

SVI360: Spherical Video Interpolation

Le-Kim Nguyen¹, Renato Martins¹,
Pascal Vasseur², and Cedric Demonceaux¹

¹ Université Bourgogne Europe, ICB UMR 6303 CNRS, France

² Université de Picardie Jules Verne, MIS UR 4290, France

Abstract. This paper addresses the problem of omnidirectional video interpolation, which plays an essential role in applications such as virtual reality and immersive video enhancement. Existing video interpolation methods are not well-suited for spherical videos, as they have difficulty handling severe distortions close to the poles. To address this issue, we propose SVI360, a dual-branch framework that combines the image frame and its rotated orthogonal view to deal with these distortions. The core methodological aspect of the approach is to reinforce equivariance of the flow displacements between the original and orthogonal views to improve intermediate frame prediction. Experiments show that our method outperforms state-of-the-art approaches in interpolation quality while maintaining accurate optical flow in four different public benchmarks. Code and pre-trained models are available at: https://icb-vision-ai.github.io/video360_interpolation/

Keywords: Video frame interpolation · Omnidirectional images · Spherical optical flow estimation.

1 Introduction

Spherical video, also known as omnidirectional or panoramic video, has attracted increasing research interest and applicability since it encapsulates 360-scene information for each image frame. This is a key capability in virtual reality, navigation and robotics. As their application increases across domains such as entertainment, education, and training, the demand for high-quality 360-degree videos becomes even more critical. In this context, achieving smooth video frame transitions is particularly important, as higher frame rates greatly enhance user immersion and help reduce cybersickness. However, streaming or transmitting high frame rate 360° videos requires substantial bandwidth [32], making it inefficient in many real-world settings. To mitigate this issue, video interpolation has emerged as a promising solution, enabling the synthesis of intermediate frames between two consecutive images. By generating additional frames locally, such methods can deliver smoother frame transitions while significantly reducing transmission requirements. Existing approaches [8, 11, 13, 20, 22] have achieved impressive results on perspective images. Yet their performance drops sharply when applied to spherical imagery (please see some examples in Fig. 1).




Fig. 1: Limitations of perspective video interpolation methods on spherical images. Red boxes indicate regions where existing methods produce severe ghosting and blurring, whereas the proposed approach better preserves fine image details.

Compared with typical perspective images, spherical data introduce severe geometric distortions, non-uniform pixel density, and large pixel flow displacements. These characteristics make correspondence estimation and motion reconstruction considerably more challenging, emphasizing the need for architectures designed to handle spherical data. Despite the growing importance of 360° video, omnidirectional video interpolation methods remain clearly overlooked. In this context, 360VFI [15] is the first method and benchmark dedicated to 360° frame interpolation. It introduces a distortion-aware design that uses spherical geometry priors for feature extraction and fusion, and provides an evaluation protocol based on flow magnitude. They stratified the test set into four distinct settings based on the vertical flow magnitude, where the flow may result from object motion and camera movement. In the Easy settings, characterized by very small pixel displacement between two frames (average flow magnitude < 2), the method can perform well. However, as the flow magnitude increases, the performance drops significantly, indicating that the approach still has limitations in handling large pixel displacements. In this work, we introduce SVI360, a novel video interpolation approach designed to improve the fidelity of interpolated frames while maintaining high-quality optical flow estimation, particularly in image regions containing larger distortions.

Our method adopts a coarse-to-fine base strategy with progressively increasing resolutions, inspired by AMT [13], where intermediate optical flows and intermediate frame features are iteratively refined using dense all-pairs bidirectional correlation volumes. To adapt this concept to spherical geometry, we incorporate the key hypothesis from PriOr-Flow [14], which suggests that generating orthogonal views of spherical images can provide complementary priors that improve optical flow estimation in the primitive (original) spherical view. Based on this insight, we design a dual-branch architecture where the orthogonal branch refines both the optical flow and the interpolated frame in the primitive branch through cross-view feature interaction, yielding more accurate frame interpolations in four widely adopted public datasets: two synthetic datasets (FlowScape [21] and Flow360 [1]) and two real-world datasets (ODV360 [2] and 360VFI [15]). The obtained results consistently demonstrate the capability of addressing distortion

near the two poles while also maintaining strong performance under large motion fields. On FlowScape, where the task is middle-frame interpolation, SVI360 surpasses the best competitor by +0.63 dB in PSNR and +0.33 dB in WS-PSNR, while also achieving higher accuracy for optical flow estimation. On Flow360, which evaluates interpolation at arbitrary temporal positions, SVI360 also has the best performance on all reported interpolation metrics, including +0.25 dB PSNR gain over the second-best method. Clear improvements are also observed when considering real-world ODV360 and 360VFI benchmarks, indicating the learned interpolation model also presents good generalization capabilities across scenes. Visual inspection on frames from several scenes also agrees with these findings. To increase interpolation accuracy while maintaining competitive computational cost, our work brings three main contributions:

1. We present a new omnidirectional video interpolation method that leverages equivariance pixel displacement priors from a rotated orthogonal view of the sphere to guide optical flow estimation and frame interpolation of the original spherical frame.
2. In order to account for non-uniform pixel density along the spherical view, we introduce a Spherical Weighted Charbonnier reconstruction loss term, enhancing performance in the central image regions that contain the majority of visual information.
3. Extensive evaluations on omnidirectional benchmarks, including FlowScape, Flow360, ODV360, and 360VFI are provided, where our method consistently outperforms state-of-the-art approaches in terms of interpolation quality while achieving competitive optical flow accuracy.

2 Related Work

2.1 Perspective Video Frame Interpolation

Video frame interpolation (VFI) for perspective images has been extensively studied and can be broadly categorized into phase-based, kernel-based, motion-based, and more recently generative approaches [12, 16–19, 34]. Phase-based methods [16, 17] represent motion as pixel-wise phase shifts and reconstruct intermediate frames by interpolating phase and amplitude across a multi-scale pyramid, but they are often limited to small pixel displacements. On the other hand, kernel-based approaches such as AdaCoF [12] further adopt deformable convolutions to adjust the receptive field. However, these methods often yield blurry results for large pixel displacements and require large kernels with many parameters to handle large displacements. Motion-based VFI approaches have achieved substantial improvements with the recent advances of optical flow methods [7, 8, 22]. Recent methods adopt a coarse-to-fine refinement strategy [11, 13, 20]. In particular, AMT [13] reconstructs dense all-pairs correlation volumes to model large-displacement correspondences, which inspires our method. In parallel, recent generative approaches have explored video inbetweening by

leveraging diffusion-based video generation models. For instance, FCVG [34] introduces a frame-wise condition to guide video generation to improve temporal stability under complex motions. Despite their strong performance on perspective images, these methods degrade considerably on omnidirectional images due to their reliance on planar projection assumptions which limits their direct applicability to omnidirectional data. In contrast, our approach presents a dual-branch AMT base model, with carefully designed major adaptations to handle larger pixel displacements and distortions of omnidirectional images.

2.2 Spherical Video Interpolation and Optical Flow

Processing omnidirectional images requires the adaptation of typical perspective signal operations [5, 6, 25]. SphereNet [5] and EquiConvs [6] propose geometry-aware spherical convolutions that account for distortion and rotational properties of the sphere. Building on a similar idea, Spherical CNNs [4] perform convolution directly on the sphere to achieve rotation-equivariant representations. Another trend is to reduce distortion by splitting spherical images into multiple perspective views. Features are extracted from tangent-plane projections and fused across views [24]; cubemap-based methods follow the same strategy but often suffer from discontinuities at view boundaries and increased system complexity. These operations have been transposed to optical flow estimation, *e.g.* SLOF [1] and PriOr-Flow [14], mitigating polar distortions by exploiting cross-view consistency. However, frame interpolation to omnidirectional imagery remains underexplored with 360VFI [15] being the only representative method. 360VFI mitigates spherical distortions by incorporating distortion-aware priors in both feature extraction and frame synthesis. Nevertheless, its learned offset design from distortion maps is less well justified than static precomputed offsets such as EquiConvs, which are effective in practice, and it does not explicitly handle large pixel displacements. Our work addresses these limitations through a coarse-to-fine strategy and spherical refinement, resulting in better interpolations consistently even for increasing flow displacements as shown in the “Hard” and “Extreme” splits in the experiments.

3 Method

Given two input spherical frames $\mathbf{I}_0$ and $\mathbf{I}_1$, our goal is to predict the intermediate frame $\mathbf{I}_t$ at time $t \in (0, 1)$. An overview of our approach is given in Fig. 2. We first generate an orthogonal view for each input frame by applying a 90° rotation around the x-axis. This design follows the distortion characteristics of equirectangular images: equatorial regions in primitive view provide more reliable motion cues, while polar regions are heavily distorted. After rotation, the original polar regions move closer to the center in the orthogonal view allowing more accurate flow estimation in these regions [14]. Once the orthogonal view is processed, its motion information is rotated back to the original orientation and used to refine the primitive branch. Unlike PriOr-Flow, which focuses on optical




Fig. 2: Overview of SVI360. The primitive (original) spherical frame and its rotated orthogonal view are processed by separate encoders, which are then reconstructed with a three-stage coarse-to-fine decoder. Correlation volumes refine optical flow and intermediate features across stages, while the “Spherical Refiner” component transfers orthogonal cues to the primitive branch to correct distortions. Finally, the spherical multi-field refinement component fuses the outputs from the two branches to produce the final interpolated frame.

flow refinement, SVI360 is designed for spherical video frame interpolation. It performs multi-stage coarse-to-fine refinement, jointly updating optical flow and intermediate features to improve the final interpolated frame. We will describe the details of each component in the following sections.

3.1 Image Rotation

The input equirectangular frame $\mathbf{I}(u, v)$ is first mapped from 2D pixel coordinates to spherical coordinates (θ, ϕ) , and then converted to the corresponding 3D Cartesian coordinate (x, y, z) . PriOr-Flow [14] found that empirically a rotation angle of $\alpha = 90^\circ$ around the x-axis optimally reduces polar distortions, allowing more accurate optical flow estimation in high-latitude regions. Finally, the rotated 3D coordinates are projected back to the equirectangular representation to obtain the orthogonal view. We denote this transformation from the primitive view to the orthogonal view as $\mathbf{T}_p^o$, and its inverse as $\mathbf{T}_o^p$. Given two consecutive equirectangular frames $\{\mathbf{I}_0^p, \mathbf{I}_1^p\}$, the corresponding orthogonal views are obtained as:

$$\{\mathbf{I}_0^o, \mathbf{I}_1^o\} = \{\mathbf{T}_p^o(\mathbf{I}_0^p), \mathbf{T}_p^o(\mathbf{I}_1^p)\}. \quad (1)$$

3.2 Correlation Volume Construction and Lookup




Fig. 3: Update modules in SVI360. (a) AMT-style recurrent update used in both primitive and orthogonal branches to refine bidirectional flow and intermediate features with DCCL cues. (b) Spherical Refiner component that fuses cross-view information. $\mathcal{L}$ denotes the correlation lookup operator that retrieves correlation features based on the current flow estimate. $\mathcal{U}_1$ and $\mathcal{U}_2$ are update networks that predict residual corrections for flow and intermediate features; Further details on the architecture are in the Supplementary Material.

Following RAFT [27], we build an all-pairs correlation volume for each branch $b \in \{p, o\}$, where p, o denotes the primitive and orthogonal branch respectively. Given the feature maps $\mathbf{g}_0^b$ and $\mathbf{g}_1^b$ extracted from the two input frames in the corresponding branch b , the correlation volume is computed as:

$$\mathcal{C}_{ijkl}^b = \sum_h (\mathbf{g}_0^b)_{hij} (\mathbf{g}_1^b)_{hkl}. \quad (2)$$

We further construct a 4-level correlation pyramid $\{\mathcal{C}_i^b\}_{i=1}^4$ via average pooling along the last two dimensions. For correlation lookup, we adopt the Dual-Cost Collaborative Lookup (DCCL) strategy from PriOr-Flow [14]. Given the current flow estimate, the lookup operation retrieves two types of correlation cues: one from the current branch and one from the corresponding rotated branch. We denote this operation as

$$\mathbf{C}^b, \mathbf{C}^{\bar{b} \rightarrow b} = \text{DCCL} \left(\{\mathcal{C}_i^b\}_{i=1}^4, \{\mathcal{C}_i^{\bar{b}}\}_{i=1}^4, \mathbf{f}^b \right), \quad (3)$$

where $b \in \{p, o\}$ is the current branch, $\bar{b}$ is the counterpart branch, and $\mathbf{f}^b$ is the estimated flow of current branch. Intuitively, DCCL first uses the current flow to locate matching regions between two primitive views. Then, it also checks the corresponding regions in the rotated view, where some distorted areas may become easier to match. The resulting same-view and cross-view cues are jointly fed into the AMT-style update component (Fig. 3(a)), enabling the refinement of both optical flow and intermediate features, instead of only refining the optical flow as in PriOr-Flow. Specifically, following the update component as in AMT [13], each branch applies a lightweight recurrent update block that takes the current flow, correlation features, and intermediate features as input, and predicts residuals for both optical flow and intermediate features. The use of these cues in the Spherical Refiner is described in the next section.

3.3 Spherical Refiner

After updating the optical flow and feature maps in both branches, we feed them into the Spherical Refiner. This component is designed as a residual update that jointly refines the optical flow and intermediate representations solely for the primitive branch by taking as input the primitive branch cues $(\mathbf{f}^p, \mathbf{X}_t^p)$ and the orthogonal branch cues $(\mathbf{f}^o, \mathbf{X}_t^o)$. To align both views in a common reference frame, we first rotate the orthogonal flow and intermediate features back to the primitive domain:

$$\mathbf{f}^{o \rightarrow p} = \mathbf{T}_o^p(\mathbf{f}^o), \quad \mathbf{X}_t^{o \rightarrow p} = \mathbf{T}_o^p(\mathbf{X}_t^o), \quad (4)$$

Given the updated flow estimates, we perform correlation lookup again to obtain the primitive correlation feature maps $\mathbf{C}^p$ and $\mathbf{C}^{o \rightarrow p}$. All cues are then fed into the update block $\mathcal{U}_2$, as illustrated in Fig. 3(b), to predict residual corrections for the primitive branch:

$$\Delta \mathbf{f}^p, \Delta \mathbf{X}_t^p = \mathcal{U}_2(\mathbf{f}^p, \mathbf{X}_t^p, \mathbf{C}^p, \mathbf{C}^{o \rightarrow p}, \mathbf{X}_t^{o \rightarrow p}, \mathbf{f}^{o \rightarrow p}). \quad (5)$$

The refined primitive flow and feature maps are then updated as:

$$\mathbf{f}^p \leftarrow \mathbf{f}^p + \Delta \mathbf{f}^p, \quad \mathbf{X}_t^p \leftarrow \mathbf{X}_t^p + \Delta \mathbf{X}_t^p. \quad (6)$$

Then, we upsample the flows and intermediate features using the AMT decoder design [13], but applied symmetrically to both the primitive and orthogonal branches. These upsampled representations are then passed to the next stage, where refinement is repeated at a higher resolution. This coarse-to-fine design allows the model to first estimate large motions at low resolution, where displacements are easier to track, and then progressively refine local details at higher resolutions.

3.4 Spherical Multi-field Refinement

Flow-based VFI methods synthesize the intermediate frame by backward warping the input frames using bilateral flows, combined with an occlusion mask and a residual term:

$$\mathbf{I}_t = \mathbf{M} \odot \mathcal{W}(\mathbf{I}_0, \mathbf{F}_{t \rightarrow 0}) + (1 - \mathbf{M}) \odot \mathcal{W}(\mathbf{I}_1, \mathbf{F}_{t \rightarrow 1}) + \mathbf{R}. \quad (7)$$

where $\mathbf{M}$ is the occlusion mask, $\mathbf{R}$ is the residual term, $\mathcal{W}$ is backward warping operation, and $\odot$ denotes element-wise multiplication. Predicting only a single pair of flows limits the possible solutions in occluded or highly distorted regions [13]. Therefore, the last decoder of each branch does not output only one flow field, occlusion mask, and residual term. Instead, it predicts N candidate sets of flow fields, occlusion masks and residuals. Each candidate set is then used in the frame synthesis formulation above to produce one intermediate frame candidate. In our implementation, we set $N = 5$, so each branch produces five

intermediate frame candidates. Different from PriOr-Flow, where the orthogonal view is mainly used to refine solely the flow, our orthogonal branch also contributes directly to image synthesis. The primitive and orthogonal branches focus on different distortion characteristics of panoramic imagery, so their candidate frames provide complementary cues. In the orthogonal branch, the network predicts $\{\mathbf{I}_t^{o,1}, \dots, \mathbf{I}_t^{o,N}\}$. These frames are then rotated back to the primitive spherical view:

$$\mathbf{I}_t^{o2p,n} = \mathbf{T}_o^p(\mathbf{I}_t^{o,n}), \quad n = 1, \dots, N. \quad (8)$$

After alignment, the primitive candidates $\{\mathbf{I}_t^{p,1}, \dots, \mathbf{I}_t^{p,N}\}$ and the rotated orthogonal candidates $\{\mathbf{I}_t^{o2p,1}, \dots, \mathbf{I}_t^{o2p,N}\}$ are jointly fused to produce the final interpolated frame:

$$\mathbf{I}_t = \mathcal{F}_{\text{fuse}}\left(\{\mathbf{I}_t^{p,1}, \dots, \mathbf{I}_t^{p,N}\}, \{\mathbf{I}_t^{o2p,1}, \dots, \mathbf{I}_t^{o2p,N}\}\right), \quad (9)$$

where $\mathbf{I}_t^{p,n}$ denotes the n -th candidate from the primitive branch and $\mathbf{I}_t^{o2p,n}$ denotes the corresponding candidate from the orthogonal branch after inverse rotation. $\mathcal{F}_{\text{fuse}}(\cdot)$ is implemented as a lightweight two-layer convolutional fusion head. This refinement aggregates complementary information across both spherical views.

3.5 Loss Functions

We supervise intermediate frame synthesis with a combination of two reconstruction error terms: census [23] and Charbonnier [3]:

$$\mathcal{L}_r = \mathcal{L}_{\text{cen}} + \lambda \mathcal{L}_{\text{char}}. \quad (10)$$

For structure preservation under illumination changes, $\mathcal{L}_{\text{cen}}$ [23] computes a soft Hamming distance between census-transformed patches (size 7×7) of $\hat{\mathbf{I}}_t$ and $\mathbf{I}_t^{gt}$. We also adapt a Charbonnier reconstruction term because it is a smooth, robust alternative to ℓ_1 and is less sensitive to outliers in occluded and highly distorted regions. In the original Charbonnier formulation [3], all pixels contribute equally to the loss. However, for spherical images, distortions near the poles often lead to significantly larger errors compared to equatorial regions. Consequently, the loss can be dominated by polar regions, causing the model to prioritize minimizing polar errors while underemphasizing central regions that contain the majority of visual information. To mitigate this effect, we propose a spherical-weighted Charbonnier loss with a latitude-dependent weight defined as:

$$\mathcal{L}_{\text{char}} = \sum_{\mathbf{p}} \rho\left(\omega(\mathbf{p}) \odot \left(\hat{\mathbf{I}}_t(\mathbf{p}) - \mathbf{I}_t^{gt}(\mathbf{p})\right)\right), \text{ and } \rho(x) = (x^2 + \epsilon^2)^\alpha, \quad (11)$$

where $\alpha = 0.5$, $\epsilon = 10^{-3}$. The weight function is defined as:

$$\omega(\mathbf{p}) = \cos(\theta_{\mathbf{p}}), \quad \theta_{\mathbf{p}} = \pi \left(\frac{v_{\mathbf{p}}}{H} - \frac{1}{2} \right), \quad (12)$$

where $v_{\mathbf{p}}$ is the vertical pixel coordinate of $\mathbf{p}$ and H is image height. This weighting reduces the dominance of heavily distorted polar regions while preserving stronger supervision near the equator. Please see the Supplementary Material for the details of the impact of this loss design and differences to AMT.

4 Experiments

4.1 Benchmarks and Evaluation Metrics

We evaluate SVI360 on two synthetic omnidirectional datasets (FlowScape [21] and Flow360 [1]) and two real-world datasets (ODV360 [2] and 360VFI [15]). The synthetic datasets provide dense ground-truth optical flow, making them well suited for both qualitative and quantitative evaluation of optical flow accuracy and frame interpolation quality. ODV360 and 360VFI consist of real-world 360° videos captured in unconstrained environments with complex object motion. They serve as challenging benchmarks for assessing the robustness of interpolation methods in realistic scenarios.

- **FlowScape**: middle-frame interpolation benchmark with 4,900 training and 1,372 testing triplets.
- **Flow360**: arbitrary-timestep interpolation benchmark with 9-frame tuples (1,286 training, 712 testing).
- **ODV360**: real-world triplets split into 18,344/1,894/1,894 for train/val/test.
- **360VFI**: real-world benchmark with 10,074 training and 930 testing triplets, with test subsets Easy/Middle/Hard/Extreme by motion magnitude.

We measure interpolation quality using standard perspective-image metrics, including PSNR and SSIM, and assess optical-flow accuracy with EPE and AE. We also consider spherical-aware metrics: Spherical end-point error (SEPE [14]) and weighted-to-spherically-uniform image-quality metrics, namely WS-PSNR [26] and WS-SSIM [33]. These weighted quality metrics account for the non-uniform spherical distortion, allowing a fairer assessment of each model’s true performance by reducing the bias from larger errors near the poles. For optical flow, SEPE is defined as the average geodesic distance between predicted and ground-truth end points of flow vectors on the unit sphere.

4.2 Implementation Details

All experiments were conducted using two H100 GPUs (80 GB). Models are optimized with AdamW with a cosine learning-rate schedule that decays from 2×10^{-4} to 2×10^{-5} over 300 epochs with a batch size of 16. The loss weight λ is fixed to 1. Data augmentation is done with horizontal and vertical flips, color jitter, and temporal reversal (swapping I_0/I_1), and a random erasing module that masks regions in the third frame to mimic occlusions caused by optical-flow warping. Besides that, we also apply spherical-specific augmentation method: random roll, pitch, yaw rotations, to create diverse view angles while preserving scene content. These augmentation methods better expose the model to occluded regions, wrap-around discontinuities, and viewpoint variations.

Table 1: Results on FlowScape. Higher is better for PSNR/WS-PSNR/SSIM/WS-SSIM; lower is better for EPE/SEPE/AE. For each group, the best result is shown in **bold**, and the second best is underlined.

Method	PSNR	WS-PSNR	SSIM	WS-SSIM	EPE	SEPE	AE	Params (M)	Latency (ms/f)
SuperSloMo	33.11	32.48	0.9034	0.9045	15.79	23.74	29.76	19.8	62
FILM	35.98	35.84	0.9369	0.9282	14.65	20.26	23.19	34.4	393
RIFE	35.97	35.66	0.9540	0.9438	15.16	21.70	23.56	9.8	26
VFIFormer	36.90	36.78	0.9614	0.9546	14.07	20.12	22.84	24.1	1293
IFRNet-L	37.03	36.58	0.9651	0.9564	12.64	<u>15.86</u>	<u>20.23</u>	19.7	79
AMT-G	<u>38.05</u>	<u>37.40</u>	<u>0.9704</u>	<u>0.9622</u>	12.26	25.91	20.90	30.6	250
SVI360 (Ours)	38.68	37.73	0.9754	0.9647	<u>12.36</u>	15.22	19.74	50.1	365

4.3 Comparison and Results

We compare SVI360 against SuperSloMo [8], FILM [20], RIFE [7], VFIFormer [22], IFRNet-L [11], and AMT-G [13]. For fair comparison, these methods are re-trained on our main benchmarks (FlowScape, Flow360, ODV360). We report PSNR, SSIM, WS-PSNR, and WS-SSIM on all datasets, while optical flow metrics (EPE, AE, and SEPE) are reported only on FlowScape. We do not retrain 360VFI [15] because reproducible training code is unavailable. On the 360VFI dataset we train our method and the strong baseline AMT-G, then compare against published baselines, including 360VFI [15], DQBC [31], EMA-VFI [30], EBME [10], UPR-Net [9], and IFRNet [11], under the official benchmark split.

FlowScape – Middle timestep interpolation ($t = 0.5$).

In this setting, we report both interpolation quality and optical flow accuracy in Tab. 1. Overall, SVI360 consistently outperforms the competitors across most evaluated metrics. AMT-G remains slightly better on EPE, but SVI360 provides a stronger overall balance between interpolation fidelity and spherical motion consistency. Compared with the second-best method, SVI360 surpasses it by 0.63 dB in PSNR and 0.33 dB in WS-PSNR. Beyond the global gains, the improvements in WS-PSNR and WS-SSIM indicate better perceptual quality over the full panoramic frame under distortion-aware weighting, rather than gains limited to specific regions. AMT-G remains a strong competitor, but it uses explicit flow supervision, whereas SVI360 does not rely on this auxiliary supervision. Earlier baselines show larger degradation on panoramic data. Due to its dual-branch design, SVI360 uses more parameters than most baselines, yet it still maintains a reasonable latency, yielding a competitive accuracy-efficiency tradeoff.

As shown in Fig. 4, SVI360 also preserves finer structural details and sharper motion boundaries, with fewer ghosting and blending artifacts than competing methods, especially near highly distorted polar regions. Benefiting from multi-view priors, where complementary views refine each other, SVI360 better maintains fine object structures at the poles while keeping high-quality reconstruction in other regions. For optical flow, SVI360 performs on par with strong base-




Fig. 4: Qualitative comparison on FlowScape. For each method, the upper image is the interpolated middle frame and the lower image is the estimated optical flow. The values shown below images report PSNR/SSIM for interpolation and EPE/AE for the flow of each image sample (not dataset-level metrics).

lines such as AMT-G and IFRNet-L, while other methods show more severe degradation in difficult areas with large pixel displacement and occlusion. It is also worth noting that estimated motion does not need to be maximally similar to the ground-truth flow at every local detail. As discussed in [29], highly precise optical flow estimation can be intractable and even suboptimal for a specific downstream video task. Similar to the design philosophy of AMT [13] and IFRNet [11], our objective is to estimate flow that is globally consistent with ground truth while allowing task-beneficial local diversity. The predicted flow of SVI360 satisfies this objective. Overall, the qualitative and quantitative comparison suggests that SVI360 improves both visual reconstruction fidelity and motion consistency on FlowScape.

Flow360 – Arbitrary timestep interpolation.

Using 9-frame tuples, target timestamps are $t \in \{1/8, 2/8, \dots, 7/8\}$. Methods marked with (*) can only predict the center frame at $t = 0.5$. For these methods, we use recursive inference to obtain $8\times$ interpolation: first predict $I_{0.5}$ from (I_0, I_1) , then predict $I_{0.25}$ from $(I_0, I_{0.5})$ and $I_{0.75}$ from $(I_{0.5}, I_1)$, and continue similarly to obtain the remaining intermediate frames. On Flow360 (Table 2), SVI360 achieves the best performance across all interpolation metrics, showing strong robustness for arbitrary-timestep prediction. It maintains at least a 0.25 dB PSNR gain against strong competitors. FILM and VFIFormer are ad-

Table 2: Results on Flow360 and ODV360. Higher is better for PSNR/SSIM/WS-PSNR/WS-SSIM. For each dataset, the best result is shown in **bold**, and the second best is underlined. (*) means the methods can only predict the center frame.

Method	Flow360				ODV360			
	PSNR	WS-PSNR	SSIM	WS-SSIM	PSNR	WS-PSNR	SSIM	WS-SSIM
SuperSloMo	31.43	31.24	0.9481	0.9440	26.41	26.95	0.8816	0.8689
FILM (*)	32.73	32.14	0.9514	0.9451	27.16	27.81	0.9059	0.8703
RIFE	32.83	32.23	0.9631	0.9515	28.74	29.51	0.9243	0.8992
VFIFormer (*)	32.54	31.93	0.9614	0.9502	27.85	28.36	0.9116	0.8725
IFRNet-L	33.16	32.60	0.9656	0.9529	28.18	28.69	0.9167	0.8830
AMT-G	<u>33.77</u>	<u>33.22</u>	<u>0.9708</u>	<u>0.9594</u>	<u>29.07</u>	<u>29.75</u>	<u>0.9303</u>	<u>0.9067</u>
Ours	34.02	33.30	0.9752	0.9615	29.25	29.93	0.9332	0.9105




Fig. 5: Qualitative comparison on ODV360. The highlighted top region marks a challenging distorted area.

ditionally affected by recursive inference for non-center timestamps, which can accumulate errors and reduce quality. Overall, these results confirm that SVI360 is very effective for multiple frame interpolation.

ODV360 – Middle frame interpolation in real-world scenes.

The model must remain robust to real capture artifacts. On the real-world ODV360 benchmark (Table 2), SVI360 also achieves the best results, indicating strong generalization to unconstrained panoramic scenes with illumination variation and complex non-rigid motion. As can be seen in Fig. 5, SVI360 yields the highest per-image PSNR/SSIM and better preserves the ceiling structure in the top region, which is consistent with the GT image.

360VFI – Interpolation across varying difficulty levels.

We train our method and strong baseline AMT-G on the 360VFI training set and compare it with published results reported in [15] and other prior meth-

Table 3: Quantitative comparison on the 360VFI benchmark across increasing difficulty levels. Higher is better for both metrics.

Method	Easy	Middle	Hard	Extreme
	WS-PSNR/WS-SSIM	WS-PSNR/WS-SSIM	WS-PSNR/WS-SSIM	WS-PSNR/WS-SSIM
IFRNet [11]	33.90/ <u>0.9538</u>	28.89/0.9047	27.74/0.8724	24.80/0.7999
DQBC [31]	33.51/0.9449	27.35/0.8776	26.02/0.8329	23.26/0.7410
EMA-VFI [30]	33.40/0.9529	28.37/0.9024	27.64/0.8817	24.89/0.8161
EBME [10]	33.48/0.9447	27.35/0.8772	25.97/0.8318	23.23/0.7403
UPR-Net [9]	33.45/0.9535	28.42/0.9024	27.72/0.8816	24.85/0.8157
360VFI [15]	33.95/0.9537	28.96/0.9060	27.81/0.8879	25.63/0.8517
AMT-G [13]	34.84 /0.9726	29.74/0.9376	28.58/0.9209	25.55/0.8817
Ours	34.81/ 0.9731	29.98 / 0.9409	28.95 / 0.9264	25.92 / 0.8880

ods listed in Table 3. We observe that SVI360 achieves the best performance across all metrics under different difficulty settings. In the Easy setting (flow magnitude < 2), where pixel displacement is minimal, most methods achieve strong performance. However, starting from the Middle setting, as the flow magnitude increases, the performance of all methods declines noticeably. Although our method also experiences some degradation, the drop-off is less significant than that of 360VFI and AMT-G, demonstrating stronger robustness to large motion fields. Qualitative results comparing SVI360 and AMT-G on the 360VFI benchmark are shown in Fig. 6. Since the official implementation of 360VFI is not yet publicly available, we could not generate its corresponding qualitative results for comparison.

4.4 Ablation Study

All ablation models are trained and evaluated on FlowScape dataset.

1) *Ablation on the Loss Functions.* We evaluate the contribution of the spherical-weighted Charbonnier and the census term. This analysis quantifies how each loss element affects interpolation quality on spherical data.

2) *Ablation on SVI360 Architecture.* We analyze the impact of key architectural modules, including the Spherical Refiner (SR) block and the Spherical Multi-Field Refinement module. For the latter, instead of using two branches as in Eq. (9), we remove orthogonal-branch candidates and fuse only primitive-branch candidates, i.e., $\mathbf{I}_t = \mathcal{F}_{\text{fuse}}(\{\mathbf{I}_t^{p,1}, \dots, \mathbf{I}_t^{p,N}\})$. This setting isolates the contribution of orthogonal-branch multi-field cues. Together with the SR-block removal, it shows the effect of each component on overall interpolation performance.

Table 4 confirms that each design choice contributes to the final performance. In the loss ablation, removing $\mathcal{L}_{\text{char}}$ causes a clear performance drop, showing that the Charbonnier term is the main reconstruction driver. Without $\mathcal{L}_{\text{cen}}$ term, training is more unstable, therefore, we recommend using both losses together. We can also notice the relevance of the spherical weighted Charbonnier loss to preserve image details of regions near the equator, as shown in Fig. 7.




Fig. 6: Qualitative results on the 360VFI benchmark. Red bounding boxes highlight regions where AMT-G fails while SVI360 produces more accurate results.

Table 4: Ablation on loss functions and SVI360 architecture. Higher is better.

Settings	PSNR	WS-PSNR	SSIM	WS-SSIM
Loss Functions				
w/o sph. Charb. Loss	33.99	33.85	0.9099	0.9471
w/o Census Loss	38.34	37.63	0.9734	0.9644
Full loss	38.68	37.73	0.9754	0.9647
SVI360 Architecture				
w/o sph. multi-field Refinement	38.08	37.29	0.9715	0.9627
w/o Spherical Refiner (SR)	38.59	37.65	0.9748	0.9640
Full model	38.68	37.73	0.9754	0.9647

For architecture ablation, using only primitive-branch candidates reduces all metrics compared with the full model, which demonstrates the benefit of cross-branch multi-field fusion for spherical interpolation quality. Removing the Spherical Refiner also degrades the interpolation metrics, showing that the aligned orthogonal cues such as estimated flow and intermediate features help improve the synthesized intermediate frame. To verify that improvements do not simply come from a larger number of parameters, we also evaluate a lightweight variant, LiteSVI360, in the Supplementary Material. LiteSVI360 uses smaller AMT backbones and has substantially fewer parameters, while maintaining competitive



Fig. 7: Qualitative comparison of reconstruction losses on FlowScape. The images show a center-region zoom of the ground truth (left) and reconstructions produced with standard Charbonnier loss (center) and spherical weighted Charbonnier loss (right). Please notice the spherical Charbonnier loss better preserves the firetruck cabin structure.

performance, showing that the gains mainly come from the proposed spherical dual-branch refinement and fusion design.

Discussion and limitations. The dual-branch equivariance feature learning design has shown effectiveness in 360° optical flow [14], and our results show that the same principle also transfers to interpolation. This suggests a broader potential for spherical reconstruction tasks, including segmentation and depth estimation. Limitations: Prior VR findings report that frame rates above 120 fps can improve user comfort and reduce simulator sickness [28]. In contrast, our model currently runs at 365 ms/frame (nearly 3 fps), so it is not suitable for real-time VR interpolation. However, many 360° contents are not rendered online. In these cases, SVI360 is better aligned with offline pipelines such as post-production video enhancement and compression-oriented preprocessing where quality is prioritized over latency.

5 Conclusion

This paper presents SVI360, a dual-branch omnidirectional video interpolation method that exploits complementary priors from primitive and orthogonal spherical views. By combining cross-view refinement with distortion-aware spherical reconstruction, SVI360 improves both interpolation fidelity and motion consistency in panoramic scenes. Extensive experiments on FlowScape, Flow360, ODV360 and 360VFI benchmarks show that SVI360 achieves strong and consistent performance, with clear gains in interpolation quality and robust behavior in challenging real-world scenes. High-quality spherical video interpolation remains fundamentally more challenging than perspective video interpolation, and pushing visual fidelity in omnidirectional domains currently requires accepting higher computational cost. Future work will focus on reducing model complexity to meet real-time constraints of VR downstream tasks.

Acknowledgements. This work was funded by the project ANR TSIA DEVIN (ANR-23-IAS2-0001), ANER MOVIS from “Region Bourgogne-Franche-Comte” and ANR MANYVIS (ANR-23-CE23-0003-01), to whom we are grateful. We thank the access to the HPC computational resources of IDRIS under the allocation AD011016829.

References

1. Bhandari, K., Duan, B., Liu, G., Latapie, H., Zong, Z., Yan, Y.: Learning omnidirectional flow in 360° video via siamese representation. In: Avidan, S., Brostow, G., Cissé, M., Farinella, G.M., Hassner, T. (eds.) *Computer Vision – ECCV 2022*. pp. 557–574. Springer Nature Switzerland, Cham (2022)
2. Cao, M., Mou, C., Yu, F., Wang, X., Zheng, Y., Zhang, J., Dong, C., Li, G., Shan, Y., Timofte, R., Sun, X., Li, W., Zhang, Z., Sheng, X., Chen, B., Ma, H., Cheng, M., Zhao, S., Cui, W., Xu, T., Li, C., Bao, L., Sun, H., Huang, H., Zhou, X., Ai, Y., He, R., Wu, R., Yang, Y., Zhang, Z., Zhang, S., Li, J., Chen, Y., Ren, D., Zuo, W., Wang, Q., Yang, H.H., Chen, Y.C., Huang, Z.K., Chen, W.T., Chiang, Y.C., Chang, H.E., Chen, I.H., Hsieh, C.H., Kuo, S.Y., Zhang, Z., Zhang, J., Wang, Y., Cui, S., Huang, J., Zhu, L., Tian, S., Yu, W., Luo, B.: Ntire 2023 challenge on 360deg omnidirectional image and video super-resolution: Datasets, methods and results. In: *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR) Workshops*. pp. 1731–1745 (June 2023)
3. Charbonnier, P., Blanc-Feraud, L., Aubert, G., Barlaud, M.: Two deterministic half-quadratic regularization algorithms for computed imaging. In: *Proceedings of 1st International Conference on Image Processing*. vol. 2, pp. 168–172 vol.2 (1994)
4. Cohen, T.S., Geiger, M., Köhler, J., Welling, M.: Spherical cnns. In: *International Conference on Learning Representations (ICLR)* (2018)
5. Coors, B., Condurache, A.P., Geiger, A.: Spherenet: Learning spherical representations for detection and classification in omnidirectional images. In: *Proceedings of the European Conference on Computer Vision (ECCV)* (September 2018)
6. Fernandez-Labrador, C., Fácil, J.M., Pérez-Yus, A., Démonceaux, C., Civera, J., Guerrero, J.: Corners for layout: End-to-end layout recovery from 360 images. *IEEE Robotics and Automation Letters* **5**(2), 1255–1262 (2020)
7. Huang, Z., Zhang, T., Heng, W., Shi, B., Zhou, S.: Real-time intermediate flow estimation for video frame interpolation. In: Avidan, S., Brostow, G., Cissé, M., Farinella, G.M., Hassner, T. (eds.) *Computer Vision – ECCV 2022*. pp. 624–642. Springer Nature Switzerland, Cham (2022)
8. Jiang, H., Sun, D., Jampani, V., Yang, M.H., Learned-Miller, E., Kautz, J.: Super slo-mo: High quality estimation of multiple intermediate frames for video interpolation. In: *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition (CVPR)* (June 2018)
9. Jin, X., Wu, L., Chen, J., Chen, Y., Koo, J., Hahm, C.h.: A unified pyramid recurrent network for video frame interpolation. In: *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*. pp. 1578–1587 (June 2023)
10. Jin, X., Wu, L., Shen, G., Chen, Y., Chen, J., Koo, J., Hahm, C.h.: Enhanced bi-directional motion estimation for video frame interpolation. In: *Proceedings of the IEEE/CVF Winter Conference on Applications of Computer Vision (WACV)*. pp. 5049–5057 (January 2023)
11. Kong, L., Jiang, B., Luo, D., Chu, W., Huang, X., Tai, Y., Wang, C., Yang, J.: Ifrnet: Intermediate feature refine network for efficient frame interpolation. In: *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*. pp. 1969–1978 (June 2022)
12. Lee, H., Kim, T., Chung, T.y., Pak, D., Ban, Y., Lee, S.: Adacof: Adaptive collaboration of flows for video frame interpolation. In: *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)* (June 2020)

13. Li, Z., Zhu, Z.L., Han, L.H., Hou, Q., Guo, C.L., Cheng, M.M.: Amt: All-pairs multi-field transforms for efficient frame interpolation. In: Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR). pp. 9801–9810 (June 2023)
14. Liu, L., Feng, M., Cheng, J., Xiang, J., Zhu, X., Yang, X.: Prior-flow: Enhancing primitive panoramic optical flow with orthogonal view. In: Proceedings of the IEEE/CVF International Conference on Computer Vision (ICCV). pp. 5326–5336 (October 2025)
15. Lu, W., Hu, M., Qiu, Y., Liao, L., Wang, Z.: 360vfi: A dataset and benchmark for omnidirectional video frame interpolation (2024)
16. Meyer, S., Djelouah, A., McWilliams, B., Sorkine-Hornung, A., Gross, M., Schroers, C.: Phasenet for video frame interpolation. In: Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition (CVPR) (June 2018)
17. Meyer, S., Wang, O., Zimmer, H., Grosse, M., Sorkine-Hornung, A.: Phase-based frame interpolation for video. In: Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition (CVPR) (June 2015)
18. Niklaus, S., Mai, L., Liu, F.: Video frame interpolation via adaptive convolution. In: Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition (CVPR) (July 2017)
19. Niklaus, S., Mai, L., Liu, F.: Video frame interpolation via adaptive separable convolution. In: Proceedings of the IEEE International Conference on Computer Vision (ICCV) (Oct 2017)
20. Reda, F., Kontkanen, J., Tabellion, E., Sun, D., Pantofaru, C., Curless, B.: Film: Frame interpolation for large motion. In: Avidan, S., Brostow, G., Cissé, M., Farinella, G.M., Hassner, T. (eds.) *Computer Vision – ECCV 2022*. pp. 250–266. Springer Nature Switzerland, Cham (2022)
21. Shi, H., Zhou, Y., Yang, K., Yin, X., Wang, Z., Ye, Y., Yin, Z., Meng, S., Li, P., Wang, K.: Panoflow: Learning 360° optical flow for surrounding temporal understanding. *Trans. Intell. Transport. Sys.* **24**(5), 5570–5585 (May 2023)
22. Shi, Z., Xu, X., Liu, X., Chen, J., Yang, M.H.: Video frame interpolation transformer. In: Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR). pp. 17482–17491 (June 2022)
23. Simon, M., Junhwa, H., Stefan, R.: Unflow: Unsupervised learning of optical flow with a bidirectional census loss. *Proceedings of the AAAI Conference on Artificial Intelligence* **32**(1) (04 2018)
24. Su, Y.C., Grauman, K.: Learning spherical convolution for fast features from 360°imagery. In: Guyon, I., Luxburg, U.V., Bengio, S., Wallach, H., Fergus, R., Vishwanathan, S., Garnett, R. (eds.) *Advances in Neural Information Processing Systems*. vol. 30. Curran Associates, Inc. (2017)
25. Su, Y.C., Grauman, K.: Kernel transformer networks for compact spherical convolution. In: Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR) (June 2019)
26. Sun, Y., Lu, A., Yu, L.: Weighted-to-spherically-uniform quality evaluation for omnidirectional video. *IEEE Signal Processing Letters* **24**(9), 1408–1412 (2017)
27. Teed, Z., Deng, J.: Raft: Recurrent all-pairs field transforms for optical flow. In: Vedaldi, A., Bischof, H., Brox, T., Frahm, J.M. (eds.) *Computer Vision – ECCV 2020*. pp. 402–419. Springer International Publishing, Cham (2020)
28. Wang, J., Shi, R., Zheng, W., Xie, W., Kao, D., Liang, H.N.: Effect of frame rate on user experience, performance, and simulator sickness in virtual reality. *IEEE Transactions on Visualization and Computer Graphics* **29**(5), 2478–2488 (2023)

29. Xue, T., Chen, B., Wu, J., Wei, D., Freeman, W.T.: Video enhancement with task-oriented flow. *International Journal of Computer Vision (IJCV)* **127**(8), 1106–1125 (2019)
30. Zhang, G., Zhu, Y., Wang, H., Chen, Y., Wu, G., Wang, L.: Extracting motion and appearance via inter-frame attention for efficient video frame interpolation. In: *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*. pp. 5682–5692 (2023)
31. Zhou, C., Liu, J., Tang, J., Wu, G.: Video frame interpolation with densely queried bilateral correlation. In: Elkind, E. (ed.) *Proceedings of the Thirty-Second International Joint Conference on Artificial Intelligence, IJCAI-23*. pp. 1786–1794. *International Joint Conferences on Artificial Intelligence Organization* (8 2023), main Track
32. Zhou, C., Li, Z., Liu, Y.: A measurement study of oculus 360 degree video streaming. In: *Proceedings of the 8th ACM on Multimedia Systems Conference*. pp. 27–37. *MMSys’17*, Association for Computing Machinery, New York, NY, USA (2017)
33. Zhou, Y., Yu, M., Ma, H., Shao, H., Jiang, G.: Weighted-to-spherically-uniform ssim objective quality evaluation for panoramic video. In: *2018 14th IEEE International Conference on Signal Processing (ICSP)*. pp. 54–57 (2018)
34. Zhu, T., Ren, D., Wang, Q., Wu, X., Zuo, W.: Generative inbetweening through frame-wise conditions-driven video generation. In: *Proceedings of the Computer Vision and Pattern Recognition Conference (CVPR)*. pp. 27968–27978 (June 2025)