
PoseCrafter: Extreme Pose Estimation with Hybrid Video Synthesis

Qing Mao^{1,2*} Tianxin Huang³ Yu Zhu¹ Jinqiu Sun⁴ Yanning Zhang^{1†} Gim Hee Lee^{2†}

¹School of Computer Science, Northwestern Polytechnical University

²School of Computing, National University of Singapore

³School of Computing and Data Science, The University of Hong Kong

⁴School of Astronautics, Northwestern Polytechnical University

maoqing@mail.nwpu.edu.cn ynzhang@nwpu.edu.cn gimhee.lee@nus.edu.sg

<https://github.com/maoqingsunny/PoseCrafter>

Abstract

Pairwise camera pose estimation from sparsely overlapping image pairs remains a critical and unsolved challenge in 3D vision. Most existing methods struggle with image pairs that have small or no overlap. Recent approaches attempt to address this by synthesizing intermediate frames using video interpolation and selecting key frames via a self-consistency score. However, the generated frames are often blurry due to small overlap inputs, and the selection strategies are slow and not explicitly aligned with pose estimation. To solve these cases, we propose **Hybrid Video Generation (HVG)** to synthesize clearer intermediate frames by coupling a video interpolation model with a pose-conditioned novel view synthesis model, where we also propose a **Feature Matching Selector (FMS)** based on feature correspondence to select intermediate frames appropriate for pose estimation from the synthesized results. Extensive experiments on Cambridge Landmarks, ScanNet, DL3DV-10K, and NAVI demonstrate that, compared to existing SOTA methods, PoseCrafter can obviously enhance the pose estimation performances, especially on examples with small or no overlap.

1 Introduction

Accurately estimating the relative pose between two images is a fundamental task in 3D vision, critical for enabling autonomous navigation, robotics, and augmented reality. While modern feature-based [1, 2] or learning-based pipelines [3, 4] perform well when input views share sufficient overlap, their performance deteriorates sharply in the small or no overlap case, owing to the fundamental scarcity of reliable feature correspondences.

To address this issue, InterPose [5] introduced a two-stage strategy: intermediate views are generated using a video interpolation model, and enriched frames selected with a frame selection strategy are provided to the pose estimation model. Although InterPose demonstrated obvious improvements in extreme cases with small overlaps, it faces two core limitations. First, the quality of generated intermediate frames remains a bottleneck. Free and open-source models like DynamiCrafter [6] can generate plausible interpolations near the input image pair but often fail to maintain geometric consistency in the central frames, likely due to the limited overlap between input views. Commercial models like Runway [7] and Luma Dream Machine [8] offer sharper results but still suffer from blur and drift in challenging scenes, in addition to incurring high inference costs. Second, InterPose relies

*Work fully done while first author is a visiting PhD student at the National University of Singapore.

†Corresponding author.

on a self-consistency score mechanism to select frames, which filters out informative frames across multiple generated video clips. This purely statistical selection strategy is slow and not explicitly aligned with the objectives of pose estimation.

To address these challenges, we propose **PoseCrafter**—a training-free framework for estimating camera poses from input image pairs with small or no overlap. PoseCrafter comprises two key components: Hybrid Video Generation (HVG) and Feature Matching Selector (FMS). In HVG, we first use the pre-trained video interpolation model DynamiCrafter [6] to synthesize a coarse video and identify a few reliable “relay” frames. These selected frames are then used to estimate the initial pose and generate high-fidelity intermediate views through a pose-conditioned novel view synthesis model, ViewCrafter [9]. In this work, we find that selecting only the one frame immediately after the start frame and the one frame just before the end frame yields the most effective relay frames. Including additional frames tends to degrade performance due to increased blurriness in the central synthesized frames. More analysis can be found in Sec. 3.2.

In FMS, instead of using statistical scores for frame selection, we propose a simple but effective solution, evaluating synthesized frames based on their feature correspondences with the input image pair to determine their suitability for pose estimation. In detail, we extract local descriptors from each candidate frame and match them to the input image pair, compute RANSAC [10] inlier counts, and select the top k frames with the highest total inliers.

Extensive experiments on common benchmarks, including Cambridge Landmarks [11], ScanNet [12], DL3DV-10K [13], and NAVI [14] show that PoseCrafter can obviously improve the accuracy of pose estimation on extreme pose image pairs with small or no overlaps, without any requirements for additional training or ground-truth supervision. In summary, our contributions are as follows:

- We propose **Hybrid Video Generation(HVG)** to synthesize high-fidelity intermediate frames, by effectively coupling a video interpolation model and a pose-conditioned novel view synthesis model.
- We develop a **Feature Matching Selector(FMS)** to deterministically select the most informative frames suitable for pose estimation, eliminating the need for expensive statistical self-consistency scoring in existing works.
- We demonstrate that, by integrating HVG and FMS, our proposed **PoseCrafter** can achieve SOTA performances across four challenging benchmarks for extreme pose estimation.

2 Related Work

Extreme Pose Estimation Traditional relative pose pipelines combine local feature matching [1, 2], RANSAC-based [10] five-point or eight-point solvers, and bundle adjustment [15]. These methods perform well when images share substantial overlap but fail under extreme cases. While deep networks have strengthened individual components [16–19], they still rely on the large overlap. To address cases with small or no overlap, several approaches [20–22, 3] have been proposed. Yang et al. [20] alternate RGB-D scene completion with pose regression, iteratively fusing geometry from both scans; Cai et al. [21] construct dense correlation volumes over synthetic video clips to infer large rotations; and DUS3R [3] regresses pixel-aligned point maps to jointly recover depth, reconstruction, and pose without camera calibration. Temporal modeling has also been explored. Tang et al. [23] combine attention and recurrent networks for pose prediction. However, each of these methods requires supervised data and training resources. More recently, video-diffusion techniques have been adapted for pose recovery. JOG3R [24] fine-tunes intermediate features of a pre-trained video generation model for SfM-style estimation, InterPose [5] reuses standard diffusion interpolation and a self-consistency scoring scheme to select the most reliable frames, both without retraining the video model. These studies highlight the power of synthesized intermediate views and motivate our PoseCrafter framework.

Generative Video Models Early video synthesis methods relied on generative adversarial networks (GANs) [25] and autoregressive models. GAN-based approaches such as TGAN [26], MoCoGAN [27], and StyleGAN-V [28] extend image GANs to model temporal coherence, but often suffer from frame inconsistencies and require complex architectures. Autoregressive techniques [29, 30] generate videos frame by frame, but they incur high computational costs and error

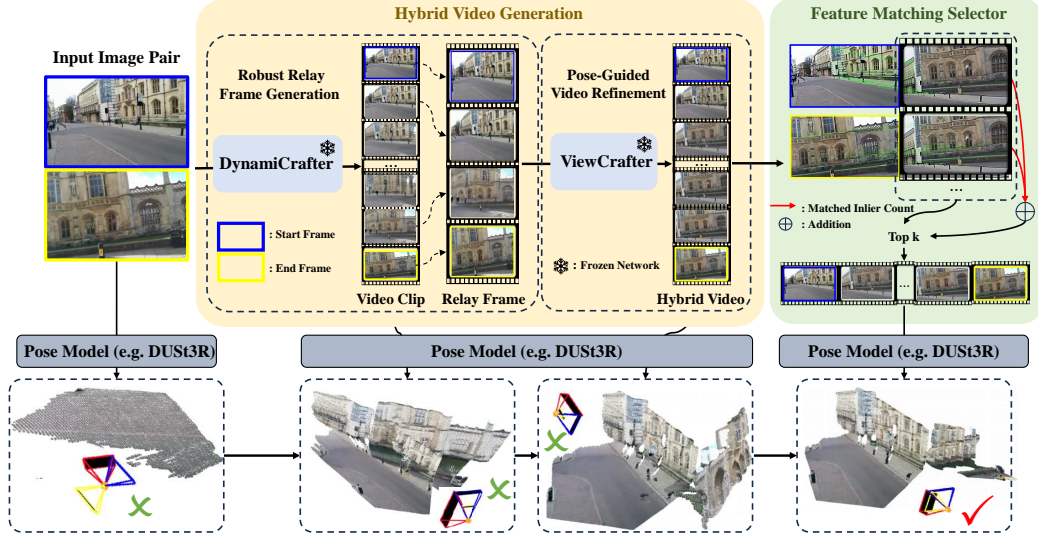


Figure 1: Overview of the PoseCrafter pipeline. Starting from an input image pair (left), the **Hybrid Video Generation** module uses DynamiCrafter to interpolate a short video clip, extracts the most reliable “relay” frames, and refines the sequence using the pose-conditioned ViewCrafter to reduce mid-sequence blur and drift. The **Feature Matching Selector** then evaluates each synthesized frame based on its feature correspondence with the input pair, selecting the most informative views for pose estimation. These selected frames are then fed into a pose estimation model (e.g. DUST3R) to estimate the final pose.

accumulation. Diffusion models have recently set a new standard for both image and video synthesis. Video-diffusion frameworks such as VDM [31], Imagen-Video [32], and Make-A-Video [33] generate high-fidelity, temporally coherent sequences via iterative denoising. ViewCrafter [9] proposed by combining video diffusion priors with coarse point-cloud cues to synthesize high-fidelity, pose-controlled novel views from one or a few images. There are also some video generation models that can be used for interpolation that use diffusion to generate smooth, geometry-aware intermediate frames. For example, DynamiCrafter [6] is based on motion-consistent editing. These diffusion-based models outperform prior GAN and autoregressive methods in both visual quality and temporal consistency, making them ideal for downstream tasks such as extreme pose estimation.

3 Our Method

As illustrated in Figure 1, given a pair of images with small or no overlap, **PoseCrafter** estimates their relative camera pose through two complementary stages. In **Hybrid Video Generation** (Section 3.2), we sequentially combine a video interpolation model, DynamiCrafter, with a pose-conditioned novel view synthesis model, ViewCrafter, to generate refined video frames that significantly reduce mid-sequence blurring and correct structural drift. In the subsequent **Feature Matching Selector** (Section 3.3), we evaluate the synthesized frames based on their feature correspondences with the input image pair and select the top k most informative frames for final pose estimation.

3.1 Preliminaries

DynamiCrafter DynamiCrafter is a publicly available diffusion model for video synthesis that employs a progressive latent denoising architecture to animate two static images into dynamic sequences [6]. Given an input image pair $(I_0, I_T \in \mathbb{R}^{H \times W \times 3})$ and an optional text prompt \mathcal{P} , it generates a dense sequence of intermediate frames using the following formula:

$$\{I_t\}_{t=0}^T = \mathcal{G}_{DC}(I_0, I_T, \mathcal{P}),$$

where \mathcal{G}_{DC} denotes the pre-trained DynamiCrafter model. The variable $t \in \mathbb{Z}^+$ denotes the frame index, where I_0 and I_T correspond to the input image pair.




Figure 2: Examples of mid-frame blurriness and confidence degradation. (a) presents frames generated by DynamiCrafter (DC) and our proposed method. (b) highlights high-confidence regions identified by DUST3R during pose estimation. **Green** boxes mark areas affected by blurriness in the DC-generated frames, while **red** boxes indicate how these regions performed in DUST3R’s confidence maps. We observe that the corresponding regions in the DynamiCrafter results exhibit notably low confidence, whereas our method generates frames with higher confidence for DUST3R.

ViewCrafter ViewCrafter is a pose-conditioned diffusion model for novel view synthesis of static 3D scenes, which can generate high-fidelity novel views along an interpolated camera trajectory through input images following the formula:

$$\{I_t\}_{t=0}^T = \mathcal{G}_{VC}(I_0, I_T, \{C_t\}_{t=0}^T),$$

where \mathcal{G}_{VC} denotes the pre-trained ViewCrafter generator. $\{C_t\}_{t=0}^T$ denotes camera poses alongside the trajectory interpolated between the input image pair.

3.2 Hybrid Video Generation

Although InterPose [5] improves pose estimation performance by directly interpolating frames between input image pairs (e.g., using DynamiCrafter [6]), the generated sequences often exhibit noticeable blurriness and inconsistencies, particularly in the middle frames. As shown in Figure 2, these blurry regions result in significantly lower confidence during pose estimation with DUST3R [3], ultimately affecting overall performance. In contrast, ViewCrafter [9] is capable of synthesizing clear, high-fidelity frames, but it requires input images that allow for the estimation of a plausible camera trajectory. To address this issue, we propose Hybrid Video Generation (HVG), which couples DynamiCrafter and ViewCrafter. Specifically, we first use DynamiCrafter to generate initial interpolated frames and then select a few reliable “relay” frames. These relay frames help ensure a plausible camera trajectory and are subsequently used by ViewCrafter to synthesize high-quality intermediate frames. As shown in Figure 2-Our, intermediate frames generated by our HVG exhibit higher confidence scores in DUST3R [3], further confirming that HVG contributes positively to pose estimation. The whole process comprises two phases: Robust Relay-Frame Generation, and Pose-Guided Video Refinement.

Robust Relay-Frame Generation In this stage, we use the pre-trained DynamiCrafter model to interpolate a short video clip from the input image pair (the start frame I_0 and the end frame I_T). Although DynamiCrafter generates plausible interpolations near the input pair, it often fails to maintain geometric consistency in the central frames, weakening feature correspondences and degrading downstream pose estimation. To isolate the most reliable “relay” frames, we conducted a study (Table 1) that evaluated various subsets of interpolated frames. We found that retaining only the start/end frames and their immediate neighbors $\{I_0, I_1, I_{T-1}, I_T\}$ minimizes rotation error and

maximizes stability. Consequently, PoseCrafter forwards only these four frames as “relay” frames to the subsequent high-fidelity novel view synthesis stage.

Pose-Guided Video Refinement Given the four reliable relay frames $\{I_0, I_1, I_{T-1}, I_T\}$, we first recover their camera poses $\{C_0, C_1, C_{T-1}, C_T\}$ using the pretrained DUS3R network. We then interpolate a smooth, dense camera trajectory $\{C_t\}_{t=0}^T$ by applying spherical linear interpolation to rotations on $SO(3)$ and linear interpolation to translations in \mathbb{R}^3 . Conditioned on this trajectory and the reliable relay frames, ViewCrafter synthesizes high-fidelity intermediate views. This pose-conditioned refinement effectively removes mid-sequence blur and geometric drift, producing crisp, structurally accurate frames that markedly improve feature correspondences for final pose estimation.

Table 1: Relay-frame sampling analysis using mean rotation error (MRE↓). The setting #Frames=2 corresponds to $\{I_0, I_T\}$, #Frames=4 corresponds to $\{I_0, I_1, I_{T-1}, I_T\}$, #Frames=6 corresponds to $\{I_0, I_1, I_2, I_{T-2}, I_{T-1}, I_T\}$, and #Frames=8 corresponds to $\{I_0, I_1, I_2, I_3, I_{T-3}, I_{T-2}, I_{T-1}, I_T\}$. The case #Frames=16 uses all frames. Results indicate that #Frames=4 consistently achieves the lowest MRE and the highest stability across datasets.

Dataset	#Frames (n)				
	2	4	6	8	16
Cambridge Landmarks	20.56	14.47	16.66	16.87	17.83
ScanNet	19.67	16.23	17.03	17.16	18.56
DL3DV-10K	15.22	14.27	14.40	14.73	14.52
NAVI	7.78	6.94	7.18	9.64	10.92

3.3 Feature Matching Selector

Although our proposed Hybrid Video Generation (HVG) produces high-fidelity and visually clear intermediate frames, not all of them are equally beneficial for pose estimation, as some generated content may be inconsistent with the input image pair. InterPose [5] addresses this problem by computing a statistical self-consistency score across multiple separate estimations. However, this approach is both slow and unstable due to the inherent uncertainty introduced by repeated sampling. Moreover, as a purely statistical measure, it does not explicitly ensure that the selected frames are well-suited for pose estimation.

To overcome these drawbacks, we employ a simple but effective deterministic selection strategy based on feature matching, named as Feature Matching Selector (FMS). For each synthesized frame I_t , we extract local descriptors and match them to both the start frame I_0 and end frame I_T , then compute RANSAC inlier counts $N_0(t)$ and $N_T(t)$. Each frame is scored by

$$S(t) = N_0(t) + N_T(t),$$

and we select the top k frames whose scores exceed a preset threshold. This inlier-driven criterion avoids the need for repeated estimations and explicitly accounts for geometric consistency with the input image pair, enabling the selection of more informative frames for subsequent pose estimation.

4 Experiments

4.1 Experiment Setup

Data Preparation Following InterPose [5], we evaluate our method on four common benchmarks: Cambridge Landmarks [11], ScanNet [12], Navi [14] and DL3DV-10K[13]. For Cambridge Landmarks and ScanNet, we select test pairs by sampling images whose relative yaw difference falls into two ranges ($[50^\circ-65^\circ]$ and $[65^\circ-90^\circ]$) to evaluate performance under small and no overlap cases. For the object-centric Navi and DL3DV-10K datasets, due to the large object overlap, we adopt a single yaw range of $[50^\circ-90^\circ]$ following InterPose’s setting. All images are resized, and center-cropped to 512×320 before processing, to adapt to the setting of DynamiCrafter [6].

Evaluation Metrics Following prior work [5], we conduct evaluations using the following metrics:

Mean Rotation Error (MRE) The average geodesic distance between the predicted and the ground-truth rotation, reported in angular.

Mean Translation Error (MTE) The average geodesic distance between the predicted and the ground-truth translation, reported in angular.

Rotation and Translation Accuracy ($R(\theta), T(\theta)$) Given an angular threshold θ (e.g., 5° , 15° , 30°), the rotation accuracy $R(\theta)$ is the fraction of estimates with rotation error below θ ; the translation accuracy $T(\theta)$ is the same definition using translation angular error.

Area Under Curve (AUC_{30}) The normalized area under the rotation or translation error curve from 0° to 30° provides a compact summary of performance across the full range of angular thresholds. In this work, we report the smaller of the rotation and translation AUC_{30} values as the final metric.

4.2 Implementation Details

In the hybrid video generation stage, we first generate 16 interpolated frames between each input image pair using DynamiCrafter. From these, we select 4 frames as reliable relay frames (as described in section 3.2), which are then used by ViewCrafter to render 25-frame sequences for subsequent selection with Feature Matching Selector (FMS). In FMS, we extract 2,000 ORB keypoints per frame and compute RANSAC inlier counts with respect to both the start and end keyframes. The generated frames are then ranked based on their inlier counts, and the top $k = 6$ frames(excluding the input image pair) are selected for the final pose estimation. All experiments were conducted on a single NVIDIA RTX 6000 GPU to ensure a consistent evaluation environment.

4.3 Comparison with State-of-the-Art

To validate our method’s effectiveness, we compare it against two closely related baselines: DUST3R [3] and InterPose [5]. DUST3r estimates relative pose directly from an input image pair, while InterPose enhances pose estimation results by using videos generated by a pre-trained video model and intermediate frames selected by a self-consistency score. Since InterPose has not been publicly available, we reproduce its pipeline based on the details provided in their paper, and denote it as InterPose[†] for subsequent comparisons. To further assess the impact of video generation quality, we also introduce two variant methods: our method without feature matching selector(Ours_{w/o FMS}), and InterPose without self-consistency score (InterPose[†]_{w/o SCS}). In these cases, all of the synthesized frames are used for pose estimation, without specific selection.

4.3.1 Quantitative Comparison on PoseCrafter

Tables 2, 3, 4, and 5 present comparative performance results of PoseCrafter against other methods across four benchmark datasets. On outward-facing benchmarks (Cambridge Landmarks and ScanNet), PoseCrafter reduces mean rotation error by approximately 9.85° and increases $R@30^\circ$ by more than 10% against InerPose[†] for data with yaw changes between $[65^\circ-90^\circ]$ of Cambridge Landmarks. These improvements demonstrate our pipeline’s superiority on extreme pose estimation. Moreover, the comparison between InterPose[†]_{w/o SCS} and PoseCrafter_{w/o FMS} further confirms that our proposed hybrid video generation approach produces frames more suitable for pose estimation than the DynamiCrafter [6] used in InterPose [5].

On center-facing benchmarks (DL3DV-10K and NAVI) with larger cross-image overlaps, PoseCrafter still yields modest but consistent reductions in rotation and translation error, alongside incremental improvements in AUC_{30} . Overall, the statistical results across multiple standard benchmarks confirm that our proposed pipeline based on Hybrid Video Generation and the Feature Matching Selector, provides a stable, training-free solution across diverse scenes and baseline conditions.

4.3.2 Qualitative Comparison on PoseCrafter

To intuitively assess the effectiveness of our approach, we present qualitative comparisons with DUST3R and InterPose[†] on the previously mentioned benchmarks in Figure 3. In Figure 3, we visualize the input image pairs (with the start frame outlined in blue and the end frame in yellow), point clouds reconstructed by DUST3R using only the input pairs, intermediate videos generated by InterPose[†] and PoseCrafter, and the corresponding estimated poses and reconstructions produced by DUST3R for each video, respectively. Compared to DUST3R, which only uses input image pairs,

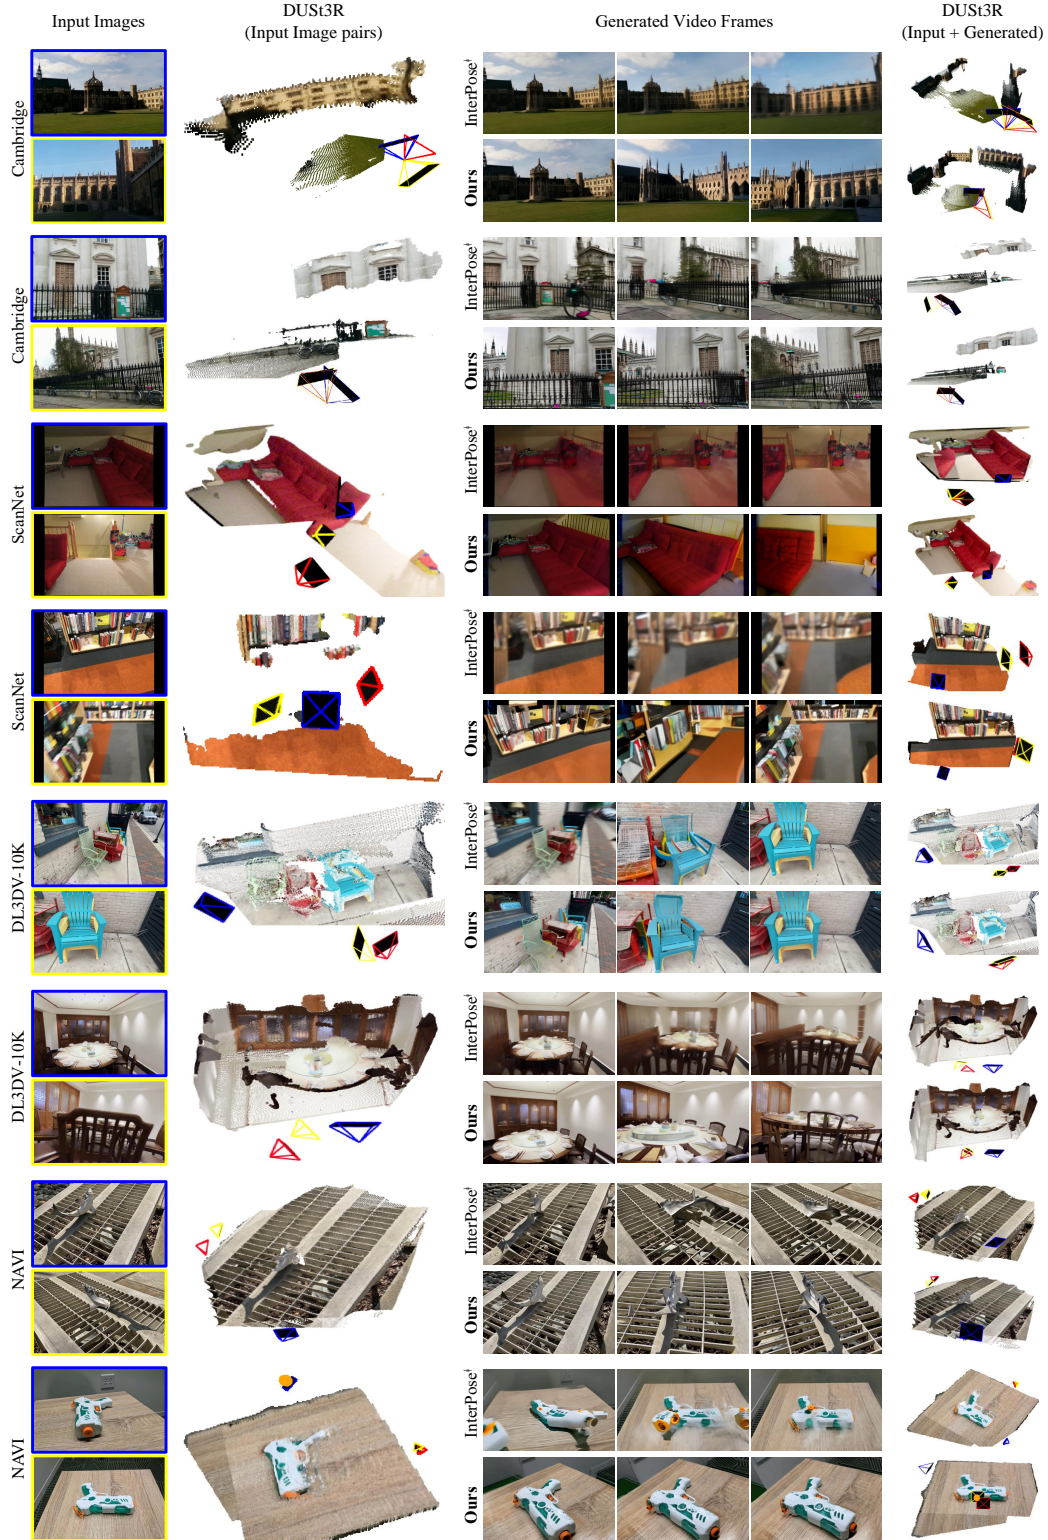


Figure 3: Qualitative comparisons on common benchmarks. Blue and yellow outlines indicate the start and end frames, respectively. In the DUST3R reconstruction results, the visualized poses in blue, yellow, and red represent the predicted pose of the start frame, predicted pose of the end frame, and the ground-truth pose of the end frame, respectively.

Table 2: Pose estimation on Cambridge Landmarks. We report rotation recall ($R@ \theta \uparrow$), translation recall ($T@ \theta \uparrow$), mean rotation error ($MRE\downarrow$), and $AUC_{30} \uparrow$.

Method	Input	Yaw range [50°-65°]					Yaw range [65°-90°]				
		$MRE\downarrow$	$R@5^\circ$	$R@15^\circ$	$R@30^\circ$	$AUC_{30} \uparrow$	$MRE\downarrow$	$R@5^\circ$	$R@15^\circ$	$R@30^\circ$	$AUC_{30} \uparrow$
DUST3R	Pair	18.14	40.34	71.25	82.99	61.98	51.24	21.67	44.67	51.67	37.93
InterPose [†] _{w/o SCS}	DynamiCrafter	16.11	42.70	75.70	87.35	65.72	42.51	30.67	42.51	61.33	47.18
InterPose [†]	DynamiCrafter	13.61	51.81	81.50	83.30	70.47	38.87	36.33	65.67	68.33	55.24
Ours _{w/o FMS}	Hybrid video	13.24	54.51	89.24	92.71	76.13	34.87	31.33	68.33	77.67	56.29
Ours	Hybrid video	11.40	55.21	89.93	93.75	77.41	29.02	36.67	71.67	78.33	60.46

Table 3: Pose estimation on ScanNet. We report rotation recall ($R@ \theta \uparrow$), translation recall ($T@ \theta \uparrow$), mean rotation error ($MRE\downarrow$), mean translation error ($MTE\downarrow$), and $AUC_{30} \uparrow$.

Yaw range	Method	Input	$R@5^\circ$	$R@15^\circ$	$R@30^\circ$	$T@5^\circ$	$T@15^\circ$	$T@30^\circ$	$MRE\downarrow$	$MTE\downarrow$	$AUC_{30} \uparrow$
50°-65°	DUST3R	Pair	43.97	74.14	79.31	25.34	52.07	78.45	19.41	25.23	47.37
	InterPose [†] _{w/o SCS}	DynamiCrafter	46.55	77.59	85.34	16.38	48.10	64.74	17.51	35.25	42.69
	InterPose [†]	DynamiCrafter	50.86	81.03	87.07	27.58	61.21	69.46	15.15	23.89	53.33
	Ours _{w/o FMS}	Hybrid video	51.72	87.07	93.10	23.28	50.62	67.07	12.38	29.02	45.53
	Ours	Hybrid video	53.45	88.79	94.83	33.62	65.52	77.69	10.77	22.14	57.03
65°-90°	DUST3R	Pair	42.05	67.05	70.45	26.59	46.59	53.86	30.82	29.99	36.50
	InterPose [†] _{w/o SCS}	DynamiCrafter	38.64	62.50	65.90	20.45	39.77	47.72	35.18	58.89	33.40
	InterPose [†]	DynamiCrafter	45.45	67.05	71.59	31.81	53.41	64.77	28.22	29.52	45.98
	Ours _{w/o FMS}	Hybrid video	46.59	76.14	82.95	23.85	48.86	57.95	22.61	35.98	41.72
	Ours	Hybrid video	50.00	77.72	84.09	37.50	63.64	73.86	17.02	29.28	56.44

Table 4: Pose estimation on DL3DV-10K with [50°-90°] yaw range.

Method	Input	$R@5^\circ$	$R@15^\circ$	$R@30^\circ$	$T@5^\circ$	$T@15^\circ$	$T@30^\circ$	$MRE\downarrow$	$MTE\downarrow$	$AUC_{30} \uparrow$
DUST3R	Pair	34.33	63.00	94.66	27.00	75.00	92.67	13.36	10.88	55.58
InterPose [†] _{w/o SCS}	DynamiCrafter	36.33	64.33	95.00	26.00	76.33	92.67	13.32	11.27	55.68
InterPose [†]	DynamiCrafter	36.11	64.33	97.66	27.66	79.67	95.33	13.17	10.76	56.05
Ours _{w/o FMS}	Hybrid video	38.33	68.33	98.33	30.66	79.61	96.67	12.89	10.71	57.16
Ours	Hybrid video	38.10	70.00	100.00	31.33	81.33	98.33	12.73	10.28	57.48

Table 5: Pose estimation on NAVI for the [50°-90°] yaw range.

Method	Input	$R@5^\circ$	$R@15^\circ$	$R@30^\circ$	$T@5^\circ$	$T@15^\circ$	$T@30^\circ$	$MRE\downarrow$	$MTE\downarrow$	$AUC_{30} \uparrow$
DUST3R	Pair	64.69	95.72	98.05	62.37	97.28	98.22	7.30	7.82	82.37
InterPose [†] _{w/o SCS}	DynamiCrafter	45.13	92.61	96.11	57.86	91.44	96.50	11.14	8.81	78.63
InterPose [†]	DynamiCrafter	66.53	97.28	98.83	67.70	96.89	98.84	6.61	6.26	82.80
Ours _{w/o FMS}	Hybrid video	59.53	97.28	98.83	72.26	95.72	98.83	6.93	6.87	81.91
Ours	Hybrid video	70.82	97.67	98.83	75.10	98.44	99.22	5.97	5.46	83.98

Table 6: Runtime and Memory Cost.

Method	Runtime		Memory Cost	
	Video Generation	Pose Estimation	Video Generation	Pose Estimation
InterPose [†]	3.2min	20.29min	14.6GB	3.1GB
Ours	3.8min	0.18min	22.8GB	3.6GB

PoseCrafter produces significantly cleaner and more complete geometry. In addition, compared to InterPose[†], our method is able to correct blur and pose errors in the middle of the sequence, reducing drift and artifacts in the reconstructed point cloud, thus providing more reliable visual evidence for subsequent pose estimation.

4.3.3 Runtime and Memory Cost Discussion

Following the quantitative and qualitative evaluations, we next analyze the efficiency of our approach. Specifically, we compare runtime and memory cost with InterPose[†]. InterPose generates 8 video sequences for each image pair and samples 11 frames from each sequence to compute the self-consistency score, resulting in a computationally heavy pipeline. In contrast, our method adopts a streamlined design: we generate only a single hybrid video and apply a deterministic frame selection

Table 7: Ablation study on Hybrid Video Generation. SCS and FMS denote the frame selection strategies from InterPose [5] and ours, respectively.

Method	Input	MRE↓	R@5°	R@15°	R@30°	AUC ₃₀ ↑
DUST3R	Pair	18.14	40.34	71.25	82.99	61.98
InterPose [‡] _{w/o SCS}	DynamicCrafter	16.11	42.70	75.70	87.50	65.72
InterPose [‡]	DynamicCrafter	13.60	51.81	81.50	83.30	70.47
InterPose [‡] _{w/ FMS}	DynamicCrafter	13.02	52.08	85.76	90.63	73.93
ViewCrafter _{w/o FMS}	ViewCrafter	13.80	52.78	82.29	88.54	71.12
ViewCrafter _{w/ FMS}	ViewCrafter	12.45	53.82	84.03	90.28	72.82
Ours _{w/o FMS}	Hybrid video	13.24	54.51	89.24	92.71	76.13
Ours _{w/ SCS}	Hybrid video	12.11	54.86	88.54	91.32	76.11
Ours	Hybrid video	11.40	55.21	89.93	93.75	77.41

Table 8: Ablation study on feature matching methods used in FMS.

Feature Matching Method	MRE↓	R@5°	R@15°	R@30°	AUC ₃₀ ↑
ORB	11.40	55.21	89.93	93.75	77.41
RoMa	13.89	53.13	88.89	91.67	74.76
LoFTR	12.41	54.51	89.24	93.40	76.30
SIFT	13.89	53.13	88.89	92.67	74.77
SuperPoint	13.36	51.74	89.93	92.71	75.48

Table 9: Ablation study on different video interpolation models used in Hybrid Video Generation.

Method	Video Interpolation Model	MRE↓	R@5°	R@15°	R@30°	AUC ₃₀ ↑
DUST3R	--	18.14	40.34	71.25	82.99	61.98
Ours	DynamiCrafter	11.40	55.21	89.93	93.75	77.41
Ours	ToonCrafter	12.08	53.99	87.16	89.19	73.11

Table 10: Ablation study on the number of intermediate frames selected by FMS.

#Frames	MRE↓	R@5°	R@15°	R@30°	AUC ₃₀ ↑
4	11.36	55.90	89.93	93.40	77.59
6	11.40	55.21	89.93	93.75	77.41
8	11.93	55.90	89.93	92.10	76.23

strategy based on feature matching. This not only avoids redundant video synthesis but also eliminates repeated self-consistency evaluations.

As summarized in Table 6, our approach substantially reduces runtime in the pose estimation stage while maintaining comparable or even superior accuracy. Although the hybrid video generation incurs slightly higher memory usage than InterPose, the overall cost remains affordable for practical deployment. These results demonstrate that our framework strikes a favorable balance between efficiency and accuracy, underscoring its scalability to real-world applications.

4.4 Ablation Studies

In this section, we present ablation studies on the two key components of our framework: Hybrid Video Generation and the Feature Matching Selector. We further investigate the impact of different video generation models and the number of selected frames. All experiments are conducted on the Cambridge Landmarks benchmark under yaw changes of [50°-65°].

Hybrid Video Generation Table 7 evaluates pose estimation performance for three video-generation variants without any frame selection: DynamiCrafter (InterPose[‡]_{w/o SCS}), ViewCrafter (Viewcrafter_{w/o FMS}), and our hybrid video generation (Ours_{w/o FMS}). Ours_{w/o FMS} achieves the lower mean rotation error and higher recall than generated results from DynamiCrafter and ViewCrafter. This confirms that our hybrid video generation strategy, combining the motion-driven

interpolation of DynamiCrafter with the pose-conditioned novel-view synthesis of ViewCrafter, effectively addresses the limitations of each method, enhancing the fidelity of intermediate frames for improved pose estimation.

Frame Matching Selector In Table 8, we evaluate the performance of different keypoint selection algorithms within the Frame Matching Selector (FMS), including RoMa [34], LoFTR [16], ORB [35], SIFT [1], and SuperPoint [36]. We observe that ORB performs best in our case, which may be attributed to its robustness against color blurring and distortions in synthesized frames, making it more effective to select the most informative frames.

Impact of Video Interpolation Model To further validate the generality of our pipeline, we experimented with ToonCrafter [37], a recent interpolation model designed for cartoon-style data but also known for handling large-motion scenarios efficiently. As shown in Table 9, ToonCrafter performs slightly worse than DynamiCrafter, which is likely due to its training bias towards non-photorealistic content. Nevertheless, it still achieves clear improvements over directly using the input image pair. These results confirm that our hybrid pose estimation framework is compatible with a variety of video interpolation backbones, highlighting its flexibility and robustness.

Impact of Selected Frame Count Table 10 analyzes the impact of varying the number of frames k selected by the FMS module, where k is 4, 6, and 8 (excluding the input image pair). The results show that $k = 6$ yields the best overall performance. Using fewer frames may not provide sufficient information for accurate pose estimation, while using more frames can introduce errors due to artifacts in the generated frames.

5 Limitation

Although the effectiveness of our method has been validated across multiple diverse benchmarks, there are still certain cases where it may not perform well. For example, in our HVG stage, both the adopted video interpolation model (DynamiCrafter) and the novel view synthesis model (ViewCrafter) face challenges when the start and end frames exhibit significant illumination differences, resulting in synthesized frames with low confidence for subsequent pose estimation. This issue could potentially be mitigated by integrating a relighting model, such as IC-Light [38], to harmonize lighting conditions. In addition, our FMS module also encounters difficulties when the image texture is overly uniform, which hinders reliable feature matching. We will explore these limitations in future work.

6 Conclusion

In this work, we propose PoseCrafter, a simple yet effective framework for estimating camera poses from input image pairs with small or no overlap—commonly referred to as extreme viewpoint changes [5]. Estimating poses under such challenging conditions remains a difficult problem, as minimal shared visual content often leads to poor matching and unreliable geometry. Rather than relying solely on video interpolation to synthesize intermediate frames, we introduce Hybrid Video Generation (HVG), which produces clearer and more geometrically consistent frames. HVG couples the vanilla interpolation model DynamiCrafter with a pose-conditioned novel view synthesis model, ViewCrafter, allowing it to refine and correct structural inconsistencies and reduce mid-sequence blur. To further improve pose estimation accuracy, we also introduce Feature Matching Selector (FMS). This module identifies the most informative frames by evaluating their feature correspondences with the input image pair, ensuring that only geometrically meaningful frames suitable for pose estimation are selected. Through extensive experiments on multiple widely used benchmarks [5], including challenging scenes with extreme viewpoint differences, we demonstrate that PoseCrafter consistently outperforms existing state-of-the-art methods in both accuracy and robustness.

Acknowledgement. This research / project is supported by the National Research Foundation (NRF) Singapore, under its NRF-Investigatorship Programme (Award ID. NRF-NRFI09-0008), and the China Scholarship Council under Grant Number 202306290143. It is also supported by the National Engineering Laboratory for Integrated Aero-Space-Ground-Ocean Big Data Application Technology and by the National Natural Science Foundation of China (NSFC) under Grant Nos. U19B2037, 61901384, and 61971356.

References

- [1] David G Lowe. Distinctive image features from scale-invariant keypoints. *International journal of computer vision*, 60:91–110, 2004.
- [2] Herbert Bay, Tinne Tuytelaars, and Luc Van Gool. Surf: Speeded up robust features. In *Computer Vision—ECCV 2006: 9th European Conference on Computer Vision, Graz, Austria, May 7–13, 2006. Proceedings, Part I 9*, pages 404–417. Springer, 2006.
- [3] Shuzhe Wang, Vincent Leroy, Yohann Cabon, Boris Chidlovskii, and Jerome Revaud. Dust3r: Geometric 3d vision made easy. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 20697–20709, 2024.
- [4] Vincent Leroy, Yohann Cabon, and Jérôme Revaud. Grounding image matching in 3d with mast3r. In *European Conference on Computer Vision*, pages 71–91. Springer, 2024.
- [5] Ruojin Cai, Jason Y Zhang, Philipp Henzler, Zhengqi Li, Noah Snavely, and Ricardo Martin-Brualla. Can generative video models help pose estimation? *arXiv preprint arXiv:2412.16155*, 2024.
- [6] Jinbo Xing, Menghan Xia, Yong Zhang, Haoxin Chen, Wangbo Yu, Hanyuan Liu, Gongye Liu, Xintao Wang, Ying Shan, and Tien-Tsin Wong. Dynamicrafter: Animating open-domain images with video diffusion priors. In *European Conference on Computer Vision*, pages 399–417. Springer, 2024.
- [7] RunwayML. Tools for human imagination. <https://runwayml.com/product>, 2024. Accessed: November 2024.
- [8] LumaAI. Luma dream machine. <https://lumalabs.ai/dream-machine>, 2024. Accessed: September 2024.
- [9] Wangbo Yu, Jinbo Xing, Li Yuan, Wenbo Hu, Xiaoyu Li, Zhipeng Huang, Xiangjun Gao, Tien-Tsin Wong, Ying Shan, and Yonghong Tian. Viewcrafter: Taming video diffusion models for high-fidelity novel view synthesis. *arXiv preprint arXiv:2409.02048*, 2024.
- [10] Martin A Fischler and Robert C Bolles. Random sample consensus: a paradigm for model fitting with applications to image analysis and automated cartography. *Communications of the ACM*, 24(6):381–395, 1981.
- [11] Alex Kendall, Matthew Grimes, and Roberto Cipolla. Convolutional networks for real-time 6-dof camera relocalization. *CoRR*, 2015.
- [12] Angela Dai, Angel X Chang, Manolis Savva, Maciej Halber, Thomas Funkhouser, and Matthias Nießner. Scannet: Richly-annotated 3d reconstructions of indoor scenes. In *Proceedings of the IEEE conference on computer vision and pattern recognition*, pages 5828–5839, 2017.
- [13] Lu Ling, Yichen Sheng, Zhi Tu, Wentian Zhao, Cheng Xin, Kun Wan, Lantao Yu, Qianyu Guo, Zixun Yu, Yawen Lu, et al. DI3dv-10k: A large-scale scene dataset for deep learning-based 3d vision. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 22160–22169, 2024.
- [14] Varun Jampani, Kevis-Kokitsi Maninis, Andreas Engelhardt, Arjun Karpur, Karen Truong, Kyle Sargent, Stefan Popov, André Araujo, Ricardo Martin Brualla, Kaushal Patel, et al. Navi: Category-agnostic image collections with high-quality 3d shape and pose annotations. *Advances in Neural Information Processing Systems*, 36:76061–76084, 2023.
- [15] Johannes L Schonberger and Jan-Michael Frahm. Structure-from-motion revisited. In *Proceedings of the IEEE conference on computer vision and pattern recognition*, pages 4104–4113, 2016.
- [16] Jiaming Sun, Zehong Shen, Yuang Wang, Hujun Bao, and Xiaowei Zhou. Loftr: Detector-free local feature matching with transformers. In *Proceedings of the IEEE/CVF conference on computer vision and pattern recognition*, pages 8922–8931, 2021.
- [17] Qunjie Zhou, Torsten Sattler, and Laura Leal-Taixe. Patch2pix: Epipolar-guided pixel-level correspondences. In *Proceedings of the IEEE/CVF conference on computer vision and pattern recognition*, pages 4669–4678, 2021.
- [18] Jason Y Zhang, Deva Ramanan, and Shubham Tulsiani. Relpose: Predicting probabilistic relative rotation for single objects in the wild. In *European Conference on Computer Vision*, pages 592–611. Springer, 2022.

- [19] Samarth Sinha, Jason Y Zhang, Andrea Tagliasacchi, Igor Gilitschenski, and David B Lindell. Sparsepose: Sparse-view camera pose regression and refinement. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 21349–21359, 2023.
- [20] Zhenpei Yang, Jeffrey Z Pan, Linjie Luo, Xiaowei Zhou, Kristen Grauman, and Qixing Huang. Extreme relative pose estimation for rgb-d scans via scene completion, 2019.
- [21] Ruojin Cai, Bharath Hariharan, Noah Snavely, and Hadar Averbuch-Elor. Extreme rotation estimation using dense correlation volumes. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 14566–14575, 2021.
- [22] Hana Bezalel, Dotan Ankri, Ruojin Cai, and Hadar Averbuch-Elor. Extreme rotation estimation in the wild. *arXiv preprint arXiv:2411.07096*, 2024.
- [23] Jianwei Tang, Jieming Wang, and Jian-Fang Hu. Predicting human poses via recurrent attention network. *Visual Intelligence*, 1(1):18, 2023.
- [24] Chun-Hao Paul Huang, Jae Shin Yoon, Hyeonho Jeong, Niloy Mitra, and Duygu Ceylan. On unifying video generation and camera pose estimation. *arXiv preprint arXiv:2501.01409*, 2025.
- [25] Ian Goodfellow, Jean Pouget-Abadie, Mehdi Mirza, Bing Xu, David Warde-Farley, Sherjil Ozair, Aaron Courville, and Yoshua Bengio. Generative adversarial networks. *Communications of the ACM*, 63(11):139–144, 2020.
- [26] Masaki Saito, Eiichi Matsumoto, and Shunta Saito. Temporal generative adversarial nets with singular value clipping. In *Proceedings of the IEEE international conference on computer vision*, pages 2830–2839, 2017.
- [27] Sergey Tulyakov, Ming-Yu Liu, Xiaodong Yang, and Jan Kautz. Mocogan: Decomposing motion and content for video generation. In *Proceedings of the IEEE conference on computer vision and pattern recognition*, pages 1526–1535, 2018.
- [28] Ivan Skorokhodov, Sergey Tulyakov, and Mohamed Elhoseiny. Stylegan-v: A continuous video generator with the price, image quality and perks of stylegan2. In *Proceedings of the IEEE/CVF conference on computer vision and pattern recognition*, pages 3626–3636, 2022.
- [29] Nal Kalchbrenner, Aaron Oord, Karen Simonyan, Ivo Danihelka, Oriol Vinyals, Alex Graves, and Koray Kavukcuoglu. Video pixel networks. In *International Conference on Machine Learning*, pages 1771–1779. PMLR, 2017.
- [30] Ruslan Rakhimov, Denis Volkhonskiy, Alexey Artemov, Denis Zorin, and Evgeny Burnaev. Latent video transformer. *arXiv preprint arXiv:2006.10704*, 2020.
- [31] Jonathan Ho, Tim Salimans, Alexey Gritsenko, William Chan, Mohammad Norouzi, and David J Fleet. Video diffusion models. *Advances in Neural Information Processing Systems*, 35:8633–8646, 2022.
- [32] Jonathan Ho, Chitwan Saharia, William Chan, David J Fleet, Mohammad Norouzi, and Tim Salimans. Cascaded diffusion models for high fidelity image generation. *Journal of Machine Learning Research*, 23(47):1–33, 2022.
- [33] Uriel Singer, Adam Polyak, Thomas Hayes, Xi Yin, Jie An, Songyang Zhang, Qiyuan Hu, Harry Yang, Oron Ashual, Oran Gafni, et al. Make-a-video: Text-to-video generation without text-video data. *arXiv preprint arXiv:2209.14792*, 2022.
- [34] Johan Edstedt, Qiyu Sun, Georg Bökman, Mårten Wadenbäck, and Michael Felsberg. Roma: Robust dense feature matching. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 19790–19800, 2024.
- [35] Ethan Rublee, Vincent Rabaud, Kurt Konolige, and Gary Bradski. Orb: An efficient alternative to sift or surf. In *2011 International conference on computer vision*, pages 2564–2571. Ieee, 2011.
- [36] Daniel DeTone, Tomasz Malisiewicz, and Andrew Rabinovich. Superpoint: Self-supervised interest point detection and description. In *Proceedings of the IEEE conference on computer vision and pattern recognition workshops*, pages 224–236, 2018.
- [37] Jinbo Xing, Hanyuan Liu, Menghan Xia, Yong Zhang, Xintao Wang, Ying Shan, and Tien-Tsin Wong. Toonrafter: Generative cartoon interpolation. *ACM Transactions on Graphics (TOG)*, 43(6):1–11, 2024.

- [38] Lvmin Zhang, Anyi Rao, and Maneesh Agrawala. Scaling in-the-wild training for diffusion-based illumination harmonization and editing by imposing consistent light transport. In *The Thirteenth International Conference on Learning Representations*, 2025.
- [39] Jianyuan Wang, Minghao Chen, Nikita Karaev, Andrea Vedaldi, Christian Rupprecht, and David Novotny. Vggt: Visual geometry grounded transformer. In *Proceedings of the Computer Vision and Pattern Recognition Conference*, pages 5294–5306, 2025.

NeurIPS Paper Checklist

1. Claims

Question: Do the main claims made in the abstract and introduction accurately reflect the paper's contributions and scope?

Answer: [\[Yes\]](#)

Justification: The main claims made in the abstract and introduction accurately reflect the paper's contributions and scope.

Guidelines:

- The answer NA means that the abstract and introduction do not include the claims made in the paper.
- The abstract and/or introduction should clearly state the claims made, including the contributions made in the paper and important assumptions and limitations. A No or NA answer to this question will not be perceived well by the reviewers.
- The claims made should match theoretical and experimental results, and reflect how much the results can be expected to generalize to other settings.
- It is fine to include aspirational goals as motivation as long as it is clear that these goals are not attained by the paper.

2. Limitations

Question: Does the paper discuss the limitations of the work performed by the authors?

Answer: [\[Yes\]](#)

Justification: The paper discusses the limitations of the work performed by the authors in the experiments section.

Guidelines:

- The answer NA means that the paper has no limitation while the answer No means that the paper has limitations, but those are not discussed in the paper.
- The authors are encouraged to create a separate "Limitations" section in their paper.
- The paper should point out any strong assumptions and how robust the results are to violations of these assumptions (e.g., independence assumptions, noiseless settings, model well-specification, asymptotic approximations only holding locally). The authors should reflect on how these assumptions might be violated in practice and what the implications would be.
- The authors should reflect on the scope of the claims made, e.g., if the approach was only tested on a few datasets or with a few runs. In general, empirical results often depend on implicit assumptions, which should be articulated.
- The authors should reflect on the factors that influence the performance of the approach. For example, a facial recognition algorithm may perform poorly when image resolution is low or images are taken in low lighting. Or a speech-to-text system might not be used reliably to provide closed captions for online lectures because it fails to handle technical jargon.
- The authors should discuss the computational efficiency of the proposed algorithms and how they scale with dataset size.
- If applicable, the authors should discuss possible limitations of their approach to address problems of privacy and fairness.
- While the authors might fear that complete honesty about limitations might be used by reviewers as grounds for rejection, a worse outcome might be that reviewers discover limitations that aren't acknowledged in the paper. The authors should use their best judgment and recognize that individual actions in favor of transparency play an important role in developing norms that preserve the integrity of the community. Reviewers will be specifically instructed to not penalize honesty concerning limitations.

3. Theory Assumptions and Proofs

Question: For each theoretical result, does the paper provide the full set of assumptions and a complete (and correct) proof?

Answer: [NA]

Justification: The theory assumptions in this paper have already been proved in previous works (Diffusion Model).

Guidelines:

- The answer NA means that the paper does not include theoretical results.
- All the theorems, formulas, and proofs in the paper should be numbered and cross-referenced.
- All assumptions should be clearly stated or referenced in the statement of any theorems.
- The proofs can either appear in the main paper or the supplemental material, but if they appear in the supplemental material, the authors are encouraged to provide a short proof sketch to provide intuition.
- Inversely, any informal proof provided in the core of the paper should be complemented by formal proofs provided in appendix or supplemental material.
- Theorems and Lemmas that the proof relies upon should be properly referenced.

4. Experimental Result Reproducibility

Question: Does the paper fully disclose all the information needed to reproduce the main experimental results of the paper to the extent that it affects the main claims and/or conclusions of the paper (regardless of whether the code and data are provided or not)?

Answer: [Yes]

Justification: Details of training and evaluation are described in the Implementation Details section.

Guidelines:

- The answer NA means that the paper does not include experiments.
- If the paper includes experiments, a No answer to this question will not be perceived well by the reviewers: Making the paper reproducible is important, regardless of whether the code and data are provided or not.
- If the contribution is a dataset and/or model, the authors should describe the steps taken to make their results reproducible or verifiable.
- Depending on the contribution, reproducibility can be accomplished in various ways. For example, if the contribution is a novel architecture, describing the architecture fully might suffice, or if the contribution is a specific model and empirical evaluation, it may be necessary to either make it possible for others to replicate the model with the same dataset, or provide access to the model. In general, releasing code and data is often one good way to accomplish this, but reproducibility can also be provided via detailed instructions for how to replicate the results, access to a hosted model (e.g., in the case of a large language model), releasing of a model checkpoint, or other means that are appropriate to the research performed.
- While NeurIPS does not require releasing code, the conference does require all submissions to provide some reasonable avenue for reproducibility, which may depend on the nature of the contribution. For example
 - (a) If the contribution is primarily a new algorithm, the paper should make it clear how to reproduce that algorithm.
 - (b) If the contribution is primarily a new model architecture, the paper should describe the architecture clearly and fully.
 - (c) If the contribution is a new model (e.g., a large language model), then there should either be a way to access this model for reproducing the results or a way to reproduce the model (e.g., with an open-source dataset or instructions for how to construct the dataset).
 - (d) We recognize that reproducibility may be tricky in some cases, in which case authors are welcome to describe the particular way they provide for reproducibility. In the case of closed-source models, it may be that access to the model is limited in some way (e.g., to registered users), but it should be possible for other researchers to have some path to reproducing or verifying the results.

5. Open access to data and code

Question: Does the paper provide open access to the data and code, with sufficient instructions to faithfully reproduce the main experimental results, as described in supplemental material?

Answer: [No]

Justification: We will open my code if our work is accepted.

Guidelines:

- The answer NA means that paper does not include experiments requiring code.
- Please see the NeurIPS code and data submission guidelines (<https://nips.cc/public/guides/CodeSubmissionPolicy>) for more details.
- While we encourage the release of code and data, we understand that this might not be possible, so “No” is an acceptable answer. Papers cannot be rejected simply for not including code, unless this is central to the contribution (e.g., for a new open-source benchmark).
- The instructions should contain the exact command and environment needed to run to reproduce the results. See the NeurIPS code and data submission guidelines (<https://nips.cc/public/guides/CodeSubmissionPolicy>) for more details.
- The authors should provide instructions on data access and preparation, including how to access the raw data, preprocessed data, intermediate data, and generated data, etc.
- The authors should provide scripts to reproduce all experimental results for the new proposed method and baselines. If only a subset of experiments are reproducible, they should state which ones are omitted from the script and why.
- At submission time, to preserve anonymity, the authors should release anonymized versions (if applicable).
- Providing as much information as possible in supplemental material (appended to the paper) is recommended, but including URLs to data and code is permitted.

6. Experimental Setting/Details

Question: Does the paper specify all the training and test details (e.g., data splits, hyperparameters, how they were chosen, type of optimizer, etc.) necessary to understand the results?

Answer: [Yes]

Justification: The paper specify all the training and test details.

Guidelines:

- The answer NA means that the paper does not include experiments.
- The experimental setting should be presented in the core of the paper to a level of detail that is necessary to appreciate the results and make sense of them.
- The full details can be provided either with the code, in appendix, or as supplemental material.

7. Experiment Statistical Significance

Question: Does the paper report error bars suitably and correctly defined or other appropriate information about the statistical significance of the experiments?

Answer: [Yes]

Justification: The paper adopts the mean metric values for 3 times testing for a fair comparison.

Guidelines:

- The answer NA means that the paper does not include experiments.
- The authors should answer “Yes” if the results are accompanied by error bars, confidence intervals, or statistical significance tests, at least for the experiments that support the main claims of the paper.
- The factors of variability that the error bars are capturing should be clearly stated (for example, train/test split, initialization, random drawing of some parameter, or overall run with given experimental conditions).

- The method for calculating the error bars should be explained (closed form formula, call to a library function, bootstrap, etc.)
- The assumptions made should be given (e.g., Normally distributed errors).
- It should be clear whether the error bar is the standard deviation or the standard error of the mean.
- It is OK to report 1-sigma error bars, but one should state it. The authors should preferably report a 2-sigma error bar than state that they have a 96% CI, if the hypothesis of Normality of errors is not verified.
- For asymmetric distributions, the authors should be careful not to show in tables or figures symmetric error bars that would yield results that are out of range (e.g. negative error rates).
- If error bars are reported in tables or plots, The authors should explain in the text how they were calculated and reference the corresponding figures or tables in the text.

8. Experiments Compute Resources

Question: For each experiment, does the paper provide sufficient information on the computer resources (type of compute workers, memory, time of execution) needed to reproduce the experiments?

Answer: [Yes]

Justification: In the Experiments section Implementation Details section, the paper introduces the detailed implementation details, including the computing resources, training and evaluation details.

Guidelines:

- The answer NA means that the paper does not include experiments.
- The paper should indicate the type of compute workers CPU or GPU, internal cluster, or cloud provider, including relevant memory and storage.
- The paper should provide the amount of compute required for each of the individual experimental runs as well as estimate the total compute.
- The paper should disclose whether the full research project required more compute than the experiments reported in the paper (e.g., preliminary or failed experiments that didn't make it into the paper).

9. Code Of Ethics

Question: Does the research conducted in the paper conform, in every respect, with the NeurIPS Code of Ethics <https://neurips.cc/public/EthicsGuidelines>?

Answer: [Yes]

Justification: The paper is sure to preserve anonymity.

Guidelines:

- The answer NA means that the authors have not reviewed the NeurIPS Code of Ethics.
- If the authors answer No, they should explain the special circumstances that require a deviation from the Code of Ethics.
- The authors should make sure to preserve anonymity (e.g., if there is a special consideration due to laws or regulations in their jurisdiction).

10. Broader Impacts

Question: Does the paper discuss both potential positive societal impacts and negative societal impacts of the work performed?

Answer: [Yes]

Justification: This paper contributes to extreme pose estimation research and extends beyond just extreme pose estimation tasks. In the future, it has the potential to enhance other research fields and benefit human society.

Guidelines:

- The answer NA means that there is no societal impact of the work performed.

- If the authors answer NA or No, they should explain why their work has no societal impact or why the paper does not address societal impact.
- Examples of negative societal impacts include potential malicious or unintended uses (e.g., disinformation, generating fake profiles, surveillance), fairness considerations (e.g., deployment of technologies that could make decisions that unfairly impact specific groups), privacy considerations, and security considerations.
- The conference expects that many papers will be foundational research and not tied to particular applications, let alone deployments. However, if there is a direct path to any negative applications, the authors should point it out. For example, it is legitimate to point out that an improvement in the quality of generative models could be used to generate deepfakes for disinformation. On the other hand, it is not needed to point out that a generic algorithm for optimizing neural networks could enable people to train models that generate Deepfakes faster.
- The authors should consider possible harms that could arise when the technology is being used as intended and functioning correctly, harms that could arise when the technology is being used as intended but gives incorrect results, and harms following from (intentional or unintentional) misuse of the technology.
- If there are negative societal impacts, the authors could also discuss possible mitigation strategies (e.g., gated release of models, providing defenses in addition to attacks, mechanisms for monitoring misuse, mechanisms to monitor how a system learns from feedback over time, improving the efficiency and accessibility of ML).

11. Safeguards

Question: Does the paper describe safeguards that have been put in place for responsible release of data or models that have a high risk for misuse (e.g., pretrained language models, image generators, or scraped datasets)?

Answer: [NA]

Justification: The paper poses no such risks.

Guidelines:

- The answer NA means that the paper poses no such risks.
- Released models that have a high risk for misuse or dual-use should be released with necessary safeguards to allow for controlled use of the model, for example by requiring that users adhere to usage guidelines or restrictions to access the model or implementing safety filters.
- Datasets that have been scraped from the Internet could pose safety risks. The authors should describe how they avoided releasing unsafe images.
- We recognize that providing effective safeguards is challenging, and many papers do not require this, but we encourage authors to take this into account and make a best faith effort.

12. Licenses for existing assets

Question: Are the creators or original owners of assets (e.g., code, data, models), used in the paper, properly credited and are the license and terms of use explicitly mentioned and properly respected?

Answer: [Yes]

Justification: The paper cites the original paper that produced the code package or dataset, and follows the licenses for existing assets.

Guidelines:

- The answer NA means that the paper does not use existing assets.
- The authors should cite the original paper that produced the code package or dataset.
- The authors should state which version of the asset is used and, if possible, include a URL.
- The name of the license (e.g., CC-BY 4.0) should be included for each asset.
- For scraped data from a particular source (e.g., website), the copyright and terms of service of that source should be provided.

- If assets are released, the license, copyright information, and terms of use in the package should be provided. For popular datasets, paperswithcode.com/datasets has curated licenses for some datasets. Their licensing guide can help determine the license of a dataset.
- For existing datasets that are re-packaged, both the original license and the license of the derived asset (if it has changed) should be provided.
- If this information is not available online, the authors are encouraged to reach out to the asset's creators.

13. **New Assets**

Question: Are new assets introduced in the paper well documented and is the documentation provided alongside the assets?

Answer: [NA]

Justification: This paper does not introduce new assets.

Guidelines:

- The answer NA means that the paper does not release new assets.
- Researchers should communicate the details of the dataset/code/model as part of their submissions via structured templates. This includes details about training, license, limitations, etc.
- The paper should discuss whether and how consent was obtained from people whose asset is used.
- At submission time, remember to anonymize your assets (if applicable). You can either create an anonymized URL or include an anonymized zip file.

14. **Crowdsourcing and Research with Human Subjects**

Question: For crowdsourcing experiments and research with human subjects, does the paper include the full text of instructions given to participants and screenshots, if applicable, as well as details about compensation (if any)?

Answer: [NA]

Justification: The paper does not involve crowdsourcing nor research with human subjects.

Guidelines:

- The answer NA means that the paper does not involve crowdsourcing nor research with human subjects.
- Including this information in the supplemental material is fine, but if the main contribution of the paper involves human subjects, then as much detail as possible should be included in the main paper.
- According to the NeurIPS Code of Ethics, workers involved in data collection, curation, or other labor should be paid at least the minimum wage in the country of the data collector.

15. **Institutional Review Board (IRB) Approvals or Equivalent for Research with Human Subjects**

Question: Does the paper describe potential risks incurred by study participants, whether such risks were disclosed to the subjects, and whether Institutional Review Board (IRB) approvals (or an equivalent approval/review based on the requirements of your country or institution) were obtained?

Answer: [NA]

Justification: The paper does not involve crowdsourcing nor research with human subjects.

Guidelines:

- The answer NA means that the paper does not involve crowdsourcing nor research with human subjects.
- Depending on the country in which research is conducted, IRB approval (or equivalent) may be required for any human subjects research. If you obtained IRB approval, you should clearly state this in the paper.

- We recognize that the procedures for this may vary significantly between institutions and locations, and we expect authors to adhere to the NeurIPS Code of Ethics and the guidelines for their institution.
- For initial submissions, do not include any information that would break anonymity (if applicable), such as the institution conducting the review.

A Camera Pose Estimation Results based on MAST3R

To further validate the generality of PoseCrafter on different pose estimation backbones, we replaced DUST3R [3] with MAST3R [4], another transformer-based framework augmented with a dense local feature head and efficient sparse matching. We evaluated this configuration on four benchmarks with varying yaw ranges: Cambridge Landmarks (50°–65°), ScanNet (50°–65°), DL3DV-10K (50°–90°), and NAVI (50°–90°). As shown in Table 11 and Table 12, using MAST3R directly for pose estimation results in lower accuracy compared to DUST3R. A possible explanation is that the sparse matching algorithm employed by MAST3R may be less effective for image pairs with small or no overlap. Nonetheless, our proposed framework consistently improves estimation accuracy over the baseline pose estimation models, including MAST3R. This demonstrates that our approach generalizes well across different pose estimation backbones by synthesizing intermediate frames that are more suitable for pose estimation.

Table 11: Camera pose estimation results on outward-facing datasets (Cambridge Landmarks and ScanNet) based on MAST3R. We report rotation recall ($R@ \theta \uparrow$), translation recall ($T@ \theta \uparrow$), mean rotation error (MRE \downarrow), mean translation error (MTE \downarrow), and AUC₃₀ \uparrow .

Method	Input	Cambridge Landmarks					ScanNet								
		R@5°	R@15°	R@30°	MRE	AUC ₃₀	R@5°	R@15°	R@30°	T@5°	T@15°	T@30°	MRE	MTE	AUC ₃₀
MASt3R	Pair	9.03	43.75	57.99	51.61	36.15	25.00	55.17	62.93	6.90	30.17	48.28	36.34	50.13	26.84
Ours(MASt3R)	Hybrid video	25.00	59.03	69.10	41.02	49.83	31.90	65.52	77.59	18.97	45.69	63.79	29.97	38.50	40.66

Table 12: Camera pose estimation results on center-facing datasets (DL3DV-10K and NAVI) based on MAST3R. We report rotation recall ($R@ \theta \uparrow$), translation recall ($T@ \theta \uparrow$), mean rotation error (MRE \downarrow), mean translation error (MTE \downarrow), and AUC₃₀ \uparrow .

Dataset	Method	Input	R@5°	R@15°	R@30°	T@5°	T@15°	T@30°	MRE	MTE	AUC ₃₀
DL3DV-10K	MASt3R	Pair	7.00	63.67	97.00	26.67	72.67	91.67	15.18	13.31	53.49
	Ours(MASt3R)	Hybrid video	7.67	68.67	97.33	30.00	76.33	92.00	14.09	12.91	55.36
NAVI	MASt3R	Pair	43.97	93.39	96.89	50.97	89.49	97.67	8.30	7.43	76.54
	Ours(MASt3R)	Hybrid video	43.97	93.77	97.28	51.75	91.83	97.28	7.71	6.98	77.26

Table 13: Additional comparison with VGGT on Cambridge Landmarks. We report mean rotation error (MRE \downarrow), rotation recall ($R@ \theta \uparrow$), and AUC₃₀ \uparrow .

Method	Input	MRE	R@5°	R@15°	R@30°	AUC ₃₀
Dust3R	Pair	18.14	40.34	71.25	82.99	61.98
Ours	Hybrid video	11.40	55.21	89.93	93.75	77.41
VGGT	Pair	20.17	40.00	70.17	82.29	60.54
VGGT _{Ours}	Hybrid video	17.88	42.43	84.40	85.76	65.15

B Additional Comparison with VGGT

To further analyze the generality of our pipeline, we conducted comparative experiments with VGGT [39] on the Cambridge Landmarks dataset under yaw changes of [50°–65°].

We first evaluated the two models in challenging cases where only image pairs with small or no overlap were provided. As shown in the second and fourth rows of Table 13, DUST3R consistently outperforms VGGT on such data.

We then integrate these two models into our pipeline and evaluate their performance. As shown in the third and fifth rows, both configurations achieve significant improvements over their directly estimated counterparts. The version using DUST3R achieves a higher accuracy than the version using VGGT. These results demonstrate that our method is compatible with different pose estimators and consistently enhances their performance.

Overall, DUS3R appears better suited than VGGT for small- or non-overlapping image pairs, and our framework provides more noticeable improvements when combined with DUS3R. While fine-tuning VGGT on small-overlap data or adjusting the hyperparameters of our pipeline may further enhance its performance, we leave such extensions for future work.

C Comparison of DUS3R Confidence-based Selection and FMS

We compare DUS3R confidence-based frame selection with our proposed Feature Matching Selector (FMS) to evaluate whether confidence scores can serve as a viable alternative. In this setting, top-ranked frames were selected from hybrid videos using four different confidence thresholds (20%, 40%, 60%, and 80%). The results are summarized in Table 14.

When the threshold was set to 20% or 40%, pose estimation accuracy decreased compared to using the full hybrid video sequence. Increasing the threshold to 60% or 80% improved performance, but the accuracy still remained lower than that achieved by our proposed FMS.

In terms of efficiency, confidence-based selection introduces substantial overhead, as all video frames must be processed by DUS3R to compute confidence maps prior to selection. This step incurs significant time and memory costs, and higher thresholds further increase runtime as more frames are selected. By contrast, our FMS requires only a single feature extraction step, immediately identifying the most informative frames with superior efficiency and accuracy. These findings highlight the practical advantages of our proposed FMS over confidence-based selection in both computational efficiency and pose estimation performance.

Table 14: Comparison of DUS3R confidence-based selection and FMS on Cambridge Landmarks. We report pose estimation results under yaw changes of $[50^\circ-65^\circ]$. Frames were selected from hybrid videos using four DUS3R confidence thresholds (20%, 40%, 60%, and 80%). We report mean rotation error (MRE \downarrow), rotation recall ($R@ \theta \uparrow$), $AUC_{30} \uparrow$ and pose estimation time.

Method	MRE \downarrow	$R@5^\circ$	$R@15^\circ$	$R@30^\circ$	$AUC_{30} \uparrow$	Pose Estimation Time
Conf(20%)	14.66	54.17	85.07	89.24	72.45	2.79min
Conf(40%)	14.36	57.30	87.15	90.97	74.24	2.91min
Conf(60%)	12.35	54.17	90.28	93.06	76.60	3.47min
Conf(80%)	12.41	53.47	90.63	93.06	76.46	4.13min
Ours_{w/o} FMS	13.24	54.51	89.24	92.71	76.13	2.56min
Ours	11.40	55.21	89.93	93.75	77.41	0.18min

D Visualization Results of Limitation

Although PoseCrafter achieves robust results across various benchmarks, certain challenging scenes still degrade its performance. In the hybrid video generation (HVG) stage, severe illumination differences between the start and end frames will introduce obvious artifacts in the synthesized intermediate views. The **red-boxed** regions in Figure 4 highlight these artifacts, which would be marked as low confidence and discarded in the downstream pose estimation backbone, e.g., DUS3R in this work. Such an obvious removal of information may significantly degrade the final performance. Moreover, in the feature match selector (FMS) module, scenes dominated by uniform or repetitive textures hinder the reliable extraction and matching of key points, resulting in fewer RANSAC [10] inliers and lower pose accuracy. The **red** lines in Figure 5 show examples of incorrect correspondences in such scenes.

E More Visualizations of Our Generated Videos

To further illustrate the superiority of our proposed hybrid video generation(HVG), we present additional side-by-side comparisons of intermediate frames produced by DynamiCrafter, ViewCrafter, and PoseCrafter. In Figure 6, DynamiCrafter [6] delivers temporally smooth transitions but exhibits progressive blur and geometric drift in the middle of generated sequences. Although ViewCrafter [9]



Figure 4: Hybrid Video Generation artifacts under severe illumination differences. **Blue** and **yellow** outlines indicate the start and end frames, respectively. Confidence images denote images filtered with the predicted confidence map in the subsequent DUS3R model. The **red** boxes highlight regions affected by artifacts in these cases. We can observe that these regions have quite low confidence.




Figure 5: Feature Matching Selection failures in low-texture regions. **Blue** and **yellow** outlines indicate the start and end frames, respectively. The **red** and **green** lines indicate incorrect and correct correspondences, respectively. Incorrect correspondences lead to errors in inlier counting, affecting the accuracy of subsequent frame selection.

can produce sharp results with minimal blur, using only input image pairs with small overlap often leads to structural artifacts and misalignments. To address these limitations, we combine DynamiCrafter and ViewCrafter to complement each other’s strengths. Specifically, we first use DynamiCrafter to synthesize intermediate “relay” frames, effectively augmenting the input image pair with frames that have larger overlaps. These relay frames are then passed to ViewCrafter to generate clearer and more geometrically consistent results. As shown in Figure 6, our proposed approach successfully produces frames that are both visually sharp and structurally reliable.




Figure 6: Comparative video synthesis results. Each row shows intermediate frames synthesized between the same start frame (blue box) and end frame (yellow box), generated by different methods: DynamiCrafter (DC), ViewCrafter (VC), and our Hybrid Video Generation (Ours). DC produces smooth motion but exhibits progressive blur and geometric drift in the middle of sequences (highlighted in red circle). Since VC is sensitive to the pose of the input image pair, it tends to produce structural misalignments in our small-overlap setting (highlighted in red box). By coupling DC and VC together, our method delivers sharp, geometrically consistent video frames throughout, correcting both the blur of DC and the misalignments of VC.