, we fuse the SAM encoder with the task-specific encoder with a Context Fusion Module (CFM). Next, we introduce a Context Adaption Module (CAM) to inject more task-specific knowledge of optical flow into the fused features via Two-Way Attention (TWA) blocks and Learned Task-Specific Embedding (LTSE) tokens. With the above designs, our proposed SAMFlow achieves remarkable performance, reaching **0.86/2.10** clean/final EPE on Sintel (Butler et al. 2012) training set and **3.55/12.32** EPE/F1-all on KITTI-15 (Geiger et al. 2013) training set, surpassing Flowformer by **8.5%/9.9%** and **13.2%/16.3%**. Furthermore, we upload our fine-tuned models to the benchmark sites of Sintel and KITTI-15, which shows significant superiority, **ranking #1** among all two-frame methods on Sintel clean pass.

In summary, our contributions are as follows:

- For the first time, we investigate the feasibility of utilizing pre-trained SAM in optical flow estimation, and we thus propose SAMFlow, a novel approach aimed at en-

hancing the accuracy of optical flow estimation by effectively addressing issues of fragmentation.

- To prevent the task mismatch from affecting the accuracy, we propose an Optical Flow Task-Specific Adaptation scheme by introducing the CFM to fuse the SAM encoder with the optical flow context encoder, and the CAM to further adapt for optical flow estimation, improving the effectiveness of SAMFlow significantly.
- Our SAMFlow achieves state-of-the-art performance on both generalization and dataset-specific evaluations, surpassing Flowformer with a large margin and ranking #1 among all two-frame methods on the clean pass of the Sintel benchmark.

Related Works

Optical Flow

Optical flow has been studied for many years as a fundamental vision task. Traditional methods such as (Lucas and Kanade 1981) and (Horn and Schunck 1981) regard optical flow estimation as an energy optimization task and use human-designed data and prior terms as optimization objectives, which cannot satisfy the complex motion in natural images. In recent years, benefiting from the emergence of deep learning and large-scale synthetic optical flow datasets, most high-performance optical flow estimation methods learn optical flow automatically in an end-to-end manner. Model design and data collection replace the data and prior term, and become the focus of today’s optical flow algorithm researchers.

For the model, the researchers successively introduced convolutional network as FlowNet (Dosovitskiy et al. 2015), multi-scale network as PWC-Net (Sun et al. 2018a), recurrent network as RAFT (Teed and Deng 2020), Transformer as FlowFormer (Huang et al. 2022a) and other structures, making the inherent learning ability of the optical flow model enhanced gradually. Meanwhile, reducing model time-consuming is also a research direction like (Cheng et al. 2023).

For data, from Chairs (Dosovitskiy et al. 2015), Things (Mayer et al. 2016) to later AutoFlow (Sun et al. 2021), Spring (Mehl et al. 2023), the diversity and authenticity of synthetic datasets are increasing, while the richness of real datasets is also slowly increasing (KITTI (Geiger et al. 2013) and HD1K (Kondermann et al. 2016)).

These efforts open up the possibility of increasingly powerful optical flow estimation models. However, the scarcity of data and the limitation of model design are still the core problems of optical flow.

Pre-trained Large Vision Model

The introduction to the pre-trained large vision model can be divided into the model architectures and the pre-training methods. We start by introducing the model architectures. In the early years, researchers use convolutional neural networks (CNN) as the basic architecture of computer vision, proposing VGG (Simonyan and Zisserman 2014), ResNet (He et al. 2016), *etc.* Recently, inspired by the success of

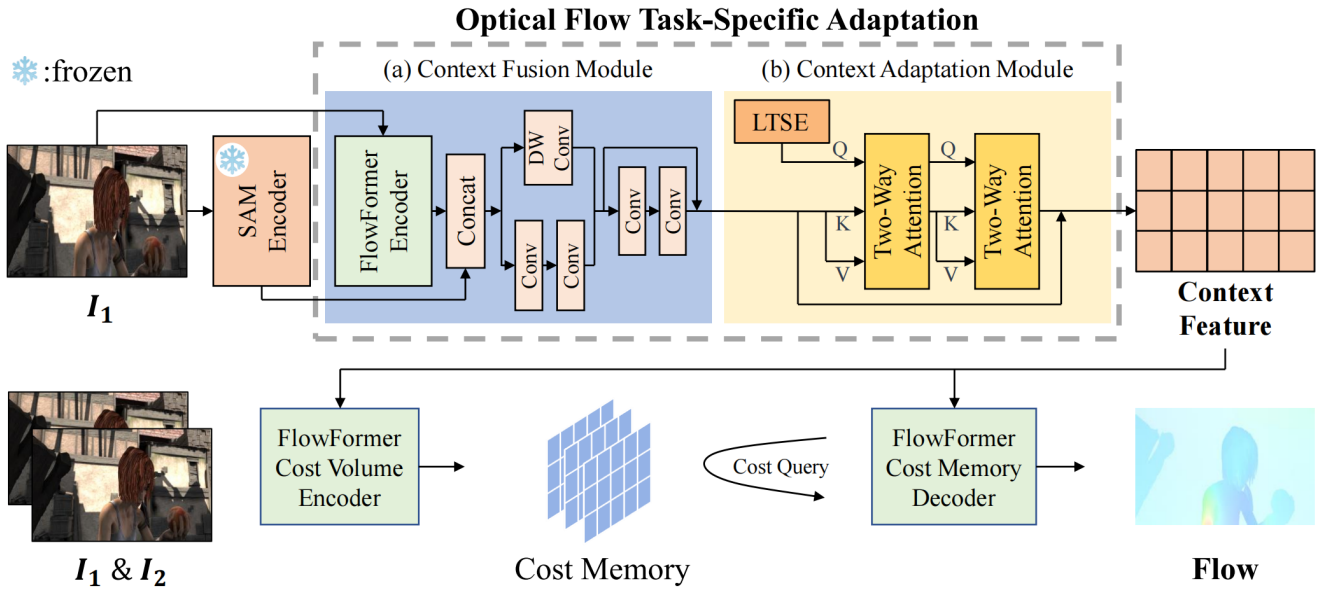


Figure 2: The overview of our SAMFlow, which utilizes the frozen SAM image encoder to boost the object perception of the optical flow model FlowFormer. We design two modules for in-depth utilizing SAM, including: (a) the CFM, which fuses SAM features with FlowFormer encoder, and (b) the CAM, which adapts the features with the Learned Task-Specific Embedding.

Transformer in natural language processing, Vision Transformer (ViT) (Dosovitskiy et al. 2020) is proposed, which has a stronger representational ability and can show obvious advantages under large-scale datasets.

Next, we introduce the pre-training methods. The early pre-training models using labeled data of pretext tasks, such as classification on ImageNet (Krizhevsky, Sutskever, and Hinton 2012). To use large-scale unlabeled data, researchers propose self-supervised pre-training methods, including contrastive learning (Chen et al. 2020), auto-encoding (Vincent et al. 2008), *etc.* A recent highlight paper (He et al. 2022) proposes Masked Auto-Encoder (MAE), which improves the traditional auto-encoder by dropping some patches when encoding to force the model to understand the image content.

Segment Anything Model

Segment Anything Model (SAM) (Kirillov et al. 2023) is a prompt-based segmentation model. The structure of SAM is divided into three parts: image encoder, prompt encoder, and decoder. The image encoder is a variant of ViT (Dosovitskiy et al. 2020), which has a large number of parameters, while the prompt encoder and the decoder are lightweight. SAM is fine-tuned from MAE with a large amount of labeled segmentation data. This work also presents an impressive training data generation scheme: it uses manual labeling/correction and model learning/prediction as two complementary processes, which can create billions of segmentation labels with low labor costs. Large-scale labeled training data endows SAM with robust understanding and segmentation capabilities, which we find suitable for eliminating fragmentation in optical flow estimation.

Proposed Method

Theoretical Analysis

Optical flow estimation methods aim to find the mapping $\zeta : (I_1, I_2) \Rightarrow F$, where I_1 and I_2 are two adjacent frames from a video, and F is the 2D optical flow field. From a probabilistic point of view, an optical flow network can be expressed as:

$$F^* = \zeta_{\theta}(I_1, I_2) = \underset{F}{\operatorname{argmax}} p(F | I_1, I_2) \quad (1)$$

where F^* is the estimated most likely optical flow, ζ_{θ} is the optical flow network with parameters θ , and $p(F | I_1, I_2)$ is the posterior distribution of optical flow.

In order to facilitate a comprehensive analysis, we use Bayes' theorem to extend $p(F | I_1, I_2)$:

$$\begin{aligned} p(F | I_1, I_2) &= \frac{p(F)p(I_1 | F)p(I_2 | I_1, F)}{p(I_1, I_2)} \\ &= \frac{p(I_1)p(F | I_1)p(I_2 | I_1, F)}{p(I_1, I_2)} \end{aligned} \quad (2)$$

To find the optimal F , we omit the unrelated term $p(I_1)$ and $p(I_1, I_2)$, and take the logarithm to separate the multiplication terms. Thus, we get another formation of F^* as in Formula 3, which consists of a **cost query term** and a **context term**.

$$F^* = \underset{F}{\operatorname{argmax}} \left\{ \underbrace{\log p(I_2 | I_1, F)}_{\text{cost query}} + \underbrace{\log p(F | I_1)}_{\text{context}} \right\} \quad (3)$$

The cost query term encompasses the interrelation between I_2 and I_1 with F . To achieve this, optical flow models such as (Sun et al. 2018a; Teed and Deng 2020; Jiang et al.

2021a) construct the 4D cost volumes and make cost queries with flow guidance.

The contextual term provides an alternative information source for optical flow estimation, requiring the model to comprehend the image context more deeply. Earlier approaches like (Teed and Deng 2020; Jiang et al. 2021a; Sun et al. 2022) overlook this aspect, with the extracted contextual features being constrained to local cues. In contrast, our endeavor involves integrating pre-trained large vision models to enhance higher-level understanding. However, not all pre-trained large vision models prove apt for optical flow estimation, as certain models exclusively capture global semantics at the cost of losing spatial detail features, resulting in limited contributions to optical flow. Empirically, SAM is a suitable candidate with its capability to generate pixel-level outputs akin to optical flow, as shown in Figure 1.

Overview

As shown in Figure 2, we redesign the context feature extraction process of the backbone model FlowFormer by utilizing the image encoder of SAM, which has powerful object perception to solve the fragmentation of optical flow estimation. We call the proposed new model as SAMFlow. Moreover, as SAM does not acquire task-specific prior knowledge related to optical flow, we design an Optical Flow Task-Specific Adaption scheme, which includes a Context Fusion Module and a Context Adaption Module. In the following two subsections, we first introduce some minor modifications to unlock the resolution requirement of the SAM encoder; then, we introduce the CFM and CAM in detail.

Modifications for Resolution

The image encoder of SAM is a ViT that only accepts a fixed input resolution of 1024×1024 . Considering the memory and time cost, this resolution is unable to use in the optical flow training framework. Therefore, we slightly modify the SAM encoder to unlock the resolution limitation. The details are provided in our supplementary material.

Optical Flow Task-Specific Adaption

Context Fusion Module: Utilizing pre-trained large vision models for dissimilar tasks faces the challenge of knowledge unmatching. For example, the local details are essential clues for optical flow, while they are dropped for understanding tasks. Thus, we propose the CFM to combine the high-level understanding of SAM and the low-level clues for optical flow by using the SAM encoder and FlowFormer encoder simultaneously.

As shown in Figure 2(a), we first concatenate the SAM and Flowformer features. Subsequently, we mix them with two residual convolutional blocks. In the former block, the features will be processed by two branches: the main branch contains two 3×3 convolutional layers, which fuse the features and reduce channels, and the other branch uses depth-wise convolution directly to reduce the channels number and keep it consistent to the main branch. We add the results of the two branches as the output of the block. The latter residual block has almost the same structure except

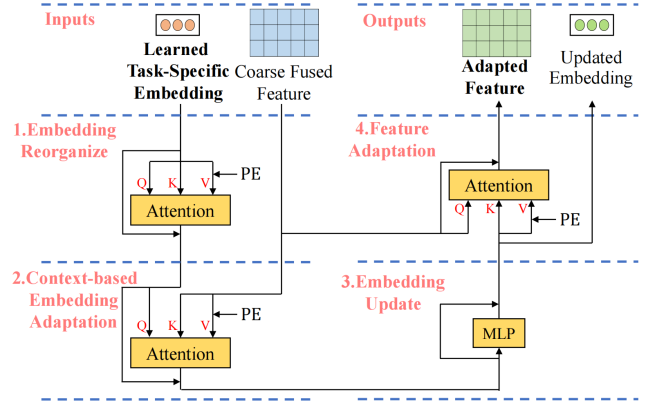


Figure 3: Our Context Adaptation Module to adapt SAM features for optical flow with the Learned Task-Specific Embedding. For the sake of brevity, only one Two-Way Attention is shown. PE is the positional embedding.

for depth-wise convolution since there is no difference in channel numbers between the input and output. Overall, this module can be represented by Formula 4, 5, 6 and 7.

$$\Phi_S = E_S(I), \Phi_F = E_F(I) \quad (4)$$

$$\Phi_{\parallel} = \Phi_S \parallel \Phi_F \quad (5)$$

$$\bar{\Phi}_C = Conv_2(Conv_1(\Phi_{\parallel})) + \Delta(\Phi_{\parallel}) \quad (6)$$

$$\Phi_C = Conv_4(Conv_3(\bar{\Phi}_C)) + \bar{\Phi}_C \quad (7)$$

where E_S and E_F are the SAM encoder and FlowFormer Encoder, Φ_S and Φ_F are the extracted features, \parallel is concatenation operator, $Conv_k$ corresponds to the k -th convolution layer. Δ is depth-wise convolution. Φ_{\parallel} and $\bar{\Phi}_C$ are intermediate variables, and Φ_C is the output of the CFM. The normalization and activation layers are omitted for brevity.

Context Adaption Module: To better utilize the task-specific knowledge to accomplish task adaptation of optical flow, we propose the Context Adaption Module, as shown in Figure 2(b) and 3. Inspired by Perceiver IO (Jaegle et al. 2021) and the mask decoder of SAM, we make the following design in the Context Adaption Module: we use Learned Task-Specific Embedding (LTSE) tokens to store some task-specific priors of optical flow, and use Two-Way Attention (TWA) blocks to inject those priors into the context feature for adaptation.

The LTSE is implemented as a set of learnable offsets of shape $K \times D$, which will be automatically optimized during the training process. We empirically set K to 3 and D to 256. Meanwhile, each TWA contains four steps:

1) **Embedding Reorganize:** as shown in Formula 8, a self-attention layer is used to reorganize the embedding of optical flow estimation task Ω_T , which is the LTSE for the first TWA block.

$$\bar{\Omega}_T = \Omega_T + Att_1(\Omega_T, \Omega_T, \Omega_T + PE) \quad (8)$$

where PE is the positional embedding. Att_1 is the first attention layer, which requires query, key, and value to be fed in order. We omit the normalization and activation layers

Training Stage	Method	Sintel(train)		KITTI-15(train)	
		clean	final	EPE	F1
C+T	HD3	3.84	8.77	13.17	24.0
	LiteFlowNet	2.48	4.04	10.39	28.5
	PWC-Net	2.55	3.93	10.35	33.7
	LiteFlowNet2	2.24	3.78	8.97	25.9
	S-Flow	1.30	2.59	4.60	15.9
	RAFT	1.43	2.71	5.04	17.4
	FM-RAFT	1.29	2.95	6.80	19.3
	GMA	1.30	2.74	4.69	17.1
	GMFlow	1.08	2.48	-	-
	GMFlowNet	1.14	2.71	4.24	15.4
	CRAFT	1.27	2.79	4.88	17.5
	SKFlow	1.22	2.46	4.47	15.5
	FlowFormer	0.94	2.33	4.09	14.72
	FlowFormer++	0.90	2.30	3.93	14.13
	Ours	0.87	2.11	3.44	12.28

Table 1: Generalization performance evaluation on Sintel and KITTI-15 train sets.

here for brevity and do not expand the attention layer in detail. $\bar{\Omega}_T$ is the intermediate result of this step.

2) **Context-based Embedding Adaption**: as shown in Formula 9, we use a cross-attention layer to adapt the embedding with the context feature for better handling the input cases.

$$\hat{\Omega}_T = \bar{\Omega}_T + \text{Att}_2(\bar{\Omega}_T, \Phi_C, \Phi_C + PE) \quad (9)$$

where $\hat{\Omega}$ is the adaptation result of this step.

3) **Embedding Update**: as shown in Formula 10, we use a Multi-layer Perceptron (MLP) to update the query.

$$\Omega_U = \hat{\Omega}_T + \text{MLP}(\hat{\Omega}_T) \quad (10)$$

where MLP is the Multi-layer Perceptron, and Ω_U is the updated embedding.

4) **Feature Adaption**: we use the updated embedding to adapt the context feature for optical flow tasks with a cross-attention layer, as shown in Formula 11.

$$\Phi_C^A = \Phi_C + \text{Att}_3(\Phi_C, \Omega_U, \Omega_U + PE) \quad (11)$$

where Φ_C^A is the adapted context feature under the guidance of optical flow task-specific queries.

Finally, we use an addition operation to blend the results of the two modules.

Experiment

Settings

Training Settings We follow the setup of previous work (Huang et al. 2022a) and divide the training into two stages: C+T-Stage and C+T+S+K+H-stage. To speed up training, we skip the stage of training on the Chairs dataset by using FlowFormer-things checkpoint as initialization, and the SAM encoder is kept frozen during training.

Test Settings For testing, we adopt the tiling strategy (Jaegle et al. 2021) to bridge the resolution gap between training and testing data.

Training Stage	Method	Sintel(test)		KITTI-15(test)
		clean	final	F1-all
C+T+S+K+H	PWC-Net+	3.45	4.60	7.72
	VCN	2.81	4.40	6.30
	MaskFlowNet	2.52	4.17	6.10
	S-Flow	1.50	2.67	4.64
	RAFT	1.94	3.18	5.10
	RAFT*	1.61	2.86	5.10
	FM-RAFT	1.72	3.60	6.17
	GMA	1.40	2.88	5.15
	GMA*	1.39	2.47	5.15
	GMFlow	1.74	2.90	9.32
	GMFlowNet	1.39	2.65	4.79
	CRAFT	1.45	2.42	4.79
	SKFlow*	1.28	2.23	4.84
	FlowFormer	1.16	2.09	4.68
	FlowFormer++	1.07	1.94	4.52
	Ours	1.00	2.08	4.49

Table 2: Benchmark evaluation on Sintel and KITTI-15 test sets. The models with * adopt the warm-start strategy proposed in (Teed and Deng 2020).

Methods	Sintel (train) Occ.		Sintel (test) Occ.	
	clean	final	clean	final
RAFT	5.36	7.09	9.65	14.68
GMA	4.25	6.22	7.96	12.50
SKFlow	3.44	4.52	7.25	11.42
FlowFormer	2.76	3.60	7.16	11.30
FlowFormer++	2.54	3.41	6.64	10.63
Ours	2.24	2.99	5.97	10.60

Table 3: Evaluation in occluded area of Sintel train and test sets.

Quantitative Comparison

We first use the model trained in the C+T-stage for evaluating the generalization performance on the training sets of Sintel and KITTI. Then, we uploaded the results of the C+T+S+K+H-stage model and the K-stage model to the Sintel Benchmark website and the KITTI Benchmark website to compare the dataset-specific accuracy with the SOTA methods, including HD3 (Yin, Darrell, and Yu 2019), LiteFlowNet (Hui, Tang, and Loy 2018), PWC-Net (Sun et al. 2018a), PWC-Net++ (Sun et al. 2019), LiteFlowNet2 (Hui, Tang, and Loy 2020), S-Flow (Zhang et al. 2021), RAFT (Teed and Deng 2020), FM-RAFT (Jiang et al. 2021b), GMA (Jiang et al. 2021a), GMFlow (Xu et al. 2022), GMFlowNet (Zhao et al. 2022), CRAFT (Sui et al. 2022), SKFlow (Sun et al. 2022), FlowFormer (Huang et al. 2022a) and FlowFormer++ (Shi et al. 2023).

Generalization Performance As shown in Table 1, for the C+T-stage, our model achieves the best performance on all metrics on the training set of Sintel and KITTI-15 datasets. The EPEs of our SAMFlow on Sintel clean and fi-

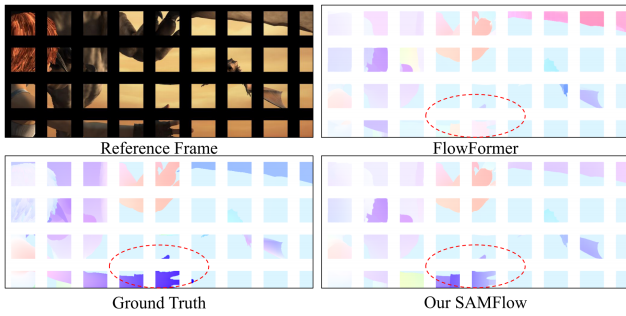


Figure 4: An example of the fragmentation attack, where our SAMFlow shows robustness over FlowFormer.




Figure 5: The average EPE of sintel clean and final pass under the fragmentation attack with different masked ratios.

nal pass reach **0.86** and **2.10**. SAMFlow also achieve **3.55** EPE and **12.32** F1 on KITTI-15 datasets. It is worth noting that FlowFormer uses two different model checkpoints with different training patch-size to obtain better performance on Sintel and KITTI. In contrast, our method uses the same checkpoint when evaluating both datasets. Nevertheless, our SAMFlow still easily surpasses the performance of FlowFormer, reducing Sintel clean/final EPE and KITTI-15 EPE/F1 by **8.5%/9.9%** and **13.2%/16.3%**, respectively.

Comparison on Benchmarks Table 2 also proves the dataset-specific performance of our SAMFlow. On the Sintel test set, our method achieves **1.00** clean EPE and **2.08** final EPE. Meanwhile, on the KITTI-15 test set, SAMFlow achieves **4.49** F1-all. Compared with FlowFormer, our method has achieved all-around improvement and can also defeat the new SOTA method FlowFormer++ on Sintel clean pass and KITTI-15. This demonstrates that our method brings significant accuracy improvements for optical flow estimation. The results can also be found on Sintel and KITTI-15 benchmark websites, where our SAMFlow **rank #1** among all two-frame methods on Sintel clean pass.

Evaluation in Occluded Area We give the comparison under occluded area, one of the significant sources of fragmentation, on the Sintel train set and the Sintel benchmark (test) as shown in Table 3. Compared with FlowFormer, our method has improved by **18.84%/17.94%** and **16.62%/6.19%** on Sintel train and test, respectively. At the same time, it surpasses FlowFormer++ on the Sintel benchmark and reaches the best.

Fragmentation Attack

To further demonstrate the ability of our method to eliminate fragmentation, we design the fragmentation attack, which splits the images into discrete parts using a grid-style mask, as shown in Figure 4. By controlling the thickness and density of the mask grid, we can mask out the images of the Sintel dataset at different ratios, creating different degrees of fragmentation. As shown in Figure 5, we compare the robustness of GMFlow, SKFlow, FlowFormer, and our model under fragmentation attack with 0%, 20%, 30% and 40% masked ratios. Attacks greater than 40% are meaningless because too much information has been lost, so we ignore these cases. It can be observed that SKFlow is greatly affected by fragmentation attacks. Its EPE increases sharply when adding 20% mask. GMFlowNet and FlowFormer are also trapped in isolated local clues caused by fragmentation, showing a noticeable performance hit. With the object perception with SAM encoder, our method exhibits strong robustness by finding the relation between the image content in different grids, achieving better results than FlowFormer.

Visualization

Figure 6 shows two examples from Sintel and KITTI-15 for qualitative comparison, corresponding to fragmentation caused by occlusion and complex lighting/textures, respectively. We visualize the optical flow fields by mapping them to the color space and the context features by computing the feature similarity between all pixels and the chosen query points. We can find that FlowFormer context similarity is unordered and cannot guarantee the integrity of moving objects, resulting in the missing leg of the girl in the first example and the holes of the car in the second example. With the perception of objects, our SAMFlow gives better context features with apparent related objects and boundaries, thus enhancing the accuracy of optical flow greatly.

Ablation Study

We conduct a series of ablation studies to validate our SAMFlow, and the results are shown in Table 4.

Encoders and Modules: We compare different context feature settings to prove the effectiveness of our designs. The baseline model is FlowFormer, which only has the optical flow task-specific encoder. We first try using the SAM encoder instead of the FlowFormer encoder and find it brings some performance improvements on Sintel. However, the effect on authentic images (KITTI-15) is limited due to the lack of priors for optical flow estimation. Subsequently, we add our CFM, which fuses the FlowFormer and SAM features with residual convolutional blocks, significantly reducing the errors on KITTI-15 datasets. We further add our CAM to inject more task-specific knowledge into the context feature, which boosts the optical flow accuracy and achieves the best performances for both datasets. A CAM-only model was also added to the experiments to illustrate that both modules are necessary.

SAM Model Scale: We try the SAM image encoders of different SAM scales, including *SAM-H*, *SAM-B* and a tiny version *MobileSAM* (Zhang et al. 2023). The baseline model

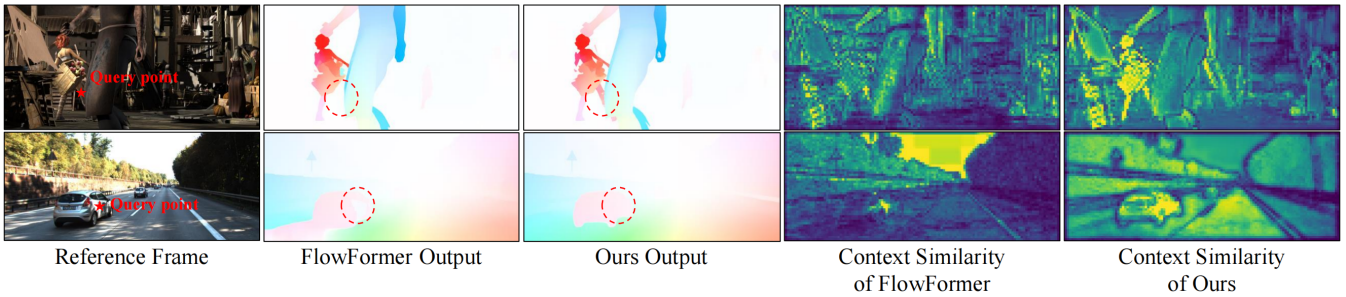


Figure 6: Two examples from Sintel and KITTI-15 for qualitative comparison corresponding to fragmentation caused by occlusion and complex lighting/textures, respectively.

Methods	Sintel(train)		KITTI-15(train)	
	clean	final	EPE	F1
FlowFormer Enc.	0.94	2.33	4.09	14.72
SAM Enc.	0.89	2.17	4.11	14.37
CFM	0.89	2.11	3.83	13.43
CAM	0.90	2.10	3.81	13.23
CFM + CAM	0.87	2.11	3.44	12.28
w/o. SAM	0.94	2.33	4.09	14.72
MobileSAM	0.88	2.19	3.78	13.22
SAM-B	0.88	2.26	3.57	12.45
SAM-H	0.87	2.11	3.44	12.28
ViT	1.08	2.38	4.01	13.47
MAE	1.01	2.35	3.91	13.56
DINO	0.94	2.26	4.07	13.76

Table 4: Ablation study of encoder type, modules, and the scale of SAM encoder. We bold the best value in each group.

is also listed, named *w/o. SAM*. All our models outperform the baseline (FlowFormer), which once again proves the effectiveness of our proposed method. Moreover, we find that larger encoders show better results in general. However, there is an exception on the final pass of the Sintel dataset, where the MobileSAM encoder performs better than SAM-B encoder. This may be due to the different architectures of MobileSAM and SAM-B encoders, which cause them to behave differently under some specific scenes.

Other Pre-trained Models Besides SAM, we also try ViT (Dosovitskiy et al. 2020), MAE (He et al. 2022) and DINO (Zhang et al. 2022). However, as shown in Table 4, they do not work well. The reasons are two-fold: on the one hand, their tasks does not need good representations of spatial content; on the other hand, the pre-trained ViT and MAE suffer from a limitation of very low-resolution (224×224 or 384×384), making them poorly suited for larger inputs.

Runtime Analysis

There might be doubts about the computational cost of utilizing the SAM encoder. In order to analyze it, we compare the performance and runtime of our three models of different scales with FlowFormer and FlowFormer++, and the




Figure 7: Runtime and accuracy comparison between FlowFormer, FlowFormer++, and our models with different SAM encoders, including SAM-B, SAM-H, and MobileSAM (MSAM). The x-axis is the average time of 100 runs of 384×1024 inputs, and the y-axis is the f1 score on KITTI.

results are presented in Figure 7. It can be found that our method can balance performance and runtime requirements by controlling the scale of the SAM encoder. Compared with FlowFormer, our SAMFlow w/. MSAM(MobileSAM) only slightly increases in runtime but shows a considerable drop in F1. Meanwhile, all of our three models are superior to FlowFormer++ in speed and performance, proving the practical significance of our method.

Conclusion

This paper focuses on the challenging fragmentation issues for optical flow estimation. We first give theoretical analysis for applying SAM feature in optical flow estimation. Thus, we propose SAMFlow, which incorporates SAM into the optical flow estimation network. Next, to address the unmatched task-specific knowledge between SAM and optical flow estimation, we introduce an Optical Flow Task-Specific Adaptation scheme, including the CFM and CAM. In experiments, we demonstrate the effectiveness of SAMFlow for fragmentation elimination and its superiority in terms of optical flow estimation accuracy, which achieves state-of-the-art performance, and ranks #1 among all two-frame methods on Sintel clean pass.

References

- Butler, D. J.; Wulff, J.; Stanley, G. B.; and Black, M. J. 2012. A naturalistic open source movie for optical flow evaluation. In *Computer Vision—ECCV 2012: 12th European Conference on Computer Vision, Florence, Italy, October 7-13, 2012, Proceedings, Part VI 12*, 611–625. Springer.
- Chen, T.; Kornblith, S.; Norouzi, M.; and Hinton, G. 2020. A simple framework for contrastive learning of visual representations. In *International conference on machine learning*, 1597–1607. PMLR.
- Cheng, R.; He, R.; Jiang, X.; Zhou, S.; Tan, W.; and Yan, B. 2023. Context-Aware Iteration Policy Network for Efficient Optical Flow Estimation. arXiv:2312.07180.
- Dong, Q.; Cao, C.; and Fu, Y. 2023. Rethinking Optical Flow from Geometric Matching Consistent Perspective. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, 1337–1347.
- Dosovitskiy, A.; Beyler, L.; Kolesnikov, A.; Weissenborn, D.; Zhai, X.; Unterthiner, T.; Dehghani, M.; Minderer, M.; Heigold, G.; Gelly, S.; et al. 2020. An image is worth 16x16 words: Transformers for image recognition at scale. *arXiv preprint arXiv:2010.11929*.
- Dosovitskiy, A.; Fischer, P.; Ilg, E.; Hausser, P.; Hazirbas, C.; Golkov, V.; Van Der Smagt, P.; Cremers, D.; and Brox, T. 2015. FlowNet: Learning optical flow with convolutional networks. In *Proceedings of the IEEE international conference on computer vision*, 2758–2766.
- Gaidon, A.; Wang, Q.; Cabon, Y.; and Vig, E. 2016. Virtual Worlds as Proxy for Multi-Object Tracking Analysis. In *CVPR*.
- Gao, C.; Saraf, A.; Huang, J.-B.; and Kopf, J. 2020. Flow-edge guided video completion. In *Computer Vision—ECCV 2020: 16th European Conference, Glasgow, UK, August 23–28, 2020, Proceedings, Part XII 16*, 713–729. Springer.
- Geiger, A.; Lenz, P.; Stiller, C.; and Urtasun, R. 2013. Vision meets robotics: The kitti dataset. *The International Journal of Robotics Research*, 32(11): 1231–1237.
- He, K.; Chen, X.; Xie, S.; Li, Y.; Dollár, P.; and Girshick, R. 2022. Masked autoencoders are scalable vision learners. In *Proceedings of the IEEE/CVF conference on computer vision and pattern recognition*, 16000–16009.
- He, K.; Zhang, X.; Ren, S.; and Sun, J. 2016. Deep residual learning for image recognition. In *Proceedings of the IEEE conference on computer vision and pattern recognition*, 770–778.
- Horn, B. K.; and Schunck, B. G. 1981. Determining optical flow. *Artificial intelligence*, 17(1-3): 185–203.
- Huang, Z.; Shi, X.; Zhang, C.; Wang, Q.; Cheung, K. C.; Qin, H.; Dai, J.; and Li, H. 2022a. Flowformer: A transformer architecture for optical flow. In *European Conference on Computer Vision*, 668–685. Springer.
- Huang, Z.; Zhang, T.; Heng, W.; Shi, B.; and Zhou, S. 2022b. Real-time intermediate flow estimation for video frame interpolation. In *European Conference on Computer Vision*, 624–642. Springer.
- Hui, T.-W.; Tang, X.; and Loy, C. C. 2018. Liteflownet: A lightweight convolutional neural network for optical flow estimation. In *Proceedings of the IEEE conference on computer vision and pattern recognition*, 8981–8989.
- Hui, T.-W.; Tang, X.; and Loy, C. C. 2020. A lightweight optical flow cnn—revisiting data fidelity and regularization. *IEEE transactions on pattern analysis and machine intelligence*, 43(8): 2555–2569.
- Jaegle, A.; Borgeaud, S.; Alayrac, J.-B.; Doersch, C.; Ionescu, C.; Ding, D.; Koppula, S.; Zoran, D.; Brock, A.; Shelhamer, E.; et al. 2021. Perceiver IO: A General Architecture for Structured Inputs & Outputs. In *International Conference on Learning Representations*.
- Jiang, S.; Campbell, D.; Lu, Y.; Li, H.; and Hartley, R. 2021a. Learning to estimate hidden motions with global motion aggregation. In *Proceedings of the IEEE/CVF International Conference on Computer Vision*, 9772–9781.
- Jiang, S.; Lu, Y.; Li, H.; and Hartley, R. 2021b. Learning optical flow from a few matches. In *Proceedings of the IEEE/CVF conference on computer vision and pattern recognition*, 16592–16600.
- Kirillov, A.; Mintun, E.; Ravi, N.; Mao, H.; Rolland, C.; Gustafson, L.; Xiao, T.; Whitehead, S.; Berg, A. C.; Lo, W.-Y.; et al. 2023. Segment anything. *arXiv preprint arXiv:2304.02643*.
- Kondermann, D.; Nair, R.; Honauer, K.; Krispin, K.; Andrusis, J.; Brock, A.; Gussfeld, B.; Rahimimoghaddam, M.; Hofmann, S.; Brenner, C.; et al. 2016. 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*, 19–28.
- Krizhevsky, A.; Sutskever, I.; and Hinton, G. E. 2012. ImageNet classification with deep convolutional neural networks. *Advances in neural information processing systems*, 25.
- Lucas, B. D.; and Kanade, T. 1981. An iterative image registration technique with an application to stereo vision. In *IJCAI’81: 7th international joint conference on Artificial intelligence*, volume 2, 674–679.
- Mayer, N.; Ilg, E.; Hausser, P.; Fischer, P.; Cremers, D.; Dosovitskiy, A.; and Brox, T. 2016. 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*, 4040–4048.
- Mehl, L.; Schmalfluss, J.; Jahedi, A.; Nalivayko, Y.; and Bruhn, A. 2023. Spring: A High-Resolution High-Detail Dataset and Benchmark for Scene Flow, Optical Flow and Stereo. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, 4981–4991.
- Shi, X.; Huang, Z.; Li, D.; Zhang, M.; Cheung, K. C.; See, S.; Qin, H.; Dai, J.; and Li, H. 2023. Flowformer++: Masked cost volume autoencoding for pretraining optical flow estimation. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, 1599–1610.

- Simonyan, K.; and Zisserman, A. 2014. Very deep convolutional networks for large-scale image recognition. *arXiv preprint arXiv:1409.1556*.
- Sui, X.; Li, S.; Geng, X.; Wu, Y.; Xu, X.; Liu, Y.; Goh, R.; and Zhu, H. 2022. Craft: Cross-attentional flow transformer for robust optical flow. In *Proceedings of the IEEE/CVF conference on Computer Vision and Pattern Recognition*, 17602–17611.
- Sun, D.; Vlasic, D.; Herrmann, C.; Jampani, V.; Krainin, M.; Chang, H.; Zabih, R.; Freeman, W. T.; and Liu, C. 2021. Autoflow: Learning a better training set for optical flow. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, 10093–10102.
- Sun, D.; Yang, X.; Liu, M.-Y.; and Kautz, J. 2018a. Pwcnet: Cnns for optical flow using pyramid, warping, and cost volume. In *Proceedings of the IEEE conference on computer vision and pattern recognition*, 8934–8943.
- Sun, D.; Yang, X.; Liu, M.-Y.; and Kautz, J. 2019. 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.
- Sun, S.; Chen, Y.; Zhu, Y.; Guo, G.; and Li, G. 2022. Sk-flow: Learning optical flow with super kernels. *Advances in Neural Information Processing Systems*, 35: 11313–11326.
- Sun, S.; Kuang, Z.; Sheng, L.; Ouyang, W.; and Zhang, W. 2018b. 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*, 1390–1399.
- Teed, Z.; and Deng, J. 2020. Raft: Recurrent all-pairs field transforms for optical flow. In *Computer Vision—ECCV 2020: 16th European Conference, Glasgow, UK, August 23–28, 2020, Proceedings, Part II 16*, 402–419. Springer.
- Vincent, P.; Larochelle, H.; Bengio, Y.; and Manzagol, P.-A. 2008. Extracting and composing robust features with denoising autoencoders. In *Proceedings of the 25th international conference on Machine learning*, 1096–1103.
- Xu, H.; Zhang, J.; Cai, J.; Rezatofighi, H.; and Tao, D. 2022. Gmflow: Learning optical flow via global matching. In *Proceedings of the IEEE/CVF conference on computer vision and pattern recognition*, 8121–8130.
- Yin, Z.; Darrell, T.; and Yu, F. 2019. Hierarchical discrete distribution decomposition for match density estimation. In *Proceedings of the IEEE/CVF conference on computer vision and pattern recognition*, 6044–6053.
- Zhang, C.; Han, D.; Qiao, Y.; Kim, J. U.; Bae, S.-H.; Lee, S.; and Hong, C. S. 2023. Faster Segment Anything: Towards Lightweight SAM for Mobile Applications. *arXiv preprint arXiv:2306.14289*.
- Zhang, F.; Woodford, O. J.; Prisacariu, V. A.; and Torr, P. H. 2021. Separable flow: Learning motion cost volumes for optical flow estimation. In *Proceedings of the IEEE/CVF international conference on computer vision*, 10807–10817.
- Zhang, H.; Li, F.; Liu, S.; Zhang, L.; Su, H.; Zhu, J.; Ni, L.; and Shum, H.-Y. 2022. DINO: DETR with Improved Denoising Anchor Boxes for End-to-End Object Detection. In *The Eleventh International Conference on Learning Representations*.
- Zhao, S.; Zhao, L.; Zhang, Z.; Zhou, E.; and Metaxas, D. 2022. Global matching with overlapping attention for optical flow estimation. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, 17592–17601.