

# A Study of Finetuning Video Transformers for Multi-view Geometry Tasks

Huimin Wu<sup>1</sup>, Kwang-Ting Cheng<sup>1</sup>, Stephen Lin<sup>2</sup>, Zhirong Wu<sup>2</sup>

<sup>1</sup>The Hong Kong University of Science and Technology

<sup>2</sup>Microsoft Research Asia

hwubl@connect.ust.hk, timcheng@ust.hk, stevelin@microsoft.com, wuzhiron@microsoft.com

## Abstract

This paper presents an investigation of vision transformer learning for multi-view geometry tasks, such as optical flow estimation, by fine-tuning video foundation models. Unlike previous methods that involve custom architectural designs and task-specific pretraining, our research finds that general-purpose models pretrained on videos can be readily transferred to multi-view problems with minimal adaptation. The core insight is that general-purpose attention between patches learns temporal and spatial information for geometric reasoning. We demonstrate that appending a linear decoder to the Transformer backbone produces satisfactory results, and iterative refinement can further elevate performance to state-of-the-art levels. This conceptually simple approach achieves top cross-dataset generalization results for optical flow estimation with end-point error (EPE) of 0.69, 1.78, and 3.15 on the Sintel clean, Sintel final, and KITTI datasets, respectively. Our method additionally establishes a new record on the online test benchmark with EPE values of 0.79, 1.88, and F1 value of 3.79. Applications to 3D depth estimation and stereo matching also show strong performance, illustrating the versatility of video-pretrained models in addressing geometric vision tasks.

## Introduction

There exists considerable interest in building general-purpose vision foundation models (He et al. 2022; Alayrac et al. 2022; Wang et al. 2023c) capable of addressing a range of tasks from low-level reasoning to high-level understanding. However, their performance on the downstream tasks remains inferior to methods with elaborate architectural designs and task-specific pretraining. Further research is needed to fully realize the potential of vision foundation models for various applications.

For multi-view geometry tasks such as optical flow estimation, complicated pipelines are generally employed to enhance performance. For example, mainstream optical flow estimators utilize a convolutional stem to extract per-frame features, a transformer encoder to extract a cost volume as an explicit motion representation, and a recurrent neural network to conduct iterative refinement (Teed and Deng 2020; Jiang et al. 2021a; Huang et al. 2022; Shi et al. 2023). The

state-of-the-art method Flowformer++ (Shi et al. 2023) additionally uses a domain-specific pretraining scheme. Towards simplifying the pipeline, CroCo (Weinzaepfel et al. 2022) and its extension (Weinzaepfel et al. 2023) use transformers as well as an adaptation module for dense prediction (Ranftl, Bochkovskiy, and Koltun 2021). However, this does not provide a general solution because of its reliance on a specialized pretraining task, an in-house dataset obtained via meticulously controlled data collection, and a much larger downstream dataset than used in common practice.

In this work, we investigate a *simpler and more general* solution for multi-view geometry tasks. Our approach is guided by the intuition that cross-frame self-attention in video-pretrained vision transformers can capture matching information essential for these multi-view problems. Based on this idea, we adapt general-purpose ViTs pre-trained on videos (Feichtenhofer et al. 2022; Wang et al. 2022) to multi-view geometry tasks with minimal adjustment. After modifying pre-trained 3D ViTs to accommodate two-frame tasks, we found that appending a linear head is sufficient to obtain high performance, and iteratively refining the results further enhances prediction quality.

Our model, called **GeoViT**, was evaluated on the tasks of optical flow estimation, stereo matching and 3D depth estimation. For optical flow estimation, it sets a new standard on the challenging Sintel (clean and final splits) and KITTI datasets, reducing EPE by 21%, 9.6%, and F1 by 11.2%, respectively. For 3D depth estimation, GeoViT achieves state-of-art results on the RGBD\_SLAM, SUN3D and Scenes11 datasets across multiple evaluation metrics. For stereo matching, our model outperforms competing methods on the ETH3D dataset for two of the three metrics.

This work makes an important contribution to the research community by demonstrating that the robust feature representations learned during video pre-training capture temporal and spatial information useful across various geometric tasks. It presents a simple yet highly effective way to harness this knowledge without custom pretraining strategies or special data collection. We believe that the connection between video understanding and geometric computer vision has been underleveraged, and we hope that our work will inspire further geometric-based research that capitalizes on the rich representations learned by video-based models.

## Related Works

**Foundation models.** The rise of self-supervised learning has enabled vision models (Wu et al. 2018; He et al. 2022) to capture rich semantics from vast raw data without human labels. These pretrained models, when finetuned to individual downstream tasks (Li et al. 2022b; Lin et al. 2023; Xu et al. 2022b), have consistently shown improved generalization and have sparked interest in building general-purpose foundation models to solve all computer vision applications within a single model.

Large-scale language models (e.g., GPT (Brown 2020)) and multi-modal models (e.g., CLIP (Radford et al. 2021)) demonstrate strong performance using language as supervision for self-supervised learning. More importantly, language as a natural interface between applications allows the foundation model to perform tasks in a zero-shot or few-shot manner. Recent vision foundation models (Wang et al. 2022, 2023c) follow this trend by integrating language and formulating heterogeneous supervisions. Despite the flexibility of the language interface, vision foundation models often need to be finetuned in order to reach optimal performance.

The promise of foundation models lies in their generality and scalability. Nevertheless, they still fall short in performance compared with domain-specific methods. Unleashing the full potential of foundation models for end tasks has become a problem of great interest.

**Optical flow estimation.** Optical flow estimation predicts per-pixel displacement from the source to the target image. Its widespread use across various vision-related tasks makes it a fundamental element of computer vision. For example, flow estimates can offer additional supervisory cues into tasks ranging from action recognition (Simonyan and Zisserman 2014) and tracking (Zheng et al. 2021) to video segmentation (Cheng et al. 2017) and self-supervised learning (Huang et al. 2023).

Methods for predicting optical flow commonly build a cost volume (Dosovitskiy et al. 2015) to characterize the dense correlation between pixels. Feature enhancement (Sun et al. 2022; Sui et al. 2022; Huang et al. 2022; Shi et al. 2023) for the cost volume further improves the prediction accuracy. Iterative refinement (Teed and Deng 2020; Huang et al. 2022) is another crucial technique for addressing flows with large displacement, where the prediction at the current timestep is fed back to the model as inputs for refinement.

A line of work (Xu et al. 2022a, 2023b; Weinzaepfel et al. 2022, 2023) attempts to simplify the pipeline using generic transformers for optical flow estimation. As transformers lack inductive biases, the model needs to be pretrained with custom pretext tasks on particular types of datasets. However, these methods have yet to reach the performance levels of heavily engineered approaches.

**Stereo matching.** Predicting disparity for stereo matching involves estimating the horizontal pixel displacement between the left and right stereo images. Given the similarities with flow prediction, recent advancements (Lipson, Teed, and Deng 2021) have adapted the flow network RAFT (Teed and Deng 2020) for stereo applications, where a 3D correlation volume is constructed instead of a 4D cost volume due to epipolar constraints. Furthermore, improvements

have been made to enhance the correlation volume’s robustness to non-ideal rectification (Li et al. 2022a) and its representative ability by introducing non-local information (Xu et al. 2023a). Additionally, efforts have been made to tame transformers to improve stereo matching (Guo et al. 2022; Li et al. 2021).

**Depth estimation.** Our work is applied to unrectified multi-view depth estimation, which predicts depth from images captured without rectifying them to a common plane. This setting allows for more natural and diverse camera configurations. A commonly used method is plane-sweep stereo (Collins 1996; Im et al. 2019), which constructs a 3D cost volume where the third dimension represents potential depth hypotheses. Subsequent improvements have enhanced the cost volume design by improving the multi-scale representation of depth information (Yang et al. 2020; Gu et al. 2020) and by using transformers to better capture global context (Ding et al. 2022; Xu et al. 2023b).

## Approach

We explore the use of large-scale video foundation models for geometric vision tasks, including optical flow, stereo correspondence, and depth estimation. Our intuition is that the cross-frame attention mechanisms learned by transformer-based architectures encode not only high-level semantic relationships but also rich correspondence cues that can be leveraged to recover fine-grained, low-level matches.

To this end, we propose GeoViT, a unified framework that takes an image pair as input and predicts the geometric quantity relating the two views. This formulation enables optical flow, stereo matching, and depth estimation to be addressed within a single architecture. GeoViT is built around a pretrained spatiotemporal Vision Transformer (ViT) encoder, which is demonstrated to provide strong geometric representations. By further integrating within an iterative refinement loop, the framework captures fine-grained geometric relationships while incrementally improving estimation accuracy through feedback and adjustment.

In this section, we will first briefly describe video foundation models based on vision transformers. Following this, we will discuss our simple yet effective adaptation scheme designed for downstream geometric tasks, emphasizing that the iterative refinement strategy empowers GeoViT to achieve superior results.

## Pretrained Video Foundation Models

Our work transfers the encoder in video foundation models to multi-view geometry tasks. Any video foundation model whose encoder follows a transformer architecture can be used. In our study, the models include MVD (Wang et al. 2023b), InternVideo (Wang et al. 2022), UMT (Li et al. 2023), MAE (He et al. 2022), and MAE<sub>st</sub> (Feichtenhofer et al. 2022). To process the video data, the transformer splits the spatio-temporal data into 3D patches, adds spatial and temporal positional encodings, and then feeds the visual tokens into the self-attention blocks.

Video foundation models are primarily benchmarked on semantic reasoning tasks, such as video action recognition (Feichtenhofer et al. 2022; Tong et al. 2022) and video

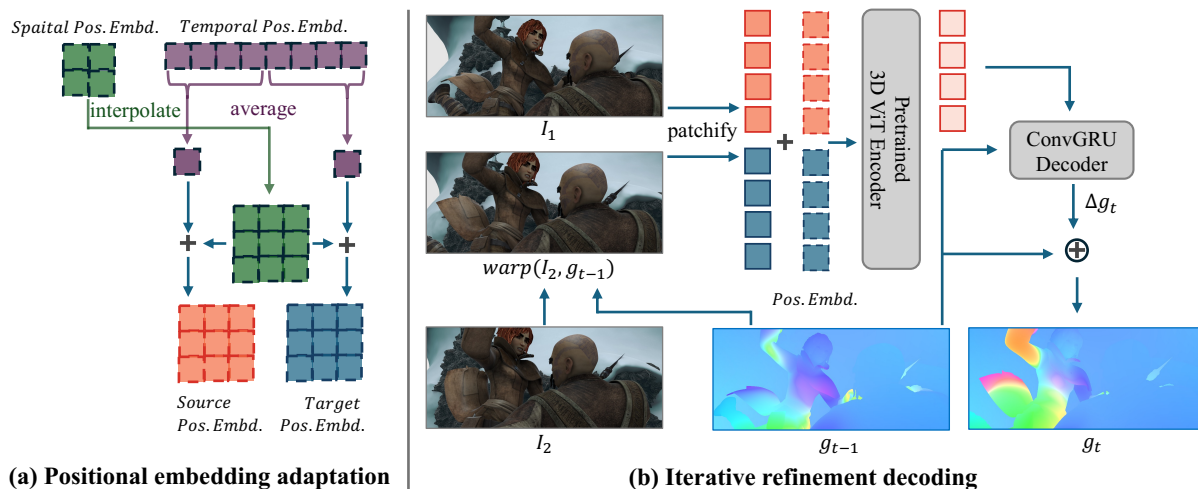


Figure 1: Overview of GeoViT. Part (a) presents adaptation of positional embeddings in pretrained 3D ViTs. The pretrained spatial Pos. Embd. (in green) are interpolated to match the desired input size. The pretrained temporal Pos. Embd. (in purple), which accounts for 8 frames, is split into two halves. The average of the first half and second half corresponds to the temporal embedding of the source image (in orange) and target image (in blue), respectively. Part (b) exhibits our iterative refinement decoding pipeline, using optical flow for illustration. The input target image is dynamically warped based on the last-step prediction  $g_{t-1}$  so that the input pair corresponds to the ground truth residual for this step. Then the source and (warped) target images are patchified, added with adapted positional embeddings, and fed to the pretrained 3D ViT for feature extraction. The decoder accepts source image features (in light orange) and the last-step prediction ( $g_{t-1}$ ) and produces its correction  $\Delta g_t$ . Adding  $g_{t-1}$  and its correction gives the current-step prediction  $g_t$ .

action detection (Wang et al. 2023a). Since the video encoder is trained on multi-frame input while most geometric tasks process a pair of frames, the patchification and position encodings need to be adjusted.

To adapt pretrained 3D ViTs for two-frame tasks, we first interpolate the 2D spatial positional encodings to match the desired input size. The temporal dimension in the pretrained 3D patch embeddings is usually 8 frames. We split the temporal position encodings into two halves and calculate the average of each half along the temporal dimension. For the patch embedding, we sum the 3D patch embedding along the temporal dimension and convert it basically to a 2D patch embedding. In this way, we adjust the position encodings and patch embeddings for the pair-wise geometry tasks. The final token representation is the addition of 2D patch embedding, 2D positional encoding, and temporal encoding. This approach ensures that the model effectively captures the necessary spatial and temporal information while maintaining compatibility with the pretrained model.

### Simple Linear Decoding

A key driver behind our approach is the inherent ability of transformers to model dense correlations, making it particularly well-suited for per-pixel predictions in downstream geometric applications. Moreover, the self-attention mechanism features a full-context receptive field, which significantly expands the range of contextual information it can capture. This design is in line with recent advancements that

focus on enhancing feature interactions across both spatial dimensions (Jiang et al. 2021a; Sun et al. 2022) and motion dimensions (Xu et al. 2022a, 2023b; Shi et al. 2023).

As an initial step, we append a linear layer to each output patch representation, and directly regress the geometrical properties. With this minimal modification to the pre-trained encoder, strong performance can already be achieved as shown in Table 2b. For optical flow estimation, our method attains an End Point Error (EPE) of 2.0 on the Sintel (final) dataset, surpassing existing state-of-the-art methods like SAMFlow (2.11). This demonstrates that the transformer encoder learned from vast video data effectively learns patch-wise correlations useful for geometrical reasoning.

### Iterative Refinement Decoding

To further enhance performance, we propose to incorporate an iterative refinement mechanism inspired by RAFT (Teed and Deng 2020). In RAFT, the optical flow prediction is decomposed and serialized to a sequence of residuals where a recurrent neural network (RNN) is used to predict the residual at each timestep. The summation of the residuals corresponds to the final prediction. The input feature at each timestep to the RNN is dynamically constructed to facilitate residual prediction. Specifically, for each source pixel, it queries the cost volume based on the current optical flow estimate. While it is feasible to migrate this refinement approach to vision transformer backbones, this would face two issues: 1) the cost volume construction is complex in de-

sign and overlaps the modeling role of self-attentions, 2) the lookup operation at each timestep is local and suboptimal (Huang et al. 2022; Shi et al. 2023).

To overcome the aforementioned limitations, we propose to remove reliance on the cost volume for a simple yet effective refinement framework. Formally, given an image pair  $I_1, I_2$ , the residual geometric property  $\Delta g_t$  at each iteration  $t$  is predicted as follows:

$$\Delta g_t = F_{\text{dec}}(F_{\text{enc}}(I_1, \text{warp}(I_2, g_{t-1})), g_{t-1}), \quad (1)$$

$$g_t = g_{t-1} + \Delta g_t, \quad (2)$$

where  $F_{\text{enc}}$  denotes a spatiotemporal ViT that takes (warped) image pairs as input and returns features corresponding to source images. The source image features together with the current prediction are fed to the decoder  $F_{\text{dec}}$  to predict the residuals. The decoder is instantiated as a ConvGRU unit following RAFT. The warping operation takes an image as input and outputs another image according to the geometrical property  $g_{t-1}$ . The warping operation for optical flow estimation and stereo matching is straightforward. For 3D depth estimation, we first convert the depth representation to pixel displacement using camera parameters, and convert the prediction results back to depth. The predicted residual is then aggregated with the prediction of the previous step.

The generated sequence of predictions  $g_1, g_2, \dots, g_T$  are regressed to minimize their L1 distance with the ground truth, weighted by exponentially increasing factors. Formally, the loss function for a pair of images is defined as:

$$L(I_1, I_2, y) = \sum_{t=1}^T \gamma^{T-t} d_{L1}(g_t, y), \quad (3)$$

where  $y$  represents the ground truth map and  $d_{L1}$  denotes L1 distance.  $\gamma$  is the weighting factor that assigns a higher weight to later predictions, which is set to 0.9 by default.  $T$  denotes the number of iterations.

The overall iterative refinement uses only a generic pretrained encoder plus a lightweight recurrent neural network. It does not rely on customized representations such as a cost volume. Looping over the encoder allows the refinement to take greater advantage of large-scale pretrained models.

A concern about the approach is that it adds significant computational complexity. However, we find this not to be a serious issue because our largest model can be trained on just 8 V100 GPUs for at most 2 weeks using techniques such as gradient checkpointing. This paper aims to develop a conceptually simple approach. Controlling factors such as data size and model size for fair comparisons is challenging.

## Experiments

In this section, we conduct comprehensive experiments qualitatively and quantitatively for optical flow estimation, stereo matching and 3D depth estimation. Ablation studies are performed on the optical flow estimation task. Most of the results are submitted to evaluation servers to obtain official results. In the following, we describe each task.

Training Data	Method	Sintel (train) KITTI-15 (train)			
		Clean	Final	F1-epe	F1-all
A	Perceiver IO (Jaegle et al. 2021)	1.81	2.42	4.98	-
	PWC-Net (Sun et al. 2018)	2.17	2.91	5.76	-
	RAFT (Teed and Deng 2020)	1.95	2.57	4.23	-
C + T	HD3 (Yin, Darrell, and Yu 2019)	3.84	8.77	13.17	24.0
	LiteFlowNet (Hui, Tang, and Loy 2018)	2.48	4.04	10.39	28.5
	PWC-Net (Sun et al. 2018)	2.55	3.93	10.35	33.7
	LiteFlowNet2 (Hui, Tang, and Loy 2020)	2.24	3.78	8.97	25.9
	S-Flow (Zhang et al. 2021)	1.30	2.59	4.60	15.9
	RAFT (Teed and Deng 2020)	1.43	2.71	5.04	17.4
	FM-RAFT (Jiang et al. 2021b)	1.29	2.95	6.80	19.3
	GMA (Jiang et al. 2021a)	1.30	2.74	4.69	17.1
	GMFlow (Xu et al. 2022a)	1.08	2.48	7.77	23.40
	GMFlowNet (Zhao et al. 2022)	1.14	2.71	4.24	15.4
	CRAFT (Sui et al. 2022)	1.27	2.79	4.88	17.5
	SKFlow (Sun et al. 2022)	1.22	2.46	4.47	15.5
	FlowFormer (Huang et al. 2022)	0.94	2.33	4.09 <sup>†</sup>	14.72 <sup>†</sup>
	FlowFormer++ (Shi et al. 2023)	0.90	2.30	3.93 <sup>†</sup>	14.13 <sup>†</sup>
	SAMFlow (Zhou et al. 2024)	0.87	2.11	3.44	12.28
	DPFlow (Morimitsu et al. 2025)	1.02	2.26	3.37	<b>11.1</b>
	GeoViT	<b>0.69<sup>†</sup></b>	<b>1.78<sup>†</sup></b>	<b>3.15<sup>†</sup></b>	11.45 <sup>†</sup>

Table 1: Experiments on Sintel (Butler et al. 2012) and KITTI (Geiger et al. 2013) datasets. ‘A’ denotes training on the autoflow dataset. ‘C + T’ denotes training sequentially on the FlyingChairs and FlyingThings datasets. <sup>†</sup> represents evaluation with tiling technique (Jaegle et al. 2021).

## Optical Flow

**Architecture.** While our framework can operate with any pretrained video transformer, we use MAE<sub>st</sub> ViT-Large (Feichtenhofer et al. 2022) as the default. The decoder is implemented as a ConvGRU with six convolutional layers: three for horizontal feature extraction and three for vertical feature extraction. It takes two inputs: a source image representation and flow features from the previous step (extracted by a motion encoder). The flow map is generated by passing the updated hidden state to a flow prediction head (two convolutional layers). To achieve full-resolution predictions, a mask layer, consisting of two convolutional layers, is also trained to indicate how to take the weighted sum of low-resolution neighbors to calculate high-resolution flow.

**Metrics.** We report average endpoint error (EPE) and F1. The EPE evaluates the average distance in pixels between the predicted and ground truth flow, using the  $l_2$  distance. Meanwhile, the F1 score measures the percentage of outlier predictions, identified as those with an error exceeding 3 pixels or over 5% of the magnitude of the ground truth flow.

**Cross Dataset Generalization** We use the one-cycle learning rate scheduler with cosine annealing. First, we train the proposed model on Chairs for 40K steps with a batch size of 8 and image size of  $368 \times 496$ , followed by 400K iterations of training on Things with a batch size of 8 on  $384 \times 768$  image crops. The highest learning rates on these two datasets are  $1 \times 10^{-4}$  and  $2.5 \times 10^{-5}$ , respectively. Our augmentation follows Flowformer (Shi et al. 2023). This stage is denoted as ‘‘C+T’’, and then the model is evaluated

Initialization	Sintel (train)		KITTI-15 (train)	
	Clean	Final	F1-epe	F1-all
Rand. Init.	13.49	13.49	36.74	85.35
MVD (Wang et al. 2023b)	0.84	1.91	3.61	14.28
InternVideo (Wang et al. 2022)	0.70	1.87	3.22	12.27
UMT (Li et al. 2023)	0.78	1.89	3.73	14.13
MAE (He et al. 2022)	0.78	1.99	4.18	14.85
MAE.st (Feichtenhofer et al. 2022)	<b>0.69</b>	<b>1.78</b>	<b>3.15</b>	<b>11.45</b>

(a) Pretraining schemes. Spatiotemporal masked autoencoder (Feichtenhofer et al. 2022) as flow estimation initialization outperforms its semantics-rich extensions (Wang et al. 2023b, 2022; Li et al. 2023) and motion-lacking predecessor (He et al. 2022).

Model	Sintel (train)		KITTI-15 (train)	
	Clean	Final	F1-epe	F1-all
ViT-Small	0.99	2.32	6.33	19.81
ViT-Base	0.87	2.22	4.72	15.81
ViT-Large	<b>0.75</b>	<b>1.82</b>	<b>3.70</b>	<b>13.78</b>

(c) Model size. A larger model leads to better results. This table uses reproduced pre-training parameters with fewer pre-training steps (800 epochs instead of the default 1600 epochs).

# Training iterations	Sintel (train)		KITTI-15 (train)	
	Clean	Final	F1-epe	F1-all
1(linear)	0.91	2.00	4.93	20.47
1	0.95	2.04	4.78	19.85
2	0.80	1.92	3.68	14.42
6	<b>0.69</b>	<b>1.78</b>	3.15	11.45
12	0.70	<b>1.78</b>	<b>3.13</b>	<b>11.29</b>

(b) The effect of training iteration number. The performance of training with more iterations grows until reaching 6 iterations.

Method	Features	Refinement	Sintel (train)	
			Clean	Final
RAFT	CNN	Queried cost volume	1.43	2.71
RAFT (ViT context)	Video ViT	Queried cost volume	0.85	1.99
GeoViT	Video ViT	Image warping	<b>0.69</b>	<b>1.78</b>

(d) The effect of warping images. Warping images for refinement performs better than using cost volumes.

Table 2: Ablation experiments for optical flow estimation on the Sintel and KITTI datasets. The best configuration is to initialize a ViT-Large encoder with MAE.st (Feichtenhofer et al. 2022) and refine for six iterations.

on the training split of Sintel and KITTI. A standard tiling scheme (Jaegle et al. 2021; Huang et al. 2022; Shi et al. 2023) with a stride step of 224 is used for evaluation.

Table 1 presents cross-dataset generalization results with comparisons to prior art. Our approach surpasses existing alternatives by a large margin on all the datasets. Compared with the current state-of-the-art SAMFlow which relies on additional large-scale segmentation labels, GeoViT reduces the error by 20.7%, 15.6 % respectively on the training set of Sintel (clean) and Sintel (final). On KITTI, the EPE and F1 is decreased by 8.4% and 6.8%.

**Ablation Study** Under the same settings used for cross-dataset generalization, we examine several design choices.

*Pretrained video foundation models.* We compare various choices of pretrained video foundation models, including MAE.st (Feichtenhofer et al. 2022), spatial MAE (He et al. 2022), MVD (Wang et al. 2023b), InternVideo (Wang et al. 2022) and UMT (Li et al. 2023). Table 2a summarizes the results. In comparison with other pretrained video models, MAE.st achieves the best performance. This is potentially because other works which aim to improve semantic capability by reconstructing learned features (Wang et al. 2023b), distilling from highly semantic features (Li et al. 2023) such as CLIP (Radford et al. 2021), or contrasting with language representations (Wang et al. 2022) do not necessarily translate well for low-level representations, leading to inferior performance. Improvement over the 2D MAE model pre-trained on images suggests that inter-frame pretraining on videos is vital for this two-frame task.

*Model Size.* Table 2c demonstrates the scalability of our adaptation. The downstream flow results show a consistent improvement when increasing the pre-trained model size.

This observation underscores the potential for further enhancement of our approach with even larger model sizes.

*Iterative Refinement Steps.* The results in Table 2b show that flow estimation performance generally improves with an increased number of refinement steps. The performance of the 12-iteration model is slightly better than that of the 6-iteration model, but the difference is minor. Therefore, for efficiency, we choose 6 as the default iteration number.

*Decoder w.r.t. RAFT.* Table 2d ablates the role of warping images as opposed to using cost volumes. We maintain RAFT’s cost volume features but use the video foundation model as the context encoder. Using the video foundation model as the context encoder improves RAFT on Sintel and performs comparably on KITTI. Our GeoViT outperforms it by a large margin, validating the role of warping images.

**Sintel benchmark** We evaluate model performance on the Sintel benchmark server. Following (Xu et al. 2023b), we finetune the model on a joint dataset composed of Things, Sintel, KITTI-2015, and HD1K for 200K iterations, with a batch size of 8, image size of  $288 \times 960$ , and highest learning rate of  $2.5 \times 10^{-5}$ . We proceeded with a Sintel finetuning stage for 5K iterations with a batch size of 8, image size of  $416 \times 1024$ , and highest learning rate of  $1 \times 10^{-5}$ .

As shown in Table 3, GeoViT sets a new performance standard on the Sintel benchmark. Our method outperforms the prior state-of-the-art SAM-Flow by 21% on the clean split and by 9.6% on the final split, despite being more generic and using less human supervision.

**KITTI benchmark** We transfer the Sintel stage-trained model (with the highest learning rate of  $2.5e-5$ ) and submit our result to the KITTI evaluation server. As shown in Ta-

Training Data	Method	Sintel (test)	
		Clean	Final
C + T + S +K + H	LiteFlowNet2 (Hui, Tang, and Loy 2020)	3.48	4.69
	PWC-Net+ (Sun et al. 2019)	3.45	4.60
	VCN (Yang and Ramanan 2019)	2.81	4.40
	MaskFlowNet (Zhao et al. 2020)	2.52	4.17
	S-Flow (Zhang et al. 2021)	1.50	2.67
	RAFT (Teed and Deng 2020)	1.94	3.18
	RAFT* (Teed and Deng 2020)	1.61	2.86
	FM-RAFT (Jiang et al. 2021a)	1.72	3.60
	GMA (Jiang et al. 2021a)	1.40	2.88
	GMA* (Jiang et al. 2021a)	1.39	2.47
	GMFlow (Xu et al. 2022a)	1.74	2.90
	GMFlowNet (Zhao et al. 2022)	1.39	2.65
	CRAFT (Sui et al. 2022)	1.45	2.42
	SKFlow* (Sun et al. 2022)	1.28	2.23
	FlowFormer (Huang et al. 2022)	1.16	2.09
	FlowFormer++ (Shi et al. 2023)	1.07	1.94
	GMFlow+(Weinzaepfel et al. 2023)	1.03	2.12
	SAM-Flow (Zhou et al. 2024)	1.00	2.08
	DPFlow (Morimitsu et al. 2025)	1.04	1.97
	GeoViT	<b>0.79<sup>†</sup></b>	<b>1.88<sup>†</sup></b>

Table 3: Experiments on Sintel (Butler et al. 2012) benchmark. ‘C + K + S + K + H’ denotes finetuning on the combined Sintel, KITTI, and HD1K training sets after ‘C + T’ training. \* denotes methods with warm-start proposed in RAFT (Teed and Deng 2020), where the flow is initialized with previous image pair flow estimation. <sup>†</sup> represents evaluation with tiling technique (Jaegle et al. 2021). Our approach ranks 1st on the Sintel benchmark.

ble 4, GeoViT achieves an F1-all score of 3.79. This significantly improves upon SOTAs, reducing the error by 11.2%.

**Qualitative Comparison** We compare the visualized flow maps on Sintel (clean) dataset with strong competitors: FlowFormer++, GMFlow+, and SAM-Flow. As demonstrated in Figure 2, our predictions show significant improvements. First, our approach provides more detailed estimates; for instance, in case #1, it accurately captures the different motion patterns of the arm and back. Second, GeoViT exhibits a higher recall for small objects, correctly predicting all bird motions in case #2. Lastly, it generates sharper motion boundaries, as seen on the left side of case #2.

## Stereo Matching

**Datasets and Evaluation.** We validate GeoViT on stereo matching with the ETH3D Stereo (Schops et al. 2017) dataset. Following (Xu et al. 2023b), we conduct two-stage training on the synthetic Scene Flow training set (Mayer et al. 2016) and a combined dataset: Scene Flow (Mayer et al. 2016), Tartan Air (Wang et al. 2020), Sintel Stereo (Butler et al. 2012), CREStereo Dataset (Li et al. 2022a), InStereo2K (Bao et al. 2020) and ETH3D (Schops et al. 2017). We report the performance of GeoViT using the online ETH benchmark (Schops et al. 2017).

**Training details.** In the first stage, GeoViT is initialized with the optical flow parameters trained on ‘‘C+T’’. This ap-

Training Data	Method	KITTI-15 (test)
		F1-all
C + T + S +K + H	LiteFlowNet2 (Hui, Tang, and Loy 2020)	7.74
	PWC-Net+ (Sun et al. 2019)	7.72
	VCN (Yang and Ramanan 2019)	6.30
	MaskFlowNet (Zhao et al. 2020)	6.10
	S-Flow (Zhang et al. 2021)	4.64
	RAFT (Teed and Deng 2020)	5.10
	RAFT* (Teed and Deng 2020)	5.10
	FM-RAFT (Jiang et al. 2021a)	6.17
	GMA (Jiang et al. 2021a)	5.15
	GMA* (Jiang et al. 2021a)	5.15
	GMFlow (Xu et al. 2022a)	9.32
	GMFlowNet (Zhao et al. 2022)	4.79
	CRAFT (Sui et al. 2022)	4.79
	SKFlow* (Sun et al. 2022)	4.84
	FlowFormer (Huang et al. 2022)	4.68 <sup>†</sup>
	FlowFormer++ (Shi et al. 2023)	4.52 <sup>†</sup>
	GMFlow+(Weinzaepfel et al. 2023)	4.27
	SAM-Flow (Zhou et al. 2024)	4.49
	DPFlow (Morimitsu et al. 2025)	<b>3.56</b>
	GeoViT	<b>3.79<sup>†</sup></b>

Table 4: Experiments on KITTI (Geiger et al. 2013) benchmark. ‘C + K + S + K + H’ denotes finetuning on the combined Sintel, KITTI, and HD1K training sets after ‘C + T’ training. <sup>†</sup> represents tiling technique (Jaegle et al. 2021).

Model	bad 1.0	bad 2.0	bad 4.0
GANet (Zhang et al. 2019)	6.56	1.10	0.54
AANet (Xu and Zhang 2020)	5.01	1.66	0.75
CFNet (Shen, Dai, and Rao 2021)	3.31	0.77	0.31
RAFT-Stereo (Lipson, Teed, and Deng 2021)	2.44	0.44	0.15
CREStereo (Li et al. 2022a)	0.98	0.22	0.10
GMStereo (Xu et al. 2023b)	1.83	0.25	0.08
MonSter (Cheng et al. 2025)	<b>0.72</b>	0.42	0.20
GeoViT	1.16	<b>0.19</b>	<b>0.03</b>

Table 5: Stereo performance on ETH3D stereo test set. Our method achieves the best results on two of the three metrics.

proach treats disparity prediction as a one-dimensional optical flow by using the first dimension of the predicted flow values as the output and concatenating zeros to convert the disparity values into two-dimensional values for refinement in the decoder. Our model is trained on the Scene Flow dataset for 100k iterations, with a batch size of 8 and an image size of 384×768. The maximum learning rate is set to 5e-5. In the second stage, we finetune our model for an additional 100k iterations, with a batch size of 8 and an image size of 416×640, with a maximum learning rate of 5e-6.

**Metrics.** The evaluation uses a widely recognized metric named bad  $np$ , which measures the percentage of pixels where the absolute error between the predicted disparity and the ground-truth disparity is greater than  $np$  pixels. We report bad 1.0, bad 2.0, and bad 4.0.

**Comparison Results.** Comparison results are presented in Table 5. Our model achieves the highest performance in terms of ‘‘bad 2.0’’ and ‘‘bad 4.0’’. For ‘‘bad 1.0’’, our re-

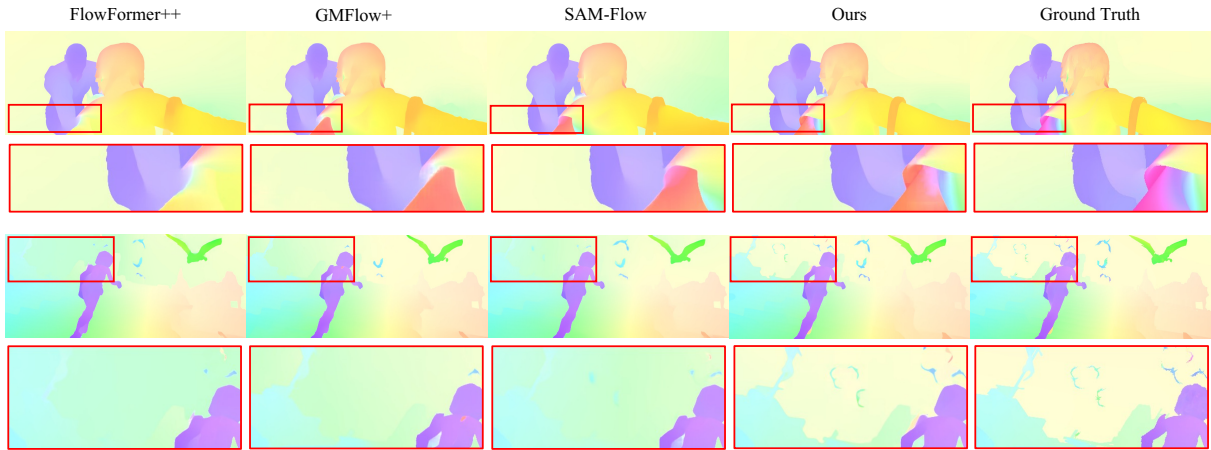


Figure 2: Visualized prediction comparison on Sintel (clean) dataset. Our approach is more accurate with more fine-grained estimates (human shoulder region in case #1), a higher recall of small objects (bird region in case #2), and a crisper motion boundary (case #2). The highlighted region is zoomed in for better visual comparison.

sults are comparable to those of CREStereo (Li et al. 2022a) and significantly outperform the prior unified solution GM-Stereo (Xu et al. 2023b) (1.16 vs. 1.83).

## Depth estimation

**Datasets and Evaluation.** We perform unrectified stereo depth estimation using a combined dataset consisting of RGBD-SLAM (Sturm et al. 2012), SUN3D (Xiao, Owens, and Torralba 2013), and Scenes11 (Ummenhofer et al. 2017) for training, and we evaluate on their respective test datasets. Our model is trained for 100k iterations with a batch size of 8, an image size of  $448 \times 576$ , and the highest learning rate of  $2.5e-5$ . The iteration number is set to 1 for depth estimation. **Metrics.** In line with (Tang and Tan 2018; Im et al. 2019; Xu et al. 2023b), we employ four evaluation metrics: Absolute Relative difference (Abs Rel), Squared Relative difference (Sq Rel), Root Mean Squared Error (RMSE), and Root Mean Squared Error in log scale (RMSE log).

**Comparison Results.** As presented in Table 6, our method exhibits performance that is either comparable to or superior to that of prior methods. On the RGBD-SLAM, the metrics in which our method underperforms indicate negligible differences, with the largest error difference being merely 0.027 in terms of Sq Rel. Conversely, our model surpasses the previous best model in terms of RMSE, achieving an improvement of 0.048. On the SUN3D, our model achieves the best results on 3 of the 4 metrics. On Scenes11, our model achieves the best performance in 2 of the 4 metrics. Our visualized failure case analysis reveals that our approach tends to produce overly smooth depth predictions, leading to slightly reduced performance on complex scene datasets like RGBD-SLAM and Scenes11 compared to SUN3D.

## Conclusion

In this work, we present a comprehensive study on fine-tuning large-scale pretrained models for multi-view geometry tasks, including optical flow estimation, stereo match-

Dataset	Model	Abs Rel	Sq Rel	RMSE	RMSE log
RGBD-SLAM	DeMoN (Ummenhofer et al. 2017)	0.157	0.524	1.780	0.202
	DeepMVS (Huang et al. 2018)	0.294	0.430	0.868	0.351
	DPSNet (Im et al. 2019)	0.154	0.215	0.723	0.226
	IIB (Yifan et al. 2022)	<b>0.095</b>	-	0.550	-
	GMDepth (Xu et al. 2023b)	0.101	<b>0.177</b>	0.556	<b>0.167</b>
	GeoViT	0.106	0.204	<b>0.508</b>	0.171
SUN3D	DeMoN (Ummenhofer et al. 2017)	0.214	1.120	2.421	0.206
	DeepMVS (Huang et al. 2018)	0.282	0.435	0.944	0.363
	DPSNet (Im et al. 2019)	0.147	0.107	0.427	0.191
	IIB (Yifan et al. 2022)	0.099	-	<b>0.293</b>	-
	GMDepth (Xu et al. 2023b)	0.112	<b>0.068</b>	0.336	0.146
	GeoViT	<b>0.095</b>	<b>0.068</b>	0.552	<b>0.124</b>
Scenes11	DeMoN (Ummenhofer et al. 2017)	0.556	3.402	2.603	0.391
	DeepMVS (Huang et al. 2018)	0.210	0.373	0.891	0.270
	DPSNet (Im et al. 2019)	0.056	0.144	0.714	0.140
	IIB (Yifan et al. 2022)	0.056	-	0.523	-
	GMDepth (Xu et al. 2023b)	0.050	0.069	0.491	<b>0.106</b>
	GeoViT	0.118	<b>0.059</b>	<b>0.318</b>	0.146

Table 6: Depth performance on DeMoN test datasets. Our approach obtains better or comparable performance.

ing, and unrectified two-view depth estimation. The core idea is to transfer a general-purpose video foundation model to the domain of geometry estimation with minimal adaptation: by either appending a linear layer as the decoder or embedding it into an iterative refinement loop. Unlike traditional methods, which often rely on complex, task-specific designs and pretraining strategies, our approach is both conceptually straightforward and effective in addressing a wide range of geometric tasks. This work not only advances the state of the art in multi-view geometry estimation but also lays the groundwork for future research in leveraging video-pretrained models in this domain.

## References

- Alayrac, J.-B.; Donahue, J.; Luc, P.; Miech, A.; Barr, I.; Hasson, Y.; Lenc, K.; Mensch, A.; Millican, K.; Reynolds, M.; et al. 2022. Flamingo: a visual language model for few-shot learning. *NeurIPS*, 35.
- Bao, W.; Wang, W.; Xu, Y.; Guo, Y.; Hong, S.; and Zhang, X. 2020. Instereo2k: a large real dataset for stereo matching in indoor scenes. *Science China Information Sciences*, 63.
- Brown, T. B. 2020. Language models are few-shot learners. *arXiv preprint arXiv:2005.14165*.
- Butler, D. J.; Wulff, J.; Stanley, G. B.; and Black, M. J. 2012. A naturalistic open source movie for optical flow evaluation. In *ECCV*. Springer.
- Cheng, J.; Liu, L.; Xu, G.; Wang, X.; Zhang, Z.; Deng, Y.; Zang, J.; Chen, Y.; Cai, Z.; and Yang, X. 2025. MonSter: Marry Monodepth to Stereo Unleashes Power. In *CVPR*.
- Cheng, J.; Tsai, Y.-H.; Wang, S.; and Yang, M.-H. 2017. Segflow: Joint learning for video object segmentation and optical flow. In *ICCV*.
- Collins, R. T. 1996. A space-sweep approach to true multi-image matching. In *CVPR*. Ieee.
- Ding, Y.; Yuan, W.; Zhu, Q.; Zhang, H.; Liu, X.; Wang, Y.; and Liu, X. 2022. Transmvsnet: Global context-aware multi-view stereo network with transformers. In *CVPR*.
- 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 *ICCV*.
- Feichtenhofer, C.; Li, Y.; He, K.; et al. 2022. Masked autoencoders as spatiotemporal learners. *NeurIPS*, 35.
- Geiger, A.; Lenz, P.; Stiller, C.; and Urtasun, R. 2013. Vision meets robotics: The kitti dataset. *The International Journal of Robotics Research*, 32(11).
- Gu, X.; Fan, Z.; Zhu, S.; Dai, Z.; Tan, F.; and Tan, P. 2020. Cascade cost volume for high-resolution multi-view stereo and stereo matching. In *CVPR*.
- Guo, W.; Li, Z.; Yang, Y.; Wang, Z.; Taylor, R. H.; Unberath, M.; Yuille, A.; and Li, Y. 2022. Context-enhanced stereo transformer. In *ECCV*. Springer.
- He, K.; Chen, X.; Xie, S.; Li, Y.; Dollár, P.; and Girshick, R. 2022. Masked autoencoders are scalable vision learners. In *CVPR*.
- Huang, B.; Zhao, Z.; Zhang, G.; Qiao, Y.; and Wang, L. 2023. MgmAE: Motion guided masking for video masked autoencoding. In *ICCV*.
- Huang, P.-H.; Matzen, K.; Kopf, J.; Ahuja, N.; and Huang, J.-B. 2018. Deepmvs: Learning multi-view stereopsis. In *CVPR*.
- Huang, Z.; Shi, X.; Zhang, C.; Wang, Q.; Cheung, K. C.; Qin, H.; Dai, J.; and Li, H. 2022. Flowformer: A transformer architecture for optical flow. In *ECCV*. Springer.
- Hui, T.-W.; Tang, X.; and Loy, C. C. 2018. LiteflowNet: A lightweight convolutional neural network for optical flow estimation. In *CVPR*.
- Hui, T.-W.; Tang, X.; and Loy, C. C. 2020. A lightweight optical flow cnn—revisiting data fidelity and regularization. *TPAMI*, 43(8).
- Im, S.; Jeon, H.-G.; Lin, S.; and Kweon, I. S. 2019. DpNet: End-to-end deep plane sweep stereo. *arXiv preprint arXiv:1905.00538*.
- Jaegle, A.; Gimeno, F.; Brock, A.; Vinyals, O.; Zisserman, A.; and Carreira, J. 2021. Perceiver: General perception with iterative attention. In *ICML*. PMLR.
- Jiang, S.; Campbell, D.; Lu, Y.; Li, H.; and Hartley, R. 2021a. Learning to estimate hidden motions with global motion aggregation. In *ICCV*.
- Jiang, S.; Lu, Y.; Li, H.; and Hartley, R. 2021b. Learning optical flow from a few matches. In *CVPR*.
- Li, J.; Wang, P.; Xiong, P.; Cai, T.; Yan, Z.; Yang, L.; Liu, J.; Fan, H.; and Liu, S. 2022a. Practical stereo matching via cascaded recurrent network with adaptive correlation. In *CVPR*.
- Li, K.; Wang, Y.; Li, Y.; Wang, Y.; He, Y.; Wang, L.; and Qiao, Y. 2023. Unmasked teacher: Towards training-efficient video foundation models. *arXiv preprint arXiv:2303.16058*.
- Li, Y.; Mao, H.; Girshick, R.; and He, K. 2022b. Exploring plain vision transformer backbones for object detection. In *ECCV*. Springer.
- Li, Z.; Liu, X.; Drenkow, N.; Ding, A.; Creighton, F. X.; Taylor, R. H.; and Unberath, M. 2021. Revisiting stereo depth estimation from a sequence-to-sequence perspective with transformers. In *ICCV*.
- Lin, Y.; Yuan, Y.; Zhang, Z.; Li, C.; Zheng, N.; and Hu, H. 2023. Dtr does not need multi-scale or locality design. In *ICCV*.
- Lipson, L.; Teed, Z.; and Deng, J. 2021. Raft-stereo: Multi-level recurrent field transforms for stereo matching. In *3DV*. IEEE.
- 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 *CVPR*.
- Morimitsu, H.; Zhu, X.; Cesar Jr, R. M.; Ji, X.; and Yin, X.-C. 2025. DPFlow: Adaptive Optical Flow Estimation with a Dual-Pyramid Framework. In *arXiv preprint arXiv:2503.14880*.
- Radford, A.; Kim, J. W.; Hallacy, C.; Ramesh, A.; Goh, G.; Agarwal, S.; Sastry, G.; Askell, A.; Mishkin, P.; Clark, J.; et al. 2021. Learning transferable visual models from natural language supervision. In *ICML*. PMLR.
- Ranftl, R.; Bochkovskiy, A.; and Koltun, V. 2021. Vision transformers for dense prediction. In *ICCV*.
- Schops, T.; Schonberger, J. L.; Galliani, S.; Sattler, T.; Schindler, K.; Pollefeys, M.; and Geiger, A. 2017. A multi-view stereo benchmark with high-resolution images and multi-camera videos. In *CVPR*.
- Shen, Z.; Dai, Y.; and Rao, Z. 2021. Cfnet: Cascade and fused cost volume for robust stereo matching. In *CVPR*.

- 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 *CVPR*.
- Simonyan, K.; and Zisserman, A. 2014. Two-stream convolutional networks for action recognition in videos. *NeurIPS*, 27.
- Sturm, J.; Engelhard, N.; Endres, F.; Burgard, W.; and Cremers, D. 2012. A benchmark for the evaluation of RGB-D SLAM systems. In *IROS*. IEEE.
- 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 *CVPR*.
- Sun, D.; Yang, X.; Liu, M.-Y.; and Kautz, J. 2018. Pwcnet: Cnns for optical flow using pyramid, warping, and cost volume. In *CVPR*.
- 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. *TPAMI*, 42(6).
- Sun, S.; Chen, Y.; Zhu, Y.; Guo, G.; and Li, G. 2022. Skflow: Learning optical flow with super kernels. *NeurIPS*, 35.
- Tang, C.; and Tan, P. 2018. BA-Net: Dense Bundle Adjustment Networks. In *ICRL*.
- Teed, Z.; and Deng, J. 2020. Raft: Recurrent all-pairs field transforms for optical flow. In *ECCV*. Springer.
- Tong, Z.; Song, Y.; Wang, J.; and Wang, L. 2022. Videomae: Masked autoencoders are data-efficient learners for self-supervised video pre-training. *arXiv preprint arXiv:2203.12602*.
- Ummenhofer, B.; Zhou, H.; Uhrig, J.; Mayer, N.; Ilg, E.; Dosovitskiy, A.; and Brox, T. 2017. Demon: Depth and motion network for learning monocular stereo. In *CVPR*.
- Wang, L.; Huang, B.; Zhao, Z.; Tong, Z.; He, Y.; Wang, Y.; Wang, Y.; and Qiao, Y. 2023a. Videomae v2: Scaling video masked autoencoders with dual masking. In *CVPR*.
- Wang, R.; Chen, D.; Wu, Z.; Chen, Y.; Dai, X.; Liu, M.; Yuan, L.; and Jiang, Y.-G. 2023b. Masked video distillation: Rethinking masked feature modeling for self-supervised video representation learning. In *CVPR*.
- Wang, W.; Dai, J.; Chen, Z.; Huang, Z.; Li, Z.; Zhu, X.; Hu, X.; Lu, T.; Lu, L.; Li, H.; et al. 2023c. Internimage: Exploring large-scale vision foundation models with deformable convolutions. In *CVPR*.
- Wang, W.; Zhu, D.; Wang, X.; Hu, Y.; Qiu, Y.; Wang, C.; Hu, Y.; Kapoor, A.; and Scherer, S. 2020. Tartanair: A dataset to push the limits of visual slam. In *IROS*. IEEE.
- Wang, Y.; Li, K.; Li, Y.; He, Y.; Huang, B.; Zhao, Z.; Zhang, H.; Xu, J.; Liu, Y.; Wang, Z.; et al. 2022. Internvideo: General video foundation models via generative and discriminative learning. *arXiv preprint arXiv:2212.03191*.
- Weinzaepfel, P.; Leroy, V.; Lucas, T.; Brégier, R.; Cabon, Y.; Arora, V.; Antsfeld, L.; Chidlovskii, B.; Csurka, G.; and Revaud, J. 2022. CroCo: Self-Supervised Pre-training for 3D Vision Tasks by Cross-View Completion. *NeurIPS*, 35.
- Weinzaepfel, P.; Lucas, T.; Leroy, V.; Cabon, Y.; Arora, V.; Brégier, R.; Csurka, G.; Antsfeld, L.; Chidlovskii, B.; and Revaud, J. 2023. CroCo v2: Improved Cross-view Completion Pre-training for Stereo Matching and Optical Flow. In *ICCV*.
- Wu, Z.; Xiong, Y.; Yu, S. X.; and Lin, D. 2018. Unsupervised feature learning via non-parametric instance discrimination. In *CVPR*.
- Xiao, J.; Owens, A.; and Torralba, A. 2013. Sun3d: A database of big spaces reconstructed using sfm and object labels. In *ICCV*.
- Xu, G.; Wang, X.; Ding, X.; and Yang, X. 2023a. Iterative geometry encoding volume for stereo matching. In *CVPR*.
- Xu, H.; and Zhang, J. 2020. Aanet: Adaptive aggregation network for efficient stereo matching. In *CVPR*.
- Xu, H.; Zhang, J.; Cai, J.; Rezatofighi, H.; and Tao, D. 2022a. Gmflow: Learning optical flow via global matching. In *CVPR*.
- Xu, H.; Zhang, J.; Cai, J.; Rezatofighi, H.; Yu, F.; Tao, D.; and Geiger, A. 2023b. Unifying flow, stereo and depth estimation. *TPAMI*.
- Xu, Y.; Zhang, J.; Zhang, Q.; and Tao, D. 2022b. Vitpose: Simple vision transformer baselines for human pose estimation. *NeurIPS*, 35.
- Yang, G.; and Ramanan, D. 2019. Volumetric correspondence networks for optical flow. *NeurIPS*, 32.
- Yang, J.; Mao, W.; Alvarez, J. M.; and Liu, M. 2020. Cost volume pyramid based depth inference for multi-view stereo. In *CVPR*.
- Yifan, W.; Doersch, C.; Arandjelović, R.; Carreira, J.; and Zisserman, A. 2022. Input-level inductive biases for 3d reconstruction. In *CVPR*.
- Yin, Z.; Darrell, T.; and Yu, F. 2019. Hierarchical discrete distribution decomposition for match density estimation. In *CVPR*.
- Zhang, F.; Prisacariu, V.; Yang, R.; and Torr, P. H. 2019. Ganet: Guided aggregation net for end-to-end stereo matching. In *CVPR*.
- Zhang, F.; Woodford, O. J.; Prisacariu, V. A.; and Torr, P. H. 2021. Separable flow: Learning motion cost volumes for optical flow estimation. In *ICCV*.
- Zhao, S.; Sheng, Y.; Dong, Y.; Chang, E. I.; Xu, Y.; et al. 2020. Maskflownet: Asymmetric feature matching with learnable occlusion mask. In *CVPR*.
- Zhao, S.; Zhao, L.; Zhang, Z.; Zhou, E.; and Metaxas, D. 2022. Global matching with overlapping attention for optical flow estimation. In *CVPR*.
- Zheng, J.; Ma, C.; Peng, H.; and Yang, X. 2021. Learning to track objects from unlabeled videos. In *ICCV*.
- Zhou, S.; He, R.; Tan, W.; and Yan, B. 2024. Samflow: Eliminating any fragmentation in optical flow with segment anything model. In *AAAI*, volume 38.