

PriOr-Flow: Enhancing Primitive Panoramic Optical Flow with Orthogonal View

Longliang Liu¹, Miaojie Feng¹, Junda Cheng¹, Jijun Xiang¹, Xuan Zhu¹, Xin Yang^{2,1†}

¹ Huazhong University of Science and Technology ² Optics Valley Laboratory
{longliangl, fmj, jundacheng, jijunx, xuanzhu, xinyang2014}@hust.edu.cn

Abstract

Panoramic optical flow enables a comprehensive understanding of temporal dynamics across wide fields of view. However, severe distortions caused by sphere-to-plane projections, such as the equirectangular projection (ERP), significantly degrade the performance of conventional perspective-based optical flow methods, especially in polar regions. To address this challenge, we propose *PriOr-Flow*, a novel dual-branch framework that leverages the low-distortion nature of the orthogonal view to enhance optical flow estimation in these regions. Specifically, we introduce the *Dual-Cost Collaborative Lookup (DCCL)* operator, which jointly retrieves correlation information from both the primitive and orthogonal cost volumes, effectively mitigating distortion noise during cost volume construction. Furthermore, our *Ortho-Driven Distortion Compensation (ODDC)* module iteratively refines motion features of the primitive branch, further suppressing polar distortions. Extensive experiments demonstrate that *PriOr-Flow* is compatible with various perspective-based iterative optical flow methods and consistently achieves state-of-the-art performance on publicly available panoramic optical flow datasets, setting a new benchmark for wide-field motion estimation. The code is publicly available at: <https://github.com/longliangLiu/PriOr-Flow>.

1. Introduction

Optical flow estimates a dense motion vector field between consecutive video frames, capturing the apparent displacement of pixels on the image plane. As a fundamental problem in computer vision, it plays a pivotal role in dynamic scene understanding and supports a wide range of applications, including frame interpolation [20, 29, 45], video inpainting [16, 24, 44], autonomous driving [17, 49–52], and 3D reconstruction and synthesis [26, 32, 53, 57].

In the realm of perspective imagery, recent approaches [8, 9, 21, 22, 39, 41, 43, 55] have achieved remarkable suc-

[†]Corresponding author.

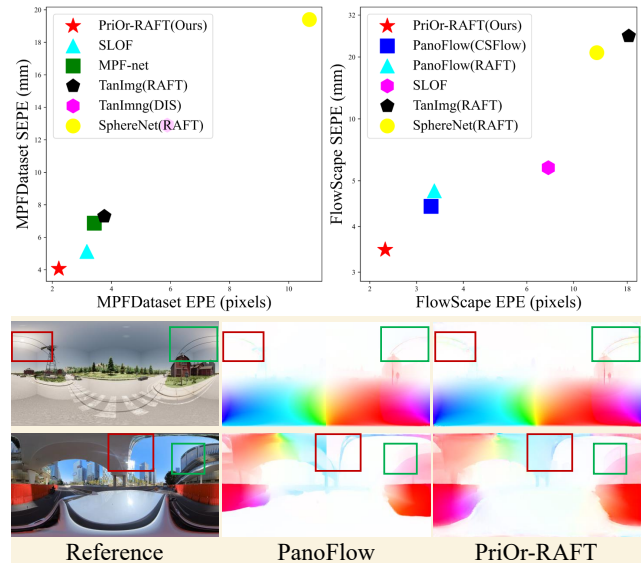


Figure 1. **Row 1:** Comparisons with state-of-the-art panoramic optical flow estimation methods on MPFDataset [28] and FlowScape [35]. **Row 2:** Visual comparison on the virtual dataset FlowScape [35]. **Row 3:** Visual comparison on the real dataset ODVista [42]. Our method performs well in the polar regions.

cess on well-established benchmarks. However, the increasing availability of panoramic cameras has created a pressing demand for panoramic optical flow estimation, which enables comprehensive motion understanding across wide fields of view. Among various projection formats, equirectangular projection (ERP) remains the most prevalent. Yet, by mapping a spherical image onto a 2D plane, ERP introduces severe distortions, particularly near the poles. As a result, applying conventional optical flow methods designed for perspective images directly to panoramic images leads to substantial performance degradation.

To address this challenge, several works have explored specialized architectures for panoramic optical flow, which can be broadly categorized into three groups: weight transformation-based, tangent plane-based, and ERP-based methods. *Weight transformation-based methods* [2, 5] adapt

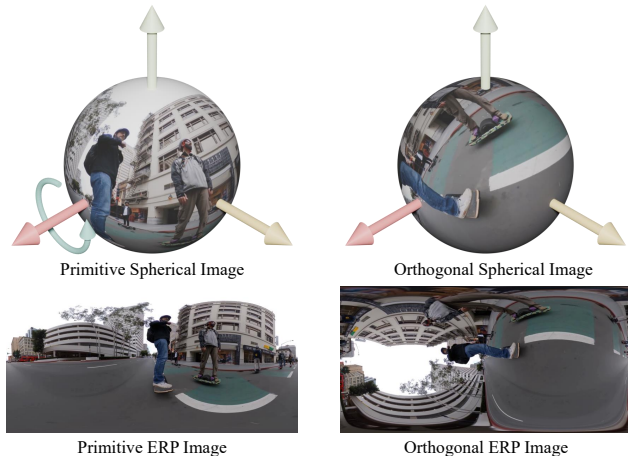


Figure 2. The orthogonal view is obtained by rotating the spherical image by 90° around the x-axis (or y-axis). The first row illustrates the spherical rotation process, while the second row shows the corresponding ERP projection images.

perspective-trained optical flow models to ERP images via weight transformations, eliminating the need for extensive panoramic optical flow datasets. However, these additional transformation layers introduce efficiency and portability concerns [6]. *Tangent plane-based methods* [48] mitigate distortions by projecting spherical images onto multiple tangent planes, allowing perspective optical flow models to operate in a locally undistorted manner. Nevertheless, this approach suffers from inter-plane discontinuities and cross-plane motion inconsistencies. *ERP-based methods* directly process ERP images and employ techniques such as deformable convolution [13] and representation learning [7] to compensate for projection distortions. However, ERP-induced distortions are most severe in polar regions, where existing approaches exhibit frequent and significant errors, as shown in Fig. 1. Despite these advancements, no existing method explicitly addresses the distortion effects in polar regions, leaving a critical gap in panoramic optical flow estimation.

To address the limitations of existing ERP-based methods, we introduce PriOr-Flow, a novel dual-branch panoramic optical flow estimation framework that uniquely exploits the low-distortion prior of the orthogonal view to explicitly compensate for severe distortions in polar regions [10, 11, 31, 47, 56]. As noted in [40, 46], the non-uniform sampling of ERP projection results in latitude-dependent distortions following a cosine pattern, with the most severe distortions occurring at the poles (Fig. 3(a)). This distortion significantly degrades optical flow accuracy in these regions. To mitigate this issue, we leverage the orthogonal view, obtained by rotating the panoramic image by 90° in spherical space (Fig. 2). Unlike the primitive ERP view, the orthogonal view exhibits an opposite distortion pattern, featuring minimal distortions in the polar regions (Fig. 3). During iterative refinement, PriOr-Flow’s primitive




Figure 3. Distortion maps of the primitive view and the orthogonal view, where brighter regions indicate higher distortion levels.

branch adaptively incorporates the low-distortion prior from the orthogonal branch, significantly enhancing optical flow estimation in polar regions.

Specifically, to mitigate distortion noise introduced during cost volume construction, we propose the Dual-Cost Collaborative Lookup (DCCL) operator. This operator leverages the estimated optical flow to jointly extract correlation information from both the primitive and orthogonal cost volumes. By incorporating both sources, DCCL prevents excessive noise in individual cost volume retrievals, thereby reducing optical flow estimation errors. To further harness the low-distortion prior of the orthogonal view, we introduce the Ortho-Driven Distortion Compensation (ODDC) module. This module explicitly guides the fusion of motion features across both branches based on their confidence estimates, ensuring that high-quality motion features from the orthogonal branch—especially in polar regions—are effectively integrated into the primitive branch. As a result, the primitive branch benefits from enhanced accuracy, ultimately enabling robust omnidirectional optical flow estimation. Additionally, we demonstrate the module’s effectiveness and broad applicability by integrating it into various iterative network architectures. Together, these components form a cohesive framework, PriOr-Flow, which establishes a new benchmark for panoramic optical flow estimation.

We evaluate PriOr-Flow on publicly available panoramic optical flow datasets, MPFDataset [28] and FlowScape [35]. As shown in Fig. 1, PriOr-RAFT establishes a new state-of-the-art, significantly outperforming previous methods. Specifically, PriOr-RAFT reduces the EPE metric by 30.0% and 29.6% on the two datasets, respectively, demonstrating a substantial performance gain. Moreover, Fig. 1 highlights that our method significantly improves accuracy in the polar regions. We further compare PriOr-Flow with the previous best approach in real-world scenarios, showing that it generalizes well across diverse practical environments. Extensive ablation studies validate the effectiveness of each proposed component.

In summary, our key contributions are:

- We propose PriOr-Flow, a novel dual-branch framework that exploits the orthogonal view’s low-distortion prior to enhance optical flow estimation in polar regions.
- We introduce the DCCL and ODDC modules to explicitly integrate low-distortion priors from the orthogonal view,

- reducing estimation errors in high-distortion areas.
- We demonstrate the broad applicability of PriOr-Flow by integrating it with multiple iterative optical flow architectures.
- PriOr-Flow sets a new state-of-the-art on multiple panoramic optical flow benchmarks and exhibits strong generalization to real-world applications.

2. Related Work

2.1. Perspective Optical Flow Estimation

Recent advances in optical flow estimation have been propelled by deep learning architectures [14, 21, 22, 38, 39, 41, 43, 55]. FlowNet [14] pioneered the use of convolutional neural networks (CNNs) to treat optical flow as a regression task, directly predicting dense motion fields from image pairs. Building on this foundation, PWC-Net [38] introduced a feature pyramid and coarse-to-fine strategy, enabling efficient estimation of large displacements by refining predictions progressively from low to high resolutions. RAFT [41] further advanced the field with a recurrent framework that constructs a 4D all-pairs cost volume and iteratively updates flow fields at fixed resolution, achieving state-of-the-art accuracy through multi-step refinement. These methods have demonstrated significant success in perspective imagery; however, their reliance on planar projection assumptions limits their direct applicability to omnidirectional data.

2.2. Panoramic Optical Flow Estimation

The increasing use of affordable panoramic cameras has created a growing demand for panoramic optical flow estimation, which enables comprehensive temporal understanding across wide fields of view. However, the spherical-to-planar projection in omnidirectional imagery (e.g., equirectangular format) introduces severe geometric distortions [2, 3], conflicting with the flat-domain priors of conventional optical flow methods. To address these challenges, recent works have developed optical flow algorithms tailored for panoramic images. LiteFlowNet360 [5] was the first deep learning-based approach for dense optical flow estimation in panoramic videos, mitigating spherical distortions using a kernel transformer network (KTN [37]) and self-supervised learning. Similarly, OmniFlowNet [2] adapts perspective-based CNNs for omnidirectional images without requiring additional training, using distortion-aware convolutions [15] to align with equirectangular projection. A key limitation of these methods is their reliance on adapting convolutional layers, which introduces computational overhead and reduces flexibility and ease of deployment. TanImg [48] estimates panoramic optical flow by dividing the spherical image into multiple tangent planes, helping to reduce distortions from equirectangular projection. However, it struggles with large displacements and

boundary continuity, limiting its effectiveness. SLOF [6] employs a siamese representation learning framework [7] with rotational augmentations and tailored losses to adapt existing flow networks for panoramic optical flow. MPF-Net [28] fuses predictions from multiple projection methods, including tri-cylinder and cube padding projections, to complement equirectangular projection. PanoFlow [35] further enhances panoramic optical flow estimation by incorporating flow distortion augmentation and deformable convolution [13], as well as a cyclic flow estimation method to handle boundary continuity. While these methods have attempted to reduce the impact of projection distortions from different angles, none have specifically addressed the most severe distortions in the polar regions. Our PriOr-Flow leverages the low-distortion prior from the orthogonal view to compensate for optical flow estimation in these polar regions, significantly improving performance in these critical areas.

2.3. Multi-view Panoramic Approaches

In the field of panoramic image processing, some studies have explored the use of multiple views of panoramic images. [19] decomposes panoramic images and stitches results from each part to avoid polar singularities. [23] simply concatenates optical flows from different views to achieve uniform flow patterns. SLOF [6] introduces additional constraints on cross-view optical flow similarity during training to learn view-invariant representations. In contrast, our proposed PriOr-Flow leverages the orthogonal view to address severe distortions in the polar regions of panoramic images.

3. Method

In this section, we present the overall architecture of PriOr-Flow. Since our method can be integrated into various networks, we use PriOr-RAFT (Fig. 4) as an example to highlight its key components. We begin by describing the orthogonal view generation process in Section 3.1. Next, we introduce the Dual-Cost Collaborative Lookup (DCCL) operator in Section 3.2. We then detail the primitive and orthogonal branches in Sections 3.4 and 3.3, respectively, with a focus on the Ortho-Driven Distortion Compensation (ODDC) module within the primitive branch. Finally, we discuss the overall loss function in Section 3.5.

3.1. Orthogonal View Generation

The equirectangular projection (ERP), the most common format for panoramic imaging, maps spherical coordinates (θ, ϕ) to a 2D plane via a linear transformation [1, 36]:

$$\begin{bmatrix} u \\ v \end{bmatrix} = \begin{bmatrix} \frac{W}{2\pi} & 0 & \frac{W}{2} \\ 0 & -\frac{H}{\pi} & \frac{H}{2} \end{bmatrix} \begin{bmatrix} \theta \\ \phi \\ 1 \end{bmatrix} \quad (1)$$

As illustrated in Fig. 2, we generate the orthogonal view by applying spherical rotation to the primitive ERP image.

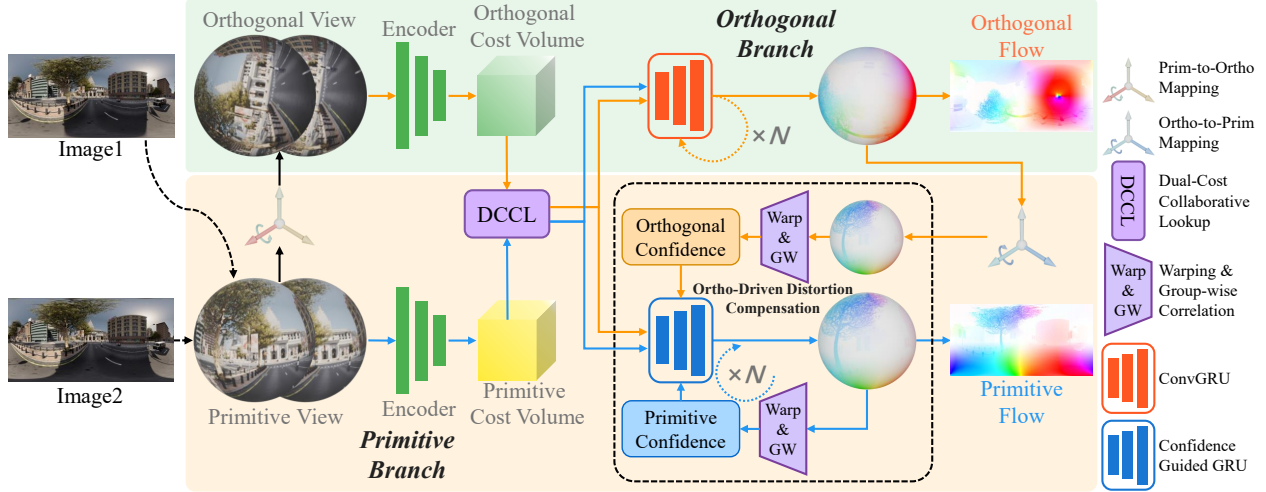


Figure 4. **Overview of Our proposed PriOr-Flow (PriOr-RAFT version).** The Dual-Cost Collaborative Lookup (DCCL) operator uses the current optical flow to jointly retrieve information from both the primitive and orthogonal cost volumes. In each iteration, the Ortho-Driven Distortion Compensation (ODDC) module aggregates motion features from the orthogonal branch to compensate for the polar regions of the primitive flow.

This spherical rotation operation \mathcal{R} is mathematically expressed as:

$$\mathbf{x}' = \mathcal{R}(\theta, \mathbf{x}) = P^{-1}(R_x(\theta) \cdot P(\mathbf{x})) \quad (2)$$

Here, \mathbf{x} represents the coordinates of an ERP pixel, and \mathbf{x}' denotes the coordinates after rotation. P is the projection from ERP image coordinates to Cartesian coordinates, and $R_x(\theta)$ is the rotation matrix about the x-axis by angle θ . The orthogonal image is obtained by re-projecting the rotated spherical surface onto the ERP plane:

$$I^o = T_p^o(I^p) = \{\text{Bi}(I^p, \mathcal{R}(-90^\circ, \mathbf{x})) \mid \mathbf{x} \in I^p\} \quad (3)$$

Where I^p denotes the primitive image, I^o denotes the orthogonal image, and Bi represents bilinear interpolation. T_p^o represents the transformation from the primitive view to the orthogonal view, while T_o^p is the inverse transformation.

3.2. Dual-Cost Collaborative Lookup (DCCL)

Due to projection distortions, significant noise may be encoded in the cost volume, especially in the polar regions, and the retrieved noisy correlation cues lead to erroneous optical flow field recovery. To address this issue, we propose the *dual-cost collaborative lookup* (DCCL) operator, which performs joint retrieval from both the primitive correlation pyramid $\{C_i^p\}$ and the orthogonal correlation pyramid $\{C_i^o\}$ (defined later in Section 3.3) to mitigate the impact of noise in the cost volume.

As shown in Fig. 5, DCCL performs correlation lookup on a unified spherical surface, mapping the indexed points to both the primitive image I_2^p and the orthogonal image I_2^o , followed by separate lookups in $\{C_i^p\}$ and $\{C_i^o\}$. Given a current estimate of the primitive optical flow $\mathcal{F}^p =$



Figure 5. Schematic diagram of the Dual-Cost Collaborative Lookup (DCCL) operator.

$\{f_1^p, f_2^p\}$, we map each pixel $\mathbf{x}^p = \{u^p, v^p\}$ in I_1^p to its estimated correspondence in I_2^p :

$$\hat{\mathbf{x}}^p = \{(u^p + f_1^p(\mathbf{x}^p)) \bmod W, v^p + f_2^p(\mathbf{x}^p)\} \quad (4)$$

where the modulo operation explicitly enforces horizontal boundary continuity for panoramic images. We then define a local grid around $\hat{\mathbf{x}}^p$:

$$\mathcal{N}(\hat{\mathbf{x}}^p)_r^p = \{\hat{\mathbf{x}}^p + \mathbf{dx} \mid \mathbf{dx} \in \mathbb{Z}^2, \|\mathbf{dx}\|_1 \leq r\} \quad (5)$$

We use $\mathcal{N}(\hat{\mathbf{x}}^p)_r^p$ to index the primitive correlation cues \mathcal{C}^p from $\{C_i^p\}$. Then we convert the primitive local grid $\mathcal{N}(\hat{\mathbf{x}}^p)_r^p$ into the orthogonal local grid:

$$\mathcal{N}(\hat{\mathbf{x}}^p)_r^o = \{\mathcal{R}(90^\circ, \mathbf{x}) \mid \mathbf{x} \in \mathcal{N}(\hat{\mathbf{x}}^p)_r^p\} \quad (6)$$

We use $\mathcal{N}(\hat{\mathbf{x}}^p)_r^o$ to index the orthogonal correlation cues \mathcal{C}^o from $\{C_i^o\}$. Then, we convert the cost map back to the primitive format:

$$\mathcal{C}^{o2p} = T_o^p(\mathcal{C}^o) \quad (7)$$

\mathcal{C}^p and \mathcal{C}^{o2p} will be fed into the confidence-guided ConvGRU to jointly guide the restoration of the primitive flow




Figure 6. Qualitative results on the test set of MPFDataset. Our PriOr-RAFT outperforms SLOF in polar regions.

field. The correlation information indexed from the orthogonal cost volume can effectively suppress the distortion noise in the polar regions of the primitive cost volume, compensating for the optical flow recovery in these regions.

3.3. Orthogonal Branch

Given two consecutive frames $\{I_1^p, I_2^p\}$, referred to as the primitive view, we first convert both to the orthogonal view:

$$\{I_1^o, I_2^o\} = \{T_p^o(I_1^p), T_p^o(I_2^p)\} \quad (8)$$

Then, following RAFT [41], we input $\{I_1^o, I_2^o\}$ into the feature network to extract the corresponding feature maps $\{f_1^o, f_2^o\}$, while feeding I_1^o into the context network to obtain contextual features. After that, we use $\{f_1^o, f_2^o\}$ to construct the orthogonal cost volume:

$$C_{ijkl}^o = \sum_h (f_1^o)_{hij} \cdot (f_2^o)_{hkl} \quad (9)$$

We then construct a 4-level correlation pyramid $\{C_i^o\} (i = 1, 2, 3, 4)$ by using 2D average pooling with a kernel size of 2 and a stride of 2 as the last two dimensions. Finally, we use a ConvGRU to iteratively update the optical flow field. In each iteration, we use the current optical flow estimation \mathcal{F}^o to retrieve correlation information from both the orthogonal correlation pyramid $\{C_i^o\}$ and the primitive correlation pyramid $\{C_i^p\}$:

$$C^o, C^{p2o} = \text{DCCL}(\{C_i^o\}, \{C_i^p\}, \mathcal{F}^o) \quad (10)$$

C^o and C^{p2o} will be fed into the ConvGRU to predict the residual flow Δf^o for updating:

$$\hat{\mathcal{F}}^o = \mathcal{F}^o + \Delta f^o \quad (11)$$

3.4. Primitive Branch

Symmetric to the orthogonal branch, the primitive branch first feeds the primitive frames $\{I_1^p, I_2^p\}$ into the encoder to

obtain the feature maps $\{f_1^p, f_2^p\}$ and constructs the correlation pyramid $\{C_i^p\} (i = 1, 2, 3, 4)$ according to Eq. (9).

Ortho-Driven Distortion Compensation (ODDC) To further mitigate the effects of polar distortion, ODDC leverages the low-distortion priors of the orthogonal view to compensate for the flow field reconstruction in the polar regions of the primitive view. To explicitly represent the advantages of each view in different regions, we compute confidence maps of the primitive flow and the orthogonal flow, respectively. Specifically, during each iteration, given the current flows \mathcal{F}^p and \mathcal{F}^o , we first convert the orthogonal flow \mathcal{F}^o to primitive format:

$$\mathcal{F}^{o2p} = T_p^o(\mathcal{F}^o) \quad (12)$$

Then we compute both the primitive and the orthogonal confidence by computing the group-wise correlation following [18, 27]:

$$\begin{aligned} G^p(\mathbf{x}) &= GW(f_1^p(\mathbf{x}), f_2^p(\mathbf{x} + \mathcal{F}^p) \bmod W) \\ G^{o2p}(\mathbf{x}) &= GW(f_1^p(\mathbf{x}), f_2^p(\mathbf{x} + \mathcal{F}^{o2p}) \bmod W) \end{aligned} \quad (13)$$

where GW denotes the group-wise correlation operation. In addition to the confidence information, we also use the current primitive flow \mathcal{F}^p to index from the primitive cost volume and the orthogonal cost volume to obtain the correlation cues to further guide the restoration of the flow field:

$$C^p, C^{o2p} = \text{DCCL}(\{C_i^p\}, \{C_i^o\}, \mathcal{F}^p) \quad (14)$$

After obtaining the correlation cues, confidence information, and flow conditions, these inputs are fed into a confidence-guided motion encoder that aggregates the primitive motion feature m^p with compensation from the orthogonal view:

$$\begin{aligned} m^p &= [\text{En}_c(C^p + C^{o2p}), \text{En}_g([G^p, G^{o2p}]), \\ &\quad \text{En}_f([\mathcal{F}^p, \mathcal{F}^{o2p}]), \mathcal{F}^p, \mathcal{F}^{o2p}] \end{aligned} \quad (15)$$

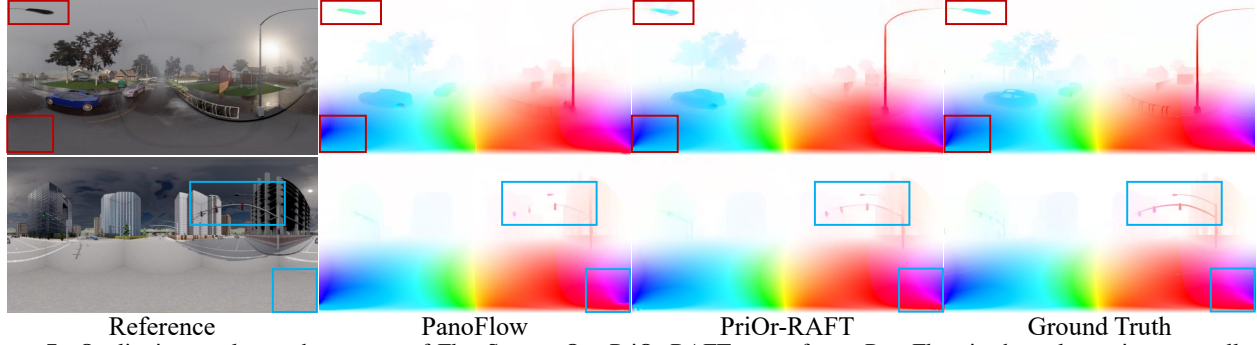


Figure 7. Qualitative results on the test set of FlowScape. Our PriOr-RAFT outperforms PanoFlow in the polar regions as well as in handling small horizontal objects.

Here, En_c , En_g and En_f are shallow convolutional layers designed to encode motion features. The aggregated feature m^p is then passed to a ConvGRU to update the hidden state of the primitive branch. Subsequently, two convolutional layers decode a residual flow Δf^p , and the primitive flow is updated in accordance with RAFT:

$$\hat{\mathcal{F}}^p = \mathcal{F}^p + \Delta f^p \quad (16)$$

3.5. Loss Function

To account for the non-uniform sampling characteristics of the ERP projection on spherical surfaces, we weight the L1 loss of each pixel by its corresponding spherical area [28, 54]. This is expressed as:

$$\|\mathcal{F}_{pr} - \mathcal{F}_{gt}\|_1^{sph} = \frac{1}{\sum \omega^j} \sum_{j \in \Omega} \|\mathcal{F}_{pr}^j - \mathcal{F}_{gt}^j\|_1 \cdot \omega^j \quad (17)$$

where ω^j denotes the spherical area weight corresponding to each pixel.

For supervision, we use the ground truth of the primitive and orthogonal views to guide the outputs of the respective branches. The loss function follows the approach in [41], where the weights are exponentially increased as the number of iterations progresses. The total loss \mathcal{L} is the sum of the primitive branch loss \mathcal{L}_p and the orthogonal branch loss \mathcal{L}_o :

$$\begin{aligned} \mathcal{L}_p &= \sum_{i=1}^N \gamma^{N-i} \|\mathcal{F}_i^p - \mathcal{F}^{gt}\|_1^{sph} \\ \mathcal{L}_o &= \sum_{i=1}^N \gamma^{N-i} \|\mathcal{F}_i^o - T_p^o(\mathcal{F}^{gt})\|_1^{sph} \\ \mathcal{L} &= \mathcal{L}_p + \mathcal{L}_o \end{aligned} \quad (18)$$

where $\gamma = 0.8$ and \mathcal{F}^{gt} represents the ground truth of the primitive view.

4. Experiment

4.1. Datasets & Evaluation Metrics

Datasets. We evaluate our model on both synthetic and real-world datasets. *MPFDataset* [28] provides dense op-

tical flow ground truth with two scenes: City (2000 training pairs, 138 testing pairs) and EFT (2211 training pairs, 99 testing pairs). *FlowScape* [35] features four weather scenarios with 5000 training pairs and 1400 testing pairs. Additionally, we test on real-world datasets *OmniPhotos* [4] and *ODVista* [42], which contain panoramic videos but lack ground truth.

Evaluation Metrics. Following [28, 48], we use end-point error (EPE) and spherical end-point error (SEPE) to evaluate panoramic optical flow:

$$SEPE = \frac{1}{|\Omega|} \sum_{i \in \Omega} d(\mathcal{F}_i^{pr}, \mathcal{F}^{gt}) \quad (19)$$

where $d(\cdot, \cdot)$ denotes the geodesic distance on the unit sphere.

4.2. Implementation Details

We implement PriOr-Flow using PyTorch and conduct experiments on NVIDIA RTX 3090 GPUs. Following the baseline approach in RAFT [41], we use the AdamW [30] optimizer and clip gradients within the range of [-1, 1]. The learning rate is scheduled using a one-cycle policy, with an initial learning rate of 1e-4. Model weights are initialized with the pre-trained RAFT weights on FlyingThings [33] before fine-tuning on the omnidirectional datasets. For the MPF-dataset, we use a batch size of 4, a learning rate of 1e-4, and train the model for 60k steps. For FlowScape, following PanoFlow [35], we use a batch size of 6, a learning rate of 1e-4, and train for 100k steps. During both training and testing, we set the number of iterations to 12.

4.3. Ablation Study

To evaluate the effectiveness of each component in PriOr-Flow, we conduct ablation experiments on the EFT scene of MPFDataset [28], which contains diverse motion patterns.

Effectiveness of proposed modules. We progressively incorporate our proposed modules into the baseline model and assess performance across different regions. As shown in Tab. 1, compared to the original RAFT operation, our DCCL module effectively suppresses distortion noise in

Model	Orthogonal View	DCCL	ODDC	Equator		Poles		All	
				EPE	SEPE	EPE	SEPE	EPE	SEPE
Baseline (RAFT)				1.09	6.29	7.90	8.56	4.49	7.43
Ortho + D	✓	✓		1.08	6.24	7.56	8.16	4.32	7.20
Full Model (PriOr-RAFT)	✓	✓	✓	1.03	5.99	5.57	6.47	3.30	6.23

Table 1. Ablation study of the effectiveness of proposed modules. DCCL stands for Dual-Cost Collaborative Lookup, and ODDC represents Ortho-Driven Distortion Compensation. We evaluate the performance of different models in both the equatorial and polar regions. Polar regions refer to areas with latitudes greater than 45° , while equatorial regions cover the rest of the sphere.




Figure 8. Qualitative results on real-world scenarios. The left two columns present the results on the OmniPhotos dataset, while the right two columns display the results on the ODVista dataset.

Model	Confidence Type	Poles		All	
		EPE	SEPE	EPE	SEPE
Ours w/o conf	-	5.65	6.60	3.36	6.37
Ours dist conf	Distortion Map	5.62	6.55	3.34	6.34
Ours (PriOr-RAFT)	Warp&GW	5.57	6.47	3.30	6.23

Table 2. Ablation study on the type of confidence in ODDC.

the cost volume, yielding notable improvements. Adding ODDC leads to significant gains, primarily benefiting polar regions (an improvement of 26.3% in the EPE metric). This suggests that motion cues from the orthogonal branch can compensate for distortions in the primitive view. We also perform an ablation study on the confidence guidance in ODDC, as shown in Tab. 2. While the distortion map provides view-specific confidence, it is fixed and independent of the input images and optical flow, limiting its ability to reflect actual flow confidence. In contrast, our method explicitly measures feature similarity before and after warping, yielding more accurate confidence estimation and enabling more effective flow fusion.

Universality of proposed modules. To verify the general applicability of our modules, we integrate them into three classic iterative perspective optical flow estimation methods as baselines. Specifically, we replace the original cost volume lookup with DCCL, introduce an additional orthogonal branch, and substitute the primitive branch’s ConvGRU with ODDC. As shown in Tab. 3, all methods achieve significant improvements in both EPE and SEPE. Furthermore, even when reducing the number of iterations to 4,

Model	EPE		SEPE	Run-time
	(px)	Reduction		
RAFT [41]	4.49	-	7.43	0.07
PriOr-RAFT (4-iter)	3.89	↓ 13.4%	7.17	0.10
PriOr-RAFT	3.30	↓ 26.5%	6.23	0.20
GMA [22]	4.26	-	7.07	0.07
PriOr-GMA (4-iter)	3.51	↓ 17.6%	6.68	0.10
PriOr-GMA	3.25	↓ 23.7%	6.17	0.20
SKFlow [39]	3.79	-	6.55	0.12
PriOr-SKFlow (4-iter)	3.33	↓ 12.1%	6.28	0.13
PriOr-SKFlow	3.19	↓ 15.8%	6.13	0.30

Table 3. Ablation study of the universality of PriOr-Flow.

View	Axis	θ	EPE(px)	SEPE(mm)
y-ortho	y	90°	3.55	6.46
x-semi-ortho	x	45°	3.41	6.46
x-ortho	x	90°	3.30	6.23

Table 4. Ablation study of the selection of the orthogonal view.

Metric	Model	Number of Iterations				
		3	4	8	12	16
EPE (px)	RAFT	6.09	5.49	4.64	4.49	4.48
	PriOr-RAFT	4.37	3.89	3.39	3.30	3.30

Table 5. Ablation study of the number of iterations.

these methods still achieve substantial accuracy gains while maintaining nearly the same inference time. In particular, PriOr-GMA accomplishes a 17.6% performance improvement with almost identical inference time.

Method	Baseline	MPFDataset [28]						FlowScape [35]									
		EFT		City		All		Sunny		Cloud		Fog		Rain		All	
		EPE	SEPE	EPE	SEPE	EPE	SEPE	EPE	SEPE	EPE	SEPE	EPE	SEPE	EPE	SEPE	EPE	SEPE
SphereNet [12]	RAFT [41]	13.2	15.7	8.28	7.44	10.7	19.4	13.7	22.9	9.94	15.6	12.5	19.3	13.3	20.2	12.9	21.0
TanImg [48]	DIS [25]	8.04	19.3	3.74	6.48	5.89	12.9	19.6	31.1	15.6	25.4	17.6	27.8	19.6	31.9	18.7	29.9
TanImg [48]	RAFT [41]	4.38	9.52	3.13	5.06	3.76	7.29	19.6	28.2	14.7	17.9	17.8	23.8	17.3	22.7	18.3	25.3
MPF-net [28]	PWC-net [38]	5.06	10.49	1.78	3.24	3.42	6.87	-	-	-	-	-	-	-	-	-	-
SLOF [6]	RAFT [41]	4.98	8.20	1.35	2.06	3.17	5.13	7.84	6.07	6.00	3.64	8.08	6.35	7.66	6.24	7.59	5.79
PanoFlow [35]	RAFT [41]	-	-	-	-	-	-	3.61	5.19	1.38	1.97	3.60	4.78	4.25	5.92	3.38	4.78
PanoFlow [35]	CSFlow [34]	-	-	-	-	-	-	3.55	4.76	1.47	1.86	3.55	4.70	3.93	5.46	3.31	4.44
PriOr-Flow	RAFT [41]	3.30	6.23	1.13	1.88	2.22	4.06	2.41	3.56	0.93	1.36	2.81	4.09	2.95	4.71	2.33	3.49

Table 6. Quantitative evaluation on MPFDataset and FlowScape. The **best** and **second best** are marked with colors. “-” denotes that the model weights are not publicly available.

Model	Equator		Poles		All	
	EPE	SEPE	EPE	SEPE	EPE	SEPE
PanoFlow	0.52	2.87	6.25	6.68	3.38	4.78
PriOr-RAFT	0.53	2.94	4.13	4.03	2.33	3.49

Table 7. Quantitative evaluation on FlowScape test set in different regions.

The selection of the orthogonal view. We evaluate different rotation axes and angles θ as shown in Tab. 4. A 90° rotation around the y-axis splits the polar region, disrupting continuity and complicating orthogonal flow estimation. A 45° rotation around the x-axis results in a “semi-orthogonal” view, but its low-distortion region covers less of the primitive image, leading to suboptimal compensation for polar regions. Therefore, we select the 90° x-axis rotation as the orthogonal view for PriOr-Flow.

Number of iterations. PriOr-Flow achieves better performance with fewer iterations. As shown in Tab. 5, compared to the baseline, PriOr-RAFT outperforms it with just 3 iterations. This demonstrates that incorporating motion features from the orthogonal branch significantly accelerates model convergence.

4.4. Comparisons with State-of-the-art

We quantitatively compare PriOr-RAFT with state-of-the-art panoramic optical flow estimation methods, including weight transformation-based, tangent plane-based [48], and ERP-based methods [6, 28, 35], on MPFDataset [28] and FlowScape [35]. For weight transformation-based methods, we adopt the more efficient SphereNet [12] weight transformation applied to RAFT, pre-trained on FlyingThings [33], instead of using equirectangular convolution [15] as in OmniFlowNet [2]. This approach achieves an EPE of 12.9 on FlowScape, significantly outperforming OmniFlowNet’s 19.6 EPE. For tangent plane-based methods, to ensure fairness, we replace the original DIS [25] used in [48] with RAFT, also pre-trained on FlyingThings. The results are summarized in Tab. 6.

MPFDataset. Our PriOr-RAFT achieves state-of-the-art performance in both the EFT and City scenarios. Specif-

ically, for planar optical flow (EPE), our method surpasses SLOF [6] by 30.0%, while for spherical optical flow (SEPE), it achieves a 20.9% improvement. Fig. 6 provides a visual comparison with SLOF [6].

FlowScape. PriOr-RAFT achieves the best performance across all weather conditions, improving EPE by 29.6% and SEPE by 21.4% over the previous SOTA. To verify its effectiveness in polar regions, we compare it with PanoFlow in different regions. As shown in Tab. 7, PriOr-RAFT outperforms PanoFlow by 39.7% in polar regions. However, due to the trade-off between distortion in different areas, our method is slightly inferior to PanoFlow in the equatorial regions. Fig. 7 presents a visual comparison between PriOr-RAFT and PanoFlow (RAFT) on FlowScape.

4.5. Generalization to Real-World Scenarios.

Due to the lack of panoramic datasets with ground truth optical flow in real-world scenarios, we evaluate our method qualitatively on OmniPhotos [4] and ODVista [42] datasets. Both PriOr-RAFT and PanoFlow [35] were trained on FlowScape and directly applied to these real-world datasets. Fig. 8 illustrates the qualitative results, demonstrating the generalization capability of our approach. Notably, the ground region in the second column and the roof region in the third column highlight our method’s ability to effectively predict optical flow in polar regions.

5. Conclusion

We propose PriOr-Flow, a novel and universal dual-branch method for panoramic optical flow estimation. By introducing the Dual-Cost Collaborative Lookup (DCCL) operator and the Ortho-Driven Distortion Compensation (ODDC) module, PriOr-Flow leverages the low-distortion prior from the orthogonal view to significantly improve performance, especially in polar regions. Our PriOr-RAFT model achieves state-of-the-art performance on both the MPFDataset and FlowScape, with notable improvements in accuracy. Additionally, PriOr-Flow demonstrates strong generalization in real-world scenarios, validating its effectiveness beyond synthetic datasets.

Acknowledgement. This research is supported by the National Key R&D Program of China (2024YFE0217700), National Natural Science Foundation of China (62472184), the Fundamental Research Funds for the Central Universities, and the Innovation Project of Optics Valley Laboratory (Grant No. OVL2025YZ005)

References

- [1] Hao Ai, Zidong Cao, and Lin Wang. A survey of representation learning, optimization strategies, and applications for omnidirectional vision. *arXiv preprint arXiv:2502.10444*, 2025. 3
- [2] Charles-Olivier Artizzu, Haozhou Zhang, Guillaume Allibert, and Cédric Demonceaux. Omniflownet: a perspective neural network adaptation for optical flow estimation in omnidirectional images. In *2020 25th International Conference on Pattern Recognition (ICPR)*, pages 2657–2662. IEEE, 2021. 1, 3, 8
- [3] Roberto G de A Azevedo, Neil Birkbeck, Francesca De Simone, Ivan Janatra, Balu Adsumilli, and Pascal Frossard. Visual distortions in 360° videos. *IEEE Transactions on Circuits and Systems for Video Technology*, 30(8):2524–2537, 2019. 3
- [4] Tobias Bertel, Mingze Yuan, Reuben Lindroos, and Christian Richardt. Omniphotos: casual 360 vr photography. *ACM Transactions on Graphics (TOG)*, 39(6):1–12, 2020. 6, 8
- [5] Keshav Bhandari, Ziliang Zong, and Yan Yan. Revisiting optical flow estimation in 360 videos. In *2020 25th International Conference on Pattern Recognition (ICPR)*, pages 8196–8203. IEEE, 2021. 1, 3
- [6] Keshav Bhandari, Bin Duan, Gaowen Liu, Hugo Latapie, Ziliang Zong, and Yan Yan. Learning omnidirectional flow in 360 video via siamese representation. In *European Conference on Computer Vision*, pages 557–574. Springer, 2022. 2, 3, 8
- [7] Xinlei Chen and Kaiming He. Exploring simple siamese representation learning. In *Proceedings of the IEEE/CVF conference on computer vision and pattern recognition*, pages 15750–15758, 2021. 2, 3
- [8] Junda Cheng, Xin Yang, Yuechuan Pu, and Peng Guo. Region separable stereo matching. *IEEE Transactions on Multimedia*, 25:4880–4893, 2022. 1
- [9] Junda Cheng, Gangwei Xu, Peng Guo, and Xin Yang. Coatsnet: Fully exploiting convolution and attention for stereo matching by region separation. *International Journal of Computer Vision*, 132(1):56–73, 2024. 1
- [10] Junda Cheng, Wei Yin, Kaixuan Wang, Xiaozhi Chen, Shijie Wang, and Xin Yang. Adaptive fusion of single-view and multi-view depth for autonomous driving. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 10138–10147, 2024. 2
- [11] Junda Cheng, Longliang Liu, Gangwei Xu, Xianqi Wang, Zhaoxing Zhang, Yong Deng, Jinliang Zang, Yurui Chen, Zhipeng Cai, and Xin Yang. Monster: Marry monodepth to stereo unleashes power. *arXiv preprint arXiv:2501.08643*, 2025. 2
- [12] Benjamin Coors, Alexandru Paul Condurache, and Andreas Geiger. Spherenet: Learning spherical representations for detection and classification in omnidirectional images. In *Proceedings of the European conference on computer vision (ECCV)*, pages 518–533, 2018. 8
- [13] Jifeng Dai, Haozhi Qi, Yuwen Xiong, Yi Li, Guodong Zhang, Han Hu, and Yichen Wei. Deformable convolutional networks. In *Proceedings of the IEEE international conference on computer vision*, pages 764–773, 2017. 2, 3
- [14] 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. 3
- [15] Clara Fernandez-Labrador, Jose M Facil, Alejandro Perez-Yus, Cédric Demonceaux, Javier Civera, and Jose J Guerrero. Corners for layout: End-to-end layout recovery from 360 images. *IEEE Robotics and Automation Letters*, 5(2): 1255–1262, 2020. 3, 8
- [16] Chen Gao, Ayush Saraf, Jia-Bin Huang, and Johannes Kopf. Flow-edge guided video completion. In *Computer Vision—ECCV 2020: 16th European Conference, Glasgow, UK, August 23–28, 2020, Proceedings, Part XII 16*, pages 713–729. Springer, 2020. 1
- [17] Andreas Geiger, Philip Lenz, and Raquel Urtasun. Are we ready for autonomous driving? the kitti vision benchmark suite. In *2012 IEEE conference on computer vision and pattern recognition*, pages 3354–3361. IEEE, 2012. 1
- [18] Xiaoyang Guo, Kai Yang, Wukui Yang, Xiaogang Wang, and Hongsheng Li. Group-wise correlation stereo network. In *Proceedings of the IEEE/CVF conference on computer vision and pattern recognition*, pages 3273–3282, 2019. 5
- [19] Kenji Hara, Kohei Inoue, and Kiichi Urahama. Gradient operators for feature extraction from omnidirectional panoramic images. *Pattern Recognition Letters*, 54:89–96, 2015. 3
- [20] Zhewei Huang, Tianyuan Zhang, Wen Heng, Boxin Shi, and Shuchang Zhou. Rife: real-time intermediate flow estimation for video frame interpolation 2020, 2011. 1
- [21] Zhaoyang Huang, Xiaoyu Shi, Chao Zhang, Qiang Wang, Ka Chun Cheung, Hongwei Qin, Jifeng Dai, and Hongsheng Li. Flowformer: A transformer architecture for optical flow. In *European conference on computer vision*, pages 668–685. Springer, 2022. 1, 3
- [22] Shihao Jiang, Dylan Campbell, Yao Lu, Hongdong Li, and Richard Hartley. Learning to estimate hidden motions with global motion aggregation. In *Proceedings of the IEEE/CVF international conference on computer vision*, pages 9772–9781, 2021. 1, 3, 7
- [23] Dabae Kim, Sarthak Pathak, Alessandro Moro, Ren Kوماتsu, Atsushi Yamashita, and Hajime Asama. E-cnn: Accurate spherical camera rotation estimation via uniformization of distorted optical flow fields. In *ICASSP 2019-2019 IEEE International Conference on Acoustics, Speech and Signal Processing (ICASSP)*, pages 2232–2236. IEEE, 2019. 3

- [24] 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
- [25] Till Kroeger, Radu Timofte, Dengxin Dai, and Luc Van Gool. Fast optical flow using dense inverse search. In *Computer Vision–ECCV 2016: 14th European Conference, Amsterdam, The Netherlands, October 11–14, 2016, Proceedings, Part IV 14*, pages 471–488. Springer, 2016. 8
- [26] Suyoung Lee, Jaeyoung Chung, Kihoon Kim, Jaeyoo Huh, Gunhee Lee, Minsoo Lee, and Kyoung Mu Lee. Omnisplat: Taming feed-forward 3d gaussian splatting for omnidirectional images with editable capabilities. *arXiv preprint arXiv:2412.16604*, 2024. 1
- [27] Jiankun Li, Peisen Wang, Pengfei Xiong, Tao Cai, Ziwei Yan, Lei Yang, Jiangyu Liu, Haoqiang Fan, and Shuaicheng Liu. Practical stereo matching via cascaded recurrent network with adaptive correlation. In *Proceedings of the IEEE/CVF conference on computer vision and pattern recognition*, pages 16263–16272, 2022. 5
- [28] Yiheng Li, Connelly Barnes, Kun Huang, and Fang-Lue Zhang. Deep 360° optical flow estimation based on multi-projection fusion. In *European Conference on Computer Vision*, pages 336–352. Springer, 2022. 1, 2, 3, 6, 8
- [29] Xiaozhang Liu, Hui Liu, and Yuxiu Lin. Video frame interpolation via optical flow estimation with image inpainting. *International Journal of Intelligent Systems*, 35(12):2087–2102, 2020. 1
- [30] Ilya Loshchilov and Frank Hutter. Decoupled weight decay regularization. *arXiv preprint arXiv:1711.05101*, 2017. 6
- [31] Jieming Lou, Weide Liu, Zhuo Chen, Fayao Liu, and Jun Cheng. Elfnet: Evidential local-global fusion for stereo matching. In *Proceedings of the IEEE/CVF International Conference on Computer Vision*, pages 17784–17793, 2023. 2
- [32] Zeyu Ma, Zachary Teed, and Jia Deng. Multiview stereo with cascaded epipolar raft. In *European Conference on Computer Vision*, pages 734–750. Springer, 2022. 1
- [33] 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. 6, 8
- [34] Hao Shi, Yifan Zhou, Kailun Yang, Xiaoting Yin, and Kaiwei Wang. Csflo: Learning optical flow via cross strip correlation for autonomous driving. In *2022 IEEE intelligent vehicles symposium (IV)*, pages 1851–1858. IEEE, 2022. 8
- [35] Hao Shi, Yifan Zhou, Kailun Yang, Xiaoting Yin, Ze Wang, Yaozu Ye, Zhe Yin, Shi Meng, Peng Li, and Kaiwei Wang. Panoflow: Learning 360° optical flow for surrounding temporal understanding. *IEEE Transactions on Intelligent Transportation Systems*, 24(5):5570–5585, 2023. 1, 2, 3, 6, 8
- [36] John P Snyder. *Flattening the earth: two thousand years of map projections*. University of Chicago Press, 1997. 3
- [37] Yu-Chuan Su and Kristen Grauman. Kernel transformer networks for compact spherical convolution. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 9442–9451, 2019. 3
- [38] Deqing Sun, Xiaodong Yang, Ming-Yu Liu, and Jan Kautz. Pwc-net: Cnns for optical flow using pyramid, warping, and cost volume. In *Proceedings of the IEEE conference on computer vision and pattern recognition*, pages 8934–8943, 2018. 3, 8
- [39] Shangkun Sun, Yuanqi Chen, Yu Zhu, Guodong Guo, and Ge Li. Skflow: Learning optical flow with super kernels. *Advances in Neural Information Processing Systems*, 35: 11313–11326, 2022. 1, 3, 7
- [40] Yule Sun, Ang Lu, and Lu Yu. Weighted-to-spherically-uniform quality evaluation for omnidirectional video. *IEEE signal processing letters*, 24(9):1408–1412, 2017. 2
- [41] Zachary Teed and Jia Deng. 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*, pages 402–419. Springer, 2020. 1, 3, 5, 6, 7, 8
- [42] Ahmed Telili, Ibrahim Farhat, Wassim Hamidouche, and Hadi Amirpour. Odvista: An omnidirectional video dataset for super-resolution and quality enhancement tasks. In *2024 IEEE International Conference on Image Processing (ICIP)*, pages 131–136. IEEE, 2024. 1, 6, 8
- [43] 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. 1, 3
- [44] 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
- [45] Xiangyu Xu, Li Siyao, Wenxiu Sun, Qian Yin, and Ming-Hsuan Yang. Quadratic video interpolation. *Advances in Neural Information Processing Systems*, 32, 2019. 1
- [46] Fanghua Yu, Xintao Wang, Mingdeng Cao, Gen Li, Ying Shan, and Chao Dong. Orst: Omnidirectional image super-resolution with distortion-aware transformer. In *Proceedings of the IEEE/CVF conference on computer vision and pattern recognition*, pages 13283–13292, 2023. 2
- [47] Hu Yu, Naishan Zheng, Man Zhou, Jie Huang, Zeyu Xiao, and Feng Zhao. Frequency and spatial dual guidance for image dehazing. In *European conference on computer vision*, pages 181–198. Springer, 2022. 2
- [48] Mingze Yuan and Christian Richardt. 360 optical flow using tangent images. In *British Machine Vision Conference: (BMVC)*. Christian Richardt, 2021. 2, 3, 6, 8
- [49] Shuai Yuan, Shuzhi Yu, Hannah Kim, and Carlo Tomasi. Semarflow: Injecting semantics into unsupervised optical flow estimation for autonomous driving. In *Proceedings of the IEEE/CVF International Conference on Computer Vision*, pages 9566–9577, 2023. 1
- [50] Zikang Yuan, Qingjie Wang, Ken Cheng, Tianyu Hao, and Xin Yang. Sdv-loam: Semi-direct visual–lidar odometry and mapping. *IEEE Transactions on Pattern Analysis and Machine Intelligence*, 45(9):11203–11220, 2023.

- [51] Zikang Yuan, Jie Deng, Ruiye Ming, Fengtian Lang, and Xin Yang. Sr-livo: Lidar-inertial-visual odometry and mapping with sweep reconstruction. *IEEE Robotics and Automation Letters*, 2024.
- [52] Zikang Yuan, Fengtian Lang, Tianle Xu, and Xin Yang. Sr-liv: Lidar-inertial odometry with sweep reconstruction. In *2024 IEEE/RSJ International Conference on Intelligent Robots and Systems (IROS)*, pages 7862–7869. IEEE, 2024. [1](#)
- [53] Zikang Yuan, Fengtian Lang, Jie Deng, Hongcheng Luo, and Xin Yang. Voxel-svio: Stereo visual-inertial odometry based on voxel map. *IEEE Robotics and Automation Letters*, 2025. [1](#)
- [54] Ziheng Zhang, Yanyu Xu, Jingyi Yu, and Shenghua Gao. Saliency detection in 360 videos. In *Proceedings of the European conference on computer vision (ECCV)*, pages 488–503, 2018. [6](#)
- [55] 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. [1](#), [3](#)
- [56] Shipeng Zhu, Zuoyan Zhao, Pengfei Fang, and Hui Xue. Improving scene text image super-resolution via dual prior modulation network. In *Proceedings of the AAAI Conference on Artificial Intelligence*, pages 3843–3851, 2023. [2](#)
- [57] Yiming Zuo and Jia Deng. View synthesis with sculpted neural points. *arXiv preprint arXiv:2205.05869*, 2022. [1](#)