

Redirect4D-Bench: A Scalable Benchmark for Camera Redirection of Monocular Dynamic Videos with Pseudo-4D Ground Truth

Wei Cao¹, Hao Zhang¹, Jiapeng Tang², Yulun Wu¹, Yingying Li¹, Ning Yu³, Yaoyao Liu¹

¹University of Illinois Urbana-Champaign, ²Technical University of Munich, ³Netflix

Camera redirection aims to replay a monocular source video along a requested camera path. Recent methods can generate plausible redirected videos, but current evaluations miss clear failures: an output may ignore the requested path, distort the foreground subject, or place it incorrectly in the new view. Reliable evaluation needs target views along the requested path, yet existing sources are limited: synthetic data has a domain gap, multi-camera capture is scarce and covers only recorded views, and monocular web videos lack observations from arbitrary camera paths. As a result, evaluations on real videos often rely on realism or semantic metrics, which do not directly measure camera following or subject placement. To tackle this, we introduce *Redirect4D-Bench*, a benchmark for camera redirection of real monocular dynamic videos, whose training-free construction pipeline scales the benchmark to more source clips, producing pseudo-4D ground truth for each requested camera path. The pipeline reconstructs a 4D scene for each curated clip, so multiple camera paths can be defined over the same dynamic event. It then renders depth and subject masks along each path as the pseudo-4D ground truth. Based on this pseudo-4D ground truth, we adapt camera-pose errors for camera accuracy and introduce object-fidelity metrics. Benchmark results show that CLIP, FID/FVD, and VBench can miss visible trajectory and subject failures, while Redirect4D-Bench metrics reveal complementary errors in camera following and subject preservation.

Project page: <https://redirect4d-bench.cvmilgroup.web.illinois.edu/>



1 Introduction

Camera redirection asks a model to replay a monocular source video along a requested camera path. As video generation moves beyond fixed viewpoints, camera redirection has become a natural interface for controlling how dynamic scenes are viewed (Gao et al., 2021; Van Hoorick et al., 2024; Wang et al., 2025e), with applications ranging from free-viewpoint replay to augmented reality (Azuma, 1997). Success requires more than visual realism: the output must follow the requested camera motion while preserving the dynamic foreground subject in the target view. Recent systems make this task increasingly plausible through camera-conditioned generation (He et al., 2024), video re-rendering (Bai et al., 2025a; Zhang et al., 2025a), and geometry-guided synthesis (Yu et al., 2025; Ren et al., 2025; Hu et al., 2025a).

However, existing evaluations often rely on generic video quality or semantic similarity, which may fail to capture whether the output follows the requested camera path and preserves the foreground subject. Consequently, a redirected video that scores well on these evaluations may still fail the actual task: it can appear realistic while ignoring the target trajectory, damaging the foreground subject, or placing the subject in the wrong target-view location.

Reliable evaluation of camera redirection on real videos is difficult because references for the requested path are usually unavailable. Without such references, it is hard to judge whether the generated video follows the requested camera path and preserves the foreground subject in the new view, as illustrated in Fig. 1. Existing metrics capture only part of this evaluation goal. Full-reference fidelity and perceptual metrics such as PSNR (Huynh-Thu and Ghanbari, 2008), SSIM (Wang et al., 2004), and LPIPS (Zhang et al., 2018) require physically captured target-view RGB frames along the same trajectory. Distributional, semantic, and quality metrics such as FID (Heusel et al., 2017), FVD (Unterthiner et al., 2019), CLIP (Radford et al., 2021), and

Input Video & 4D Point Cloud					CLIP			VBench							Visual check
					T↑	F↑	V↑	SC↑	BG↑	TF↑	MS↑	AQ↑	IQ↑	OC↑	
															
ReCamMaster					0.314	0.967	0.920	0.898	0.938	0.944	0.970	0.569	0.715	0.313	 Does not follow trajectory
TrajectoryCrafter					0.284	0.948	0.802	0.733	0.903	0.933	0.964	0.440	0.741	0.265	 Distorted Geometry
GEN3C					0.288	0.951	0.864	0.708	0.919	0.948	0.982	0.546	0.660	0.293	 Distorted Geometry
FreeOrbit4D					0.304	0.954	0.856	0.825	0.919	0.932	0.977	0.531	0.565	0.292	 Looks Good!

Figure 1 High video-metric scores do not imply successful camera redirection. CLIP-T/F/V (Radford et al., 2021) denote text alignment, adjacent-frame consistency, and source-video consistency; VBench (Huang et al., 2024) SC/BG/TF/MS/AQ/IQ/OC denote subject consistency, background consistency, temporal flickering, motion smoothness, aesthetic quality, imaging quality, and overall consistency. The shaded cells mark the best metric scores, but the corresponding outputs either do not follow the requested camera path or deform the subject. The visual check instead identifies the output that follows the target trajectory and preserves the foreground subject.

VBench (Huang et al., 2024) measure realism, similarity, or temporal quality, but they are not tied to the requested camera path. Camera-pose errors are closer to the task (He et al., 2024; Bai et al., 2025a), but they do not verify whether the foreground subject remains recognizable or correctly placed. This missing-reference problem is fundamental for real dynamic videos. No existing data source jointly provides in-the-wild realism and per-path target views: studio multi-camera capture covers only staged subjects, synthetic and hybrid corpora differ from real videos (Bai et al., 2025a; Yu et al., 2025; Ren et al., 2025), and monocular web videos observe only the source camera path. Thus, a useful in-the-wild benchmark must construct a per-case reference for each requested trajectory.

To address this, we introduce *Redirect4D-Bench*, which constructs the missing reference for each case: a training-free pipeline reconstructs a 4D scene from each curated clip, so that multiple camera paths can be defined over the same dynamic event and pseudo-4D ground truth can be rendered along each of them. For each redirection case, *Redirect4D-Bench* combines a curated source clip, the reconstructed 4D scene, a parameterized camera path, and rendered depth and subject masks as the pseudo-4D ground truth. Together, these assets specify the camera path and subject placement that each generated video should satisfy. The retained benchmark covers human, robot, animal, and other dynamic subjects, with most cases requiring large-angle camera motion.

Based on this pseudo-4D ground truth, we adapt camera-pose errors and introduce object-fidelity metrics. *Camera accuracy* measures whether the camera trajectory estimated from the generated video follows the benchmark target trajectory. *Object fidelity* measures whether the foreground subject is detected, recognizable, and spatially aligned with the target-view subject-mask pseudo ground truth. Across four representative camera-redirection systems, these metrics reveal that camera following and subject preservation are partially decoupled, while CLIP, FID/FVD, and VBench scores can rank visually apparent task failures inconsistently. Our contributions are:

- We introduce *Redirect4D-Bench*, a scalable benchmark for camera redirection of real monocular dynamic videos. Its training-free pseudo-4D ground-truth pipeline reconstructs a 4D scene per curated clip and defines multiple camera paths over one dynamic event.
- We adapt camera-pose errors into a camera-accuracy metric and introduce object-fidelity metrics, measuring whether a generated video follows the requested camera path and whether the foreground

subject remains recognizable and correctly localized in the target view.

- We benchmark representative camera-redirection methods and show that common video metrics can miss visible task failures, while Redirect4D-Bench metrics expose complementary errors in camera following and subject preservation.

2 Related Work

Camera redirection and camera-aware generation. Camera redirection aims to resynthesize a source video along a user-specified camera trajectory, enabling post-capture viewpoint control and free-viewpoint replay of dynamic scenes (Gao et al., 2021; Yu et al., 2025; Zhang et al., 2025a). Recent camera-aware generation methods use explicit camera conditioning and camera-control analysis (Wang et al., 2024; Bahmani et al., 2024, 2025), ray or relative camera encodings (He et al., 2024; Li et al., 2025a), metric-scale and reference-based camera control (Li et al., 2025b; Luo et al., 2025), multi-video and time-camera control (Kuang et al., 2024; Wang et al., 2025e), and training-free control (Hou and Chen, 2024). They operate on images or videos through generative novel-view synthesis (Van Hoorick et al., 2024; Yu et al., 2024; You et al., 2024), video re-rendering (Bai et al., 2025a; Zhang et al., 2025a; Huang et al., 2025b), spatiotemporal caches (Ren et al., 2025), geometry or 3D-tracking guidance (Yu et al., 2025; Hu et al., 2025a; Gu et al., 2025), 4D geometric control (Zheng et al., 2026; Cao et al., 2026b), and 4D point-cloud proxies (Lin et al., 2026; Cao et al., 2026a). Unlike CameraBench (Lin et al., 2025), which evaluates camera-motion recognition in existing videos, Redirect4D-Bench targets generated redirected videos and tests whether they follow the requested camera trajectory while preserving the foreground subject under that trajectory.

Video data and benchmark references. Web videos have long supported large-scale video benchmarks, from action and video understanding (Heilbron et al., 2015; Kay et al., 2017; Abu-El-Haija et al., 2016) to video-language learning and video generation data (Xue et al., 2022; Miech et al., 2019; Bain et al., 2021; Chen et al., 2024a; Nan et al., 2024; Ju et al., 2024). However, these datasets provide semantic labels, captions, or unconstrained videos rather than trajectory-specific geometric references, and object-level 3D annotations on images (Ma et al., 2024a,b; Ji et al., 2026) cover static objects rather than dynamic subjects along a camera path. Dynamic camera-pose and spatially annotated web-video datasets such as DynPose100K (Rockwell et al., 2025) and SpatialVID (Wang et al., 2025b) scale pose, depth, and motion annotations for spatial learning, but they do not assign target trajectories to source clips or provide target-view subject references for redirection. Camera-control and view-synthesis systems therefore rely on posed static-scene videos (Zhou et al., 2018; Ling et al., 2024), driving data (Sun et al., 2020; Cao et al., 2025; Zhou et al., 2023, 2025a), synthetic multi-view data (Greff et al., 2022), hybrid web/static-view corpora (Van Hoorick et al., 2024; Bai et al., 2025a; Yu et al., 2025; Ren et al., 2025), or small dynamic benchmarks limited to recorded views (Gao et al., 2022). Renderable 3D/4D representations can be built with NeRFs (Mildenhall et al., 2021; Tancik et al., 2023), Gaussian or feed-forward splatting (Kerbl et al., 2023; Charatan et al., 2024; Chen et al., 2024b; Jiang et al., 2025a; Wu et al., 2026), dynamic radiance fields (Gao et al., 2021), motion-aware Gaussians (Wang et al., 2025d; Lei et al., 2025), point maps (Zhang et al., 2024), depth/pose estimation (Teed and Deng, 2021; Yang et al., 2024a; Li et al., 2025c; Jiang et al., 2025b), multi-view 4D generation (Yao et al., 2025; Zhang et al., 2025b), and non-rigid surface reconstruction (Cao et al., 2024; Tang et al., 2026). Redirect4D-Bench curates monocular web clips and reconstructs a 4D scene for each clip to derive target-trajectory depth and subject-mask pseudo ground truth.

Evaluation metrics. Existing evaluations measure useful but incomplete aspects of camera redirection. Full-reference fidelity and perceptual metrics such as PSNR (Huynh-Thu and Ghanbari, 2008), SSIM (Wang et al., 2004), and LPIPS (Zhang et al., 2018) require captured target-view RGB frames along the same trajectory. Distributional metrics such as FID (Heusel et al., 2017) and FVD (Unterthiner et al., 2019), semantic metrics such as CLIP (Radford et al., 2021), and video-generation benchmarks such as EvalCrafter (Liu et al., 2024a) and VBench (Huang et al., 2024) measure realism, similarity, and temporal quality, but they are not tied to the requested camera path or target-view subject placement. Motion-oriented benchmarks (Ling et al., 2025) and camera-pose errors used in camera-control evaluation (He et al., 2024; Bai et al., 2025a; Wang et al., 2025f) are closer to the task, but they do not jointly evaluate camera following, subject recognition,




Figure 2 Redirect4D-Bench construction pipeline. Curated source clips and foreground masks feed two parallel branches: foreground completion from multi-view observations, and metric background reconstruction with ViPE (Huang et al., 2025a). We align the completed foreground to the metric scene through 3D–3D correspondences at foreground-mask pixels, yielding a shared 4D scene. This scene is rendered along target trajectories to produce the pseudo ground-truth target mask for evaluation.

and target-view localization for each source clip. Redirect4D-Bench therefore evaluates camera accuracy for trajectory following together with object fidelity for detection, recognition, and target-view localization.

3 Pseudo-4D Ground Truth Construction

For each source clip and target camera path, evaluating camera redirection requires a reference that specifies how the dynamic scene should be observed. No existing data source provides such references for in-the-wild dynamic scenes. Redirect4D-Bench addresses this by constructing *pseudo-4D ground truth*: for each curated source clip, we reconstruct a temporally aligned, foreground-complete 4D point-cloud representation that fills in subject geometry occluded in the source view, express target camera trajectories in its coordinate space, and render pseudo ground-truth target masks along each trajectory. Figure 2 summarizes the pipeline; the next three subsections describe source-video curation, 4D reconstruction with trajectory rendering, and the resulting dataset contents.

3.1 Source Video Collection and Curation

Redirect4D-Bench starts from in-the-wild monocular clips with moving foreground subjects that admit reliable tracking and 4D reconstruction. Adapting the web-video pipeline of Animal-in-Motion (Zhao et al., 2025), we download candidate YouTube videos from category-level queries, split them into shot-level clips, and filter out clips that are off-topic or multi-shot or that lack useful foreground motion. For each retained clip, we obtain object tracks and foreground masks with a Grounded-SAM pipeline (Ren et al., 2024) that combines Grounding DINO (Liu et al., 2023a) with SAM-based video segmentation (Ravi et al., 2024), discard tracks with severe overlap, truncation, very small subjects, unstable identity association, or infeasible crops, and standardize the remainder into object-centric sequences. After a final manual review, the retained tracks are passed to the 4D reconstruction stage described in Sec. 3.2.

3.2 4D Reconstruction and Pseudo Ground Truth Rendering

For every source clip, we build a renderable 4D point-cloud representation that combines a completed dynamic foreground with a metric static background. Following Cao et al. (2026a), we reconstruct foreground and




Figure 3 Per-case data bundle. Each retained redirection case contains a 45-frame source clip with foreground masks, a foreground-complete 4D point cloud, a parameterized target trajectory, a text prompt, a rendered target-view depth map, and the pseudo ground-truth target mask used as the per-case object-fidelity target for detection, recognition, and localization.

background in two coordinate spaces: the foreground in canonical object space for completion, and the background in global scene space, where target trajectories are defined.

Foreground completion. We use the curated masks (Sec. 3.1) to isolate the subject and generate multi-view observations with SV4D 2.0 (Yao et al., 2025). Source and synthesized views are fed to VGGT-1B (Wang et al., 2025c), whose point-map head predicts dense per-pixel 3D coordinates; masking to subject pixels gives a per-frame foreground point cloud expressed in canonical object space.

Metric-scale background reconstruction. For background reconstruction, we use ViPE (Huang et al., 2025a), which we found to produce a more stable metric scene on spatially extended web-video clips than PAGE-4D (Zhou et al., 2025b). ViPE recovers source-camera poses and intrinsics; combined with metric depth, these unproject background pixels into the global scene space in which we both specify target trajectories and evaluate the recovered camera motion (Sec. 4).

Foreground-background alignment. The completed foreground and the metric background initially live in different coordinate spaces. Using foreground-mask pixels in the source view, we establish 3D–3D correspondences between the completed foreground point map and the metric source-view point map (Cao et al., 2026a). From these correspondences, we estimate a per-frame scale-and-translation transform, apply it to the full completed foreground, and smooth the result over time. The resulting 4D scene combines the metric background and per-frame completed foreground point clouds in the shared global scene space, ready for trajectory rendering.

Pseudo ground truth rendering. Each target trajectory is a parameterized camera arc whose center is obtained by unprojecting the first-frame center pixel at the reconstructed depth; the same depth gives the base orbit radius, with the source camera at the origin. The arc is specified by yaw, pitch, roll, and a radius scale; together with the source clip, the resulting 45 per-frame poses define a camera-redirection case. We rasterize the 4D point cloud along the target poses to obtain depth and foreground-mask scaffolds (Cao et al., 2026a), which condition Wan2.2 VACE (Wang et al., 2025a) together with a reference image and text prompt to synthesize the target-view RGB. The mask scaffold is then refined via SAM (Carion et al., 2025) video propagation on the synthesized frames, yielding the pseudo ground-truth target mask.




Figure 4 Redirect4D-Bench metric protocol. (Left) The evaluated method receives the source video and target camera trajectory and produces an output video, from which SAM (Carion et al., 2025) extracts a per-frame subject mask. (#1, top right) **Camera accuracy:** ViPE (Huang et al., 2025a) recovers an estimated camera trajectory from the output video and visualizes it alongside the target in the reconstructed scene; we report rotation and translation errors. (#2, bottom right) **Object fidelity:** for each output frame, Detection checks whether the subject is present and Qwen3-VL (Bai et al., 2025b) Recognition judges whether the rendered subject is structurally plausible (e.g., catching defects like extra limbs); on recognized frames, MaskIoU is computed against the per-case pseudo ground-truth target mask.

3.3 Dataset Contents and Statistics

Figure 3 shows the per-case bundle. Redirect4D-Bench comprises 83 cases from 62 source clips at 15 fps and 832×480 resolution, spanning robot, human, and animal subjects (plus one other clip). Most cases (78/83) require camera arcs spanning at least 90° , making the benchmark a stress test of trajectory following rather than small viewpoint perturbation. Because the construction pipeline is training-free, additional cases can be added at roughly 0.5 A100-hours per case (~ 40 A100-hours for the 83 retained cases; Appendix D), with manual effort limited to a final review of curated tracks. Target-view RGB is not provided as ground truth: camera redirection reveals previously occluded background regions that admit many plausible completions. Evaluation instead relies on two well-defined targets: the requested camera trajectory (for camera accuracy) and the pseudo ground-truth target mask (for object fidelity). Detailed dataset statistics are given in Appendix A. Safety filtering and the public release policy are described in Appendix B.

4 Benchmark Metrics

Standard video metrics do not verify whether a generated video follows the requested camera path or preserves the foreground subject in the target view. Redirect4D-Bench therefore introduces two task-specific metric families: camera accuracy and object fidelity.

4.1 Task Definition

A redirection case consists of a source clip $\mathcal{V}^{\text{src}} = \{\mathbf{I}_t^{\text{src}}\}_{t=1}^T$ and a target camera trajectory $\{\boldsymbol{\pi}_t^{\text{tgt}}\}_{t=1}^T$. A method outputs a redirected video $\hat{\mathcal{V}}^{\text{tgt}}$ of length T , which Redirect4D-Bench evaluates along two axes (Fig. 4): camera accuracy (Sec. 4.3), asking whether the recovered camera path follows the target trajectory; and object fidelity (Sec. 4.4), asking whether the foreground subject remains recognizable and correctly localized in the target view.

4.2 Limitations of Existing Metrics

Full-reference fidelity metrics. PSNR (Huynh-Thu and Ghanbari, 2008), SSIM (Wang et al., 2004), and LPIPS (Zhang et al., 2018) require a captured target-view frame along the same camera path, which does

not exist for in-the-wild monocular videos (Sec. 1); pixel-level fidelity cannot rank methods on real dynamic scenes.

Distributional and quality metrics. FID (Heusel et al., 2017), FVD (Unterthiner et al., 2019), CLIP similarities (Radford et al., 2021), and VBench dimensions (Huang et al., 2024) measure realism, semantic similarity, and temporal quality at the dataset level, but none is conditioned on the requested camera path: a video can score well while ignoring the path or misplacing the subject (Fig. 1), so they diagnose general video quality rather than whether the requested redirection happened.

Camera-pose metrics. Rotation and translation errors assess whether a generated video follows the target camera trajectory, and we adopt this style of measurement in Sec. 4.3. On their own, they remain incomplete: a method can match the target trajectory while the foreground subject disappears, distorts, or appears in the wrong target-view location. Pose accuracy alone is not a faithful proxy for successful redirection, so Redirect4D-Bench complements it with object fidelity (Sec. 4.4).

4.3 Camera Accuracy

Camera accuracy asks whether the generated video $\hat{\mathcal{V}}^{\text{tgt}}$ follows the target trajectory $\{\pi_t^{\text{tgt}}\}_{t=1}^T$ (Fig. 4, camera-pose branch). Baselines do not necessarily expose their internal camera poses, so we evaluate every method through a uniform pose-recovery protocol that estimates a recovered trajectory $\{\hat{\pi}_t\}_{t=1}^T$ from the generated video itself; this makes the metric independent of any baseline-specific camera representation or output format.

Pose recovery and alignment. We run ViPE (Huang et al., 2025a) with its `lyra` configuration on each generated video to recover per-frame poses $\hat{\pi}_t = (\hat{R}_t, \hat{\mathbf{t}}_t)$. To remove the arbitrary global offset inherent to monocular pose recovery, both target and recovered trajectories are expressed relative to their first frame, so the metric reports per-frame deviation from the requested arc rather than absolute world position. Because monocular pose recovery returns translations only up to a per-clip scale, TransErr is computed in the same metric units as the target trajectory after aligning the two first frames, and it should be read alongside RotErr, which is scale-invariant. Figure 4 (top right) shows the ViPE reconstruction of the generated video with the estimated and target trajectories in separate 3D frames.

Rotation and translation error. Following camera-control works (He et al., 2024; Bai et al., 2025a), per-frame errors against the aligned target $\pi_t^{\text{tgt}} = (R_t^{\text{tgt}}, \mathbf{t}_t^{\text{tgt}})$ are

$$\text{RotErr}_t = \arccos\left(\frac{\text{tr}(R_t^{\text{tgt}} \hat{R}_t^{\top}) - 1}{2}\right), \quad \text{TransErr}_t = \|\mathbf{t}_t^{\text{tgt}} - \hat{\mathbf{t}}_t\|_2.$$

Benchmark-level RotErr and TransErr are computed by averaging per-frame errors within each case, then over the 83 retained cases.

4.4 Object Fidelity

Following the target camera path does not guarantee successful redirection: the generated video must also render a recognizable foreground subject and place it at the correct target-view location. Object fidelity (Fig. 4, object-fidelity branch) is therefore a per-frame cascade over target-view subject masks: detection and recognition measure whether the subject is rendered correctly, and MaskIoU measures whether it is correctly localized.

Reference and predicted masks. For each generated video, we apply SAM (Carion et al., 2025) to obtain a predicted subject mask at every frame and compare against the per-case pseudo ground-truth target mask from Sec. 3. Object-fidelity metrics use the clip length T defined in Sec. 4.1 ($T = 45$) as their denominator; by design, every retained case has a non-empty pseudo ground-truth target mask in all T frames (Sec. 3.2). Frames where SAM returns an empty or invalid mask thus count as failures.

Detection and recognition. Detection asks whether the subject appears in the generated frame at all: $\text{detected}(t)$ holds iff SAM returns a non-empty subject mask at frame t (the reference target mask is non-empty in every frame by construction). Detection does not check where the subject appears; localization is measured separately by MaskIoU. Recognition refines this by asking Qwen3-VL-32B (Bai et al., 2025b) whether the predicted subject is structurally plausible for the case’s category (the exact prompt is given in Fig. 7, Appendix C); $\text{recognized}(t)$ holds iff the frame is both detected and judged plausible. For instance, a human silhouette with an extra limb attached or a fragmented body fails recognition (Fig. 4, bottom right). The two rates are

$$D = \frac{|\{t : \text{detected}(t)\}|}{T}, \quad R = \frac{|\{t : \text{recognized}(t)\}|}{T}, \quad R \leq D.$$

Benchmark-level D and R are computed per case, then averaged over the 83 retained cases.

Object localization. On recognized frames, we compute MaskIoU between predicted and reference masks. The conditional score averages over recognized frames; the recognition-weighted score instead averages over all T frames, counting non-recognized frames as zero:

$$\begin{aligned} \text{cMaskIoU} &= \text{mean}_{t:\text{recognized}(t)} \text{MaskIoU}(t), \\ R \cdot \text{cMaskIoU} &= \frac{1}{T} \sum_{t=1}^T \text{MaskIoU}(t) \mathbf{1}[\text{recognized}(t)]. \end{aligned}$$

$R \cdot \text{cMaskIoU}$ is our primary localization metric: it rewards both rendering the correct subject and placing it at the correct target-view location.

5 Experiments

We apply the evaluation protocol of Sec. 4 to four camera-redirection baselines, presenting both a case-level (Fig. 5) and a benchmark-level (Table 1) comparison.

5.1 Baselines


We evaluate four camera-redirection systems on the 83 retained Redirect4D-Bench cases. All baselines receive the same source clip and target camera trajectory; we re-format the trajectory to match each method’s expected input where needed, but the camera path itself is unchanged, so output differences reflect the methods rather than the inputs.

(i) ReCamMaster (Bai et al., 2025a) is a video re-rendering model conditioned on the source video, a text prompt, and the target camera trajectory. We feed the benchmark trajectory directly into its native camera encoder, after resampling it to ReCamMaster’s expected inference length.

(ii) TrajectoryCrafter (Yu et al., 2025) estimates per-frame depth from the source video, unprojects it into a dynamic point cloud, renders the cloud along the target trajectory, and refines the result with a CogVideoX-based diffusion model (Yang et al., 2024b). For a fair comparison, we replace its DepthCrafter (Hu et al., 2025b) front-end with ViPE (Huang et al., 2025a), matching Gen3C and our construction pipeline.

(iii) Gen3C (Ren et al., 2025) builds a spatiotemporal 3D cache from the source video and renders it along the target camera path as the diffusion condition. We feed ViPE (Huang et al., 2025a) depth, poses, and intrinsics into its dynamic-video path, with each trajectory mapped to Gen3C’s native yaw/pitch/roll/scale parameters.

(iv) FreeOrbit4D (Cao et al., 2026a) performs training-free camera redirection by reconstructing a 4D scene from the source clip and refining the renders with video diffusion; we run its published configuration on each Redirect4D-Bench case without modification.



Method	Traditional metrics										Redirect4D-Bench metrics					
	CLIP ↑			VBench ↑							Camera acc. ↓		Object fidelity ↑			
	T	F	V	SC	BG	TF	MS	AQ	IQ	OC	RotErr	TransErr	D	R	cIoU	R-cIoU
ReCamMaster	0.317	0.983	0.935	0.931	0.911	0.962	0.984	0.507	0.709	0.219	27.313	0.727	1.000	0.756	0.095	0.071
TrajectoryCrafter	0.311	0.949	0.846	0.796	0.907	0.933	0.968	0.498	0.748	0.223	3.985	0.152	1.000	0.044	0.493	0.022
Gen3C	0.301	0.949	0.849	0.814	0.939	0.958	0.984	0.498	0.685	0.215	7.099	0.157	1.000	0.178	0.507	0.090
FreeOrbit4D	0.304	0.934	0.858	0.874	0.934	0.952	0.982	0.554	0.703	0.223	0.793	0.027	1.000	1.000	0.919	0.919

Figure 5 Case-level comparison on a camel case. (Left) the source clip and the four baselines’ generated target-view frames, each with a visual-check mark. (Right) (a) the foreground-complete 4D point cloud with the target trajectory arc; (b) a 3D plot overlaying the four baselines’ recovered trajectories on the target trajectory. The accompanying per-case table reports scores on the same sample. FID/FVD are omitted because they are aggregate distribution metrics; cIoU abbreviates cMaskIoU; **bold** marks the best score per column.

5.2 Results

Case-level comparison. Figure 5 shows a representative large-angle camel case. ReCamMaster produces plausible frames but does not follow the requested camera path; TrajectoryCrafter and Gen3C follow the trajectory more closely but distort the camel’s body; only FreeOrbit4D follows the path while keeping the camel intact. The traditional CLIP and VBench scores contradict this ranking, scattering “best” cells across all four methods including the visibly broken ones. The Redirect4D-Bench columns instead concentrate the best cells on FreeOrbit4D, because they separately measure the three aspects the visual check assesses: camera following, subject recognition, and target-view localization. A method that fails on any one of these is visible in the corresponding column instead of being averaged out by other strengths: ReCamMaster’s RotErr of 27.3° against FreeOrbit4D’s 0.79°, and R -cMaskIoU of 0.071 against 0.919, quantify exactly the failures visible in the frames. Seven additional per-case comparisons on bear, dancer, deer, and robot subjects (Figs. 8 to 14 in Appendix E) exhibit the same failure modes.

Benchmark-level comparison. Table 1 aggregates the comparison over all 83 retained cases. The same pattern persists. Traditional metric winners remain scattered: ReCamMaster takes FID/FVD and several CLIP/VBench columns, FreeOrbit4D wins CLIP-T, background consistency, and overall consistency, and TrajectoryCrafter wins imaging quality. These metrics reward realism, similarity, and temporal smoothness, but none reveals whether the requested redirection actually happened.

The Redirect4D-Bench metrics, in contrast, give a clear two-axis diagnosis. *Camera accuracy.* FreeOrbit4D, Gen3C, and TrajectoryCrafter cluster near 5° mean rotation error, while ReCamMaster drifts by 20.81°, roughly four times worse, confirming that ReCamMaster does not follow the benchmark trajectory at scale. *Object fidelity.* The four baselines split into two failure modes. TrajectoryCrafter and Gen3C detect the

Table 1 Benchmark-level comparison. Traditional metrics on the left, Redirect4D-Bench metrics on the right. CLIP (Radford et al., 2021): CLIP-T is text alignment, CLIP-F is adjacent-frame consistency, and CLIP-V is source-video consistency. VBench (Huang et al., 2024) columns SC/BG/TF/MS/AQ/IQ/OC stand for subject consistency, background consistency, temporal flickering, motion smoothness, aesthetic quality, imaging quality, and overall consistency. cIoU abbreviates cMaskIoU. **Bold** = best per column.

Method	Traditional metrics												Redirect4D-Bench metrics					
	CLIP $\uparrow$			Dist. $\downarrow$		VBench $\uparrow$							Camera acc. $\downarrow$		Object fidelity $\uparrow$			
	T	F	V	FID	FVD	SC	BG	TF	MS	AQ	IQ	OC	RotErr	TransErr	D	R	cIoU	R -cIoU
ReCamMaster	0.324	0.970	0.881	52.429	662.048	0.884	0.900	0.953	0.981	0.508	0.605	0.274	20.81	0.596	0.965	0.907	0.246	0.227
TrajectoryCrafter	0.324	0.958	0.832	83.559	1082.468	0.836	0.895	0.932	0.971	0.472	0.631	0.271	5.26	0.379	1.000	0.530	0.495	0.253
Gen3C	0.321	0.954	0.849	77.513	1047.264	0.839	0.902	0.942	0.980	0.480	0.543	0.272	5.25	0.312	1.000	0.596	0.486	0.289
FreeOrbit4D	0.326	0.953	0.843	67.305	925.250	0.860	0.915	0.937	0.976	0.496	0.610	0.282	5.10	0.280	1.000	0.990	0.907	0.898

subject reliably ($D = 1.000$) but fail recognition: their recognition rates fall to $R = 0.530$ and $R = 0.596$, i.e., only about half of their frames are judged structurally plausible by Qwen3-VL, consistent with the camel-case distortions. ReCamMaster shows the opposite, recognizing the subject in $R = 0.907$ of frames but placing it at the wrong target-view location (R -cMaskIoU = 0.227). Only FreeOrbit4D scores well on both axes (RotErr = 5.10°, R -cMaskIoU = 0.898), agreeing with the visual check on the camel case (Fig. 5).

Camera following and object fidelity need to be reported as separate axes: a method can follow the trajectory while breaking the subject, or preserve plausible appearance while missing the path. These two failure modes are distinct defect categories that traditional video metrics conflate into a single “visual quality” score. Ranking by traditional metrics alone risks selecting visibly broken outputs as winners, while the Redirect4D-Bench axes restore the link between metrics and visible failures.

6 Conclusion

Evaluating camera redirection on real monocular videos is held back by a missing-reference problem: no captured target view exists for an arbitrary camera path, so video-quality metrics measure realism rather than per-trajectory task success. Redirect4D-Bench addresses this with a scalable, training-free pipeline that produces per-trajectory pseudo ground truth for real dynamic scenes, and defines two task-specific metrics (camera accuracy and object fidelity) that agree with the visual check of the outputs where CLIP, FID/FVD, and VBench disagree. Across the four baselines, camera accuracy and object fidelity separate two failure modes that traditional metrics blur together: a method may follow the requested path but break the subject, or keep the subject intact but place it in the wrong target-view location. Reporting both metrics, not a single quality score, exposes each failure.

Limitations and future work. Errors from the multi-stage pipeline (foreground completion, background reconstruction, alignment) can accumulate in the pseudo-4D ground truth, which cannot be verified against real target views, and the pipeline assumes a single dominant subject. Because the pseudo ground truth is built with the same reconstruction approach as FreeOrbit4D (Cao et al., 2026a), FreeOrbit4D’s scores may benefit from this shared construction and are best read as a reference for future methods rather than as a neutral upper bound. Since construction is training-free, extending it to larger, multi-subject releases is a natural next step. Beyond dataset scalability, our benchmark could also support continual evaluation of video generation models as new categories, scenes, and camera trajectories emerge. This would enable future studies of continual learning (Liu et al., 2021a,b, 2023b,c, 2024b), examining whether models can acquire new camera-redirection capabilities while preserving their performance on previously learned scenarios.

Acknowledgments. This research is supported by the National Artificial Intelligence Research Resource Pilot Awards NAIRR250199, NAIRR260019, and NAIRR260077; AMD University Program’s AI & HPC Cluster; NVIDIA Academic Grant Program; Lambda’s Research Grant; and Google Cloud Research Credits Program. Computational resources are also provided by Delta and DeltaAI at the National Center for Supercomputing Applications through ACCESS allocations CIS250012, CIS250816, CIS250864, and CIS251188.

References

- Sami Abu-El-Haija, Nisarg Kothari, Joonseok Lee, Paul Natsev, George Toderici, Balakrishnan Varadarajan, and Sudheendra Vijayanarasimhan. YouTube-8M: A large-scale video classification benchmark. *arXiv preprint arXiv:1609.08675*, 2016.
- Ronald T. Azuma. A survey of augmented reality. *Presence: Teleoperators and Virtual Environments*, 6(4):355–385, 1997.
- Sherwin Bahmani, Ivan Skorokhodov, Aliaksandr Siarohin, Willi Menapace, Guocheng Qian, Michael Vasilkovsky, Hsin-Ying Lee, Chaoyang Wang, Jiaxu Zou, Andrea Tagliasacchi, et al. VD3D: Taming large video diffusion transformers for 3D camera control. *arXiv preprint arXiv:2407.12781*, 2024.
- Sherwin Bahmani, Ivan Skorokhodov, Guocheng Qian, Aliaksandr Siarohin, Willi Menapace, Andrea Tagliasacchi, David B. Lindell, and Sergey Tulyakov. AC3D: Analyzing and improving 3D camera control in video diffusion transformers. In *CVPR*, pages 22875–22889, 2025.
- Jianhong Bai, Menghan Xia, Xiao Fu, Xintao Wang, Lianrui Mu, Jinwen Cao, Zuozhu Liu, Haoji Hu, Xiang Bai, Pengfei Wan, et al. ReCamMaster: Camera-controlled generative rendering from a single video. In *ICCV*, pages 14834–14844, 2025a.
- Shuai Bai, Yuxuan Cai, Ruizhe Chen, Keqin Chen, Xionghui Chen, Zesen Cheng, Lianghao Deng, Wei Ding, Chang Gao, Chunjiang Ge, et al. Qwen3-VL technical report. *arXiv preprint arXiv:2511.21631*, 2025b.
- Max Bain, Arsha Nagrani, Gül Varol, and Andrew Zisserman. Frozen in time: A joint video and image encoder for end-to-end retrieval. In *ICCV*, pages 1728–1738, 2021.
- Wei Cao, Chang Luo, Biao Zhang, Matthias Nießner, and Jiapeng Tang. Motion2VecSets: 4D latent vector set diffusion for non-rigid shape reconstruction and tracking. In *CVPR*, pages 20496–20506, 2024.
- Wei Cao, Marcel Hallgarten, Tianyu Li, Daniel Dauner, Xunjiang Gu, Caojun Wang, Yakov Miron, Marco Aiello, Hongyang Li, Igor Gilitschenski, et al. Pseudo-simulation for autonomous driving. *arXiv preprint arXiv:2506.04218*, 2025.
- Wei Cao, Hao Zhang, Fengrui Tian, Yulun Wu, Yingying Li, Shenlong Wang, Ning Yu, and Yaoyao Liu. FreeOrbit4D: Training-free arbitrary camera redirection for monocular videos via foreground-complete 4D reconstruction. In *ACM SIGGRAPH Conference Papers*, 2026a.
- Wei Cao, Hao Zhang, Vikram Voleti, Yuqun Wu, Shimon Vainer, Mark Boss, Yaoyao Liu, et al. 4Director: Controlling video world models with rigid 3D geometry. *arXiv preprint arXiv:2610.02160*, 2026b.
- Nicolas Carion, Laura Gustafson, Yuan-Ting Hu, Shoubhik Debnath, Ronghang Hu, Didac Suris, Chaitanya Ryali, Kalyan Vasudev Alwala, Haitham Khedr, Andrew Huang, Jie Lei, Tengyu Ma, Baishan Guo, Arpit Kalla, Markus Marks, Joseph Greer, Meng Wang, Peize Sun, Roman Rädle, Triantafyllos Afouras, Effrosyni Mavroudi, Katherine Xu, Tsung-Han Wu, Yu Zhou, Liliane Momeni, Rishi Hazra, Shuangrui Ding, Sagar Vaze, Francois Porcher, Feng Li, Siyuan Li, Aishwarya Kamath, Ho Kei Cheng, Piotr Dollár, Nikhila Ravi, Kate Saenko, Pengchuan Zhang, and Christoph Feichtenhofer. SAM 3: Segment anything with concepts. *arXiv preprint arXiv:2511.16719*, 2025.
- David Charatan, Sizhe Lester Li, Andrea Tagliasacchi, and Vincent Sitzmann. pixelSplat: 3D gaussian splats from image pairs for scalable generalizable 3D reconstruction. In *CVPR*, pages 19457–19467, 2024.
- Tsai-Shien Chen, Aliaksandr Siarohin, Willi Menapace, Ekaterina Deyneka, Hsiang-Wei Chao, Byung Eun Jeon, Yuwei Fang, Hsin-Ying Lee, Jian Ren, Ming-Hsuan Yang, et al. Panda-70M: Captioning 70m videos with multiple cross-modality teachers. In *CVPR*, 2024a.
- Yuedong Chen, Haofei Xu, Chuanxia Zheng, Bohan Zhuang, Marc Pollefeys, Andreas Geiger, Tat-Jen Cham, and Jianfei Cai. MVSplat: Efficient 3D gaussian splatting from sparse multi-view images. In *ECCV*, pages 370–386, 2024b.
- Chen Gao, Ayush Saraf, Johannes Kopf, and Jia-Bin Huang. Dynamic view synthesis from dynamic monocular video. In *ICCV*, pages 5712–5721, 2021.
- Hang Gao, Ruilong Li, Shubham Tulsiani, Bryan Russell, and Angjoo Kanazawa. Monocular dynamic view synthesis: A reality check. In *NeurIPS*, 2022.

- Klaus Greff, Francois Belletti, Lucas Beyer, Carl Doersch, Yilun Du, Daniel Duckworth, David J. Fleet, Dan Gnanapragasam, Florian Golemo, Charles Herrmann, et al. Kubric: A scalable dataset generator. In *CVPR*, pages 3749–3761, 2022.
- Zekai Gu, Rui Yan, Jiahao Lu, Peng Li, Zhiyang Dou, Chenyang Si, Zhen Dong, Qifeng Liu, Cheng Lin, Ziwei Liu, et al. Diffusion as shader: 3D-aware video diffusion for versatile video generation control. In *ACM SIGGRAPH Conference Papers*, pages 1–12, 2025.
- Hao He, Yinghao Xu, Yuwei Guo, Gordon Wetzstein, Bo Dai, Hongsheng Li, and Ceyuan Yang. CameraCtrl: Enabling camera control for text-to-video generation. *arXiv preprint arXiv:2404.02101*, 2024.
- Fabian Caba Heilbron, Victor Escorcia, Bernard Ghanem, and Juan Carlos Niebles. Activitynet: A large-scale video benchmark for human activity understanding. In *CVPR*, pages 961–970, 2015.
- Martin Heusel, Hubert Ramsauer, Thomas Unterthiner, Bernhard Nessler, and Sepp Hochreiter. GANs trained by a two time-scale update rule converge to a local Nash equilibrium. In *NeurIPS*, volume 30, 2017.
- Chen Hou and Zhibo Chen. Training-free camera control for video generation. *arXiv preprint arXiv:2406.10126*, 2024.
- Tao Hu, Haoyang Peng, Xiao Liu, and Yuwen Ma. EX-4D: Extreme viewpoint 4D video synthesis via depth watertight mesh. *arXiv preprint arXiv:2506.05554*, 2025a.
- Wenbo Hu, Xiangjun Gao, Xiaoyu Li, Sijie Zhao, Xiaodong Cun, Yong Zhang, Long Quan, and Ying Shan. DepthCrafter: Generating consistent long depth sequences for open-world videos. In *CVPR*, 2025b.
- Jiahui Huang, Qunjie Zhou, Hesam Rabeti, Aleksandr Korovko, Huan Ling, Xuanchi Ren, Tianchang Shen, Jun Gao, Dmitry Slepichev, Chen-Hsuan Lin, Jiawei Ren, Kevin Xie, Joydeep Biswas, Laura Leal-Taixe, and Sanja Fidler. ViPE: Video pose engine for 3D geometric perception. *arXiv preprint arXiv:2508.10934*, 2025a.
- Zhenning Huang, Hyeonho Jeong, Xuelin Chen, Yulia Gryaditskaya, Tuanfeng Y. Wang, Joan Lasenby, and Chun-Hao Huang. SpaceTimePilot: Generative rendering of dynamic scenes across space and time. *arXiv preprint arXiv:2512.25075*, 2025b.
- Ziqi Huang, Yinan He, Jiashuo Yu, Fan Zhang, Chenyang Si, Yuming Jiang, Yuanhan Zhang, Tianxing Wu, Qingyang Jin, Nattapol Chanpaisit, Yaohui Wang, Xinyuan Chen, Limin Wang, Dahua Lin, Yu Qiao, and Ziwei Liu. VBench: Comprehensive benchmark suite for video generative models. In *CVPR*, pages 21807–21818, 2024.
- Quan Huynh-Thu and Mohammed Ghanbari. Scope of validity of PSNR in image/video quality assessment. *Electronics Letters*, 44(13):800–801, 2008.
- Eric Ji, Qiran Hu, Wufei Ma, Sarthak Jain, Yingying Li, Minh N Do, and Yaoyao Liu. AC3S: Adaptive conditioning for 3D-aware synthetic data generation. In *ECCV*, pages 510–525, 2026.
- Lihan Jiang, Yucheng Mao, Linning Xu, Tao Lu, Kerui Ren, Yichen Jin, Xudong Xu, Mulin Yu, Jiangmiao Pang, Feng Zhao, et al. AnySplat: Feed-forward 3D gaussian splatting from unconstrained views. *ACM Transactions on Graphics*, 44(6):1–16, 2025a.
- Zeren Jiang, Chuanxia Zheng, Iro Laina, Diane Larlus, and Andrea Vedaldi. Geo4D: Leveraging video generators for geometric 4D scene reconstruction. In *ICCV*, pages 20658–20671, 2025b.
- Xuan Ju, Yiming Gao, Zhaoyang Zhang, Ziyang Yuan, Xintao Wang, Ailing Zeng, Yu Xiong, Qiang Xu, and Ying Shan. MiraData: A large-scale video dataset with long durations and structured captions. *NeurIPS*, 37:48955–48970, 2024.
- Will Kay, Joao Carreira, Karen Simonyan, Brian Zhang, Chloe Hillier, Sudheendra Vijayanarasimhan, Fabio Viola, Tim Green, Trevor Back, Paul Natsev, et al. The kinetics human action video dataset. *arXiv preprint arXiv:1705.06950*, 2017.
- Bernhard Kerbl, Georgios Kopanas, Thomas Leimkuehler, and George Drettakis. 3D gaussian splatting for real-time radiance field rendering. *ACM Transactions on Graphics*, 42(4):139:1–139:14, 2023.
- Zhengfei Kuang, Shengqu Cai, Hao He, Yinghao Xu, Hongsheng Li, Leonidas J. Guibas, and Gordon Wetzstein. Collaborative video diffusion: Consistent multi-video generation with camera control. *NeurIPS*, 37:16240–16271, 2024.
- Jiahui Lei, Yijia Weng, Adam W. Harley, Leonidas Guibas, and Kostas Daniilidis. MoSca: Dynamic gaussian fusion from casual videos via 4D motion scaffolds. In *CVPR*, pages 6165–6177, 2025.

- Ruilong Li, Brent Yi, Junchen Liu, Hang Gao, Yi Ma, and Angjoo Kanazawa. Cameras as relative positional encoding. *arXiv preprint arXiv:2507.10496*, 2025a.
- Teng Li, Guangcong Zheng, Rui Jiang, Shuigen Zhan, Tao Wu, Yehao Lu, Yining Lin, Chuanyun Deng, Yegan Xiong, Min Chen, et al. RealCam-I2V: Real-world image-to-video generation with interactive complex camera control. In *ICCV*, pages 28785–28796, 2025b.
- Zhengqi Li, Richard Tucker, Forrester Cole, Qianqian Wang, Linyi Jin, Vickie Ye, Angjoo Kanazawa, Aleksander Holynski, and Noah Snavely. MegaSAM: Accurate, fast and robust structure and motion from casual dynamic videos. In *CVPR*, pages 10486–10496, 2025c.
- Kuan Heng Lin, Zhizheng Liu, Pablo Salamanca, Yash Kant, Ryan Burgert, Yuancheng Xu, Koichi Namekata, Yiwei Zhao, Bolei Zhou, Micah Goldblum, Paul Debevec, and Ning Yu. Vista4D: Video reshooting with 4D point clouds. In *CVPR*, 2026.
- Zhiqiu Lin, Siyuan Cen, Daniel Jiang, Jay Karhade, Hewei Wang, Chancharik Mitra, Tiffany Ling, Yuhang Huang, Sifan Liu, Mingyu Chen, et al. Towards understanding camera motions in any video. *arXiv preprint arXiv:2504.15376*, 2025.
- Lu Ling, Yichen Sheng, Zhi Tu, Wentian Zhao, Cheng Xin, Kun Wan, Lantao Yu, Qianyu Guo, Zixun Yu, Yawen Lu, et al. DL3DV-10K: A large-scale scene dataset for deep learning-based 3D vision. In *CVPR*, pages 22160–22169, 2024.
- Xinran Ling, Chen Zhu, Meiqi Wu, Hangyu Li, Xiaokun Feng, Cundian Yang, Aiming Hao, Jiashu Zhu, Jiahong Wu, and Xiangxiang Chu. VMBench: A benchmark for perception-aligned video motion generation. In *ICCV*, pages 13087–13098, 2025.
- Shilong Liu, Zhaoyang Zeng, Tianhe Ren, Feng Li, Hao Zhang, Jie Yang, Chunyuan Li, Jianwei Yang, Hang Su, Jun Zhu, et al. Grounding DINO: Marrying DINO with grounded pre-training for open-set object detection. *arXiv preprint arXiv:2303.05499*, 2023a.
- Yaofang Liu, Xiaodong Cun, Xuebo Liu, Xintao Wang, Yong Zhang, Haoxin Chen, Yang Liu, Tiejiong Zeng, Raymond Chan, and Ying Shan. EvalCrafter: Benchmarking and evaluating large video generation models. In *CVPR*, pages 22139–22149, 2024a.
- Yaoyao Liu, Bernt Schiele, and Qianru Sun. Adaptive aggregation networks for class-incremental learning. In *CVPR*, pages 2544–2553, 2021a.
- Yaoyao Liu, Bernt Schiele, and Qianru Sun. RMM: Reinforced memory management for class-incremental learning. *NeurIPS*, 34:3478–3490, 2021b.
- Yaoyao Liu, Yingying Li, Bernt Schiele, and Qianru Sun. Online hyperparameter optimization for class-incremental learning. In *AAAI*, volume 37, pages 8906–8913, 2023b.
- Yaoyao Liu, Bernt Schiele, Andrea Vedaldi, and Christian Rupprecht. Continual detection transformer for incremental object detection. In *CVPR*, pages 23799–23808, 2023c.
- Yaoyao Liu, Yingying Li, Bernt Schiele, and Qianru Sun. Wakening past concepts without past data: Class-incremental learning from online placebos. In *WACV*, pages 2215–2224, 2024b.
- Yawen Luo, Xiaoyu Shi, Jianhong Bai, Menghan Xia, Tianfan Xue, Xintao Wang, Pengfei Wan, Di Zhang, and Kun Gai. CamCloneMaster: Enabling reference-based camera control for video generation. In *ACM SIGGRAPH Asia Conference Papers*, pages 1–10, 2025.
- Wufei Ma, Qihao Liu, Jiahao Wang, Angtian Wang, Xiaoding Yuan, Yi Zhang, Zihao Xiao, Guofeng Zhang, Beijia Lu, Ruxiao Duan, et al. Generating images with 3d annotations using diffusion models. In *ICLR*, 2024a.
- Wufei Ma, Guofeng Zhang, Qihao Liu, Guanning Zeng, Adam Kortylewski, Yaoyao Liu, and Alan Yuille. Imagenet3d: Towards general-purpose object-level 3d understanding. In *NeurIPS*, pages 96127–96149, 2024b.
- Antoine Miech, Dimitri Zhukov, Jean-Baptiste Alayrac, Makarand Tapaswi, Ivan Laptev, and Josef Sivic. HowTo100M: Learning a text-video embedding by watching hundred million narrated video clips. In *ICCV*, pages 2630–2640, 2019.
- Ben Mildenhall, Pratul P. Srinivasan, Matthew Tancik, Jonathan T. Barron, Ravi Ramamoorthi, and Ren Ng. NeRF: Representing scenes as neural radiance fields for view synthesis. *Communications of the ACM*, 65(1):99–106, 2021.

- Kepan Nan, Rui Xie, Penghao Zhou, Tiehan Fan, Zhenheng Yang, Zhijie Chen, Xiang Li, Jian Yang, and Ying Tai. OpenVid-1M: A large-scale high-quality dataset for text-to-video generation. *arXiv preprint arXiv:2407.02371*, 2024.
- Alec Radford, Jong Wook Kim, Chris Hallacy, Aditya Ramesh, Gabriel Goh, Sandhini Agarwal, Girish Sastry, Amanda Askell, Pamela Mishkin, Jack Clark, et al. Learning transferable visual models from natural language supervision. In *ICML*, pages 8748–8763. PMLR, 2021.
- Nikhila Ravi, Valentin Gabeur, Yuan-Ting Hu, Ronghang Hu, Chaitanya Ryali, Tengyu Ma, Haitham Khedr, Roman Rädle, Chloe Rolland, Laura Gustafson, Eric Mintun, Junting Pan, Kalyan Vasudev Alwala, Nicolas Carion, Chao-Yuan Wu, Ross Girshick, Piotr Dollár, and Christoph Feichtenhofer. SAM 2: Segment anything in images and videos. *arXiv preprint arXiv:2408.00714*, 2024.
- Tianhe Ren, Shilong Liu, Ailing Zeng, Jing Lin, Kunchang Li, He Cao, Jiayu Chen, Xinyu Huang, Yukang Chen, Feng Yan, et al. Grounded SAM: Assembling open-world models for diverse visual tasks. *arXiv preprint arXiv:2401.14159*, 2024.
- Xuanchi Ren, Tianchang Shen, Jiahui Huang, Huan Ling, Yifan Lu, Merlin Nimier-David, Thomas Müller, Alexander Keller, Sanja Fidler, and Jun Gao. Gen3C: 3D-informed world-consistent video generation with precise camera control. In *CVPR*, pages 6121–6132, 2025.
- Chris Rockwell, Joseph Tung, Tsung-Yi Lin, Ming-Yu Liu, David F. Fouhey, and Chen-Hsuan Lin. Dynamic camera poses and where to find them. In *CVPR*, pages 12444–12455, 2025.
- Pei Sun, Henrik Kretschmar, Xerxes Dotiwalla, Aurelien Chouard, Vijaysai Patnaik, Paul Tsui, James Guo, Yin Zhou, Yuning Chai, Benjamin Caine, et al. Scalability in perception for autonomous driving: Waymo open dataset. In *CVPR*, pages 2446–2454, 2020.
- Matthew Tancik, Ethan Weber, Evonne Ng, Ruilong Li, Brent Yi, Terrance Wang, Alexander Kristoffersen, Jake Austin, Kamyar Salahi, Abhik Ahuja, et al. Nerfstudio: A modular framework for neural radiance field development. In *ACM SIGGRAPH 2023 Conference Proceedings*, pages 1–12, 2023.
- Jiapeng Tang, Wei Cao, Biao Zhang, Chang Luo, Yaoyao Liu, and Matthias Nießner. Motion2VecSets: Non-rigid shape reconstruction and tracking with 4D latent set diffusion. *IEEE Transactions on Pattern Analysis and Machine Intelligence*, 2026.
- Zachary Teed and Jia Deng. Droid-slam: Deep visual slam for monocular, stereo, and rgb-d cameras. *NeurIPS*, 34: 16558–16569, 2021.
- Thomas Unterthiner, Sjoerd van Steenkiste, Karol Kurach, Raphael Marinier, Marcin Michalski, and Sylvain Gelly. FVD: A new metric for video generation. In *ICLR 2019 Workshops*, 2019.
- Basile Van Hoorick, Rundu Wu, Ege Ozguroglu, Kyle Sargent, Ruoshi Liu, Pavel Tokmakov, Achal Dave, Changxi Zheng, and Carl Vondrick. Generative camera dolly: Extreme monocular dynamic novel view synthesis. In *ECCV*, pages 313–331. Springer, 2024.
- Ang Wang, Baole Ai, Bin Wen, Chaojie Mao, Chen-Wei Xie, Di Chen, Fei Wu Yu, Haiming Zhao, Jianxiao Yang, Jianyuan Zeng, et al. Wan: Open and advanced large-scale video generative models. *arXiv preprint arXiv:2503.20314*, 2025a.
- Jiahao Wang, Yufeng Yuan, Rujie Zheng, Youtian Lin, Jian Gao, Lin-Zhuo Chen, Yajie Bao, Yi Zhang, Chang Zeng, Yanxi Zhou, et al. SpatialVID: A large-scale video dataset with spatial annotations. *arXiv preprint arXiv:2509.09676*, 2025b.
- Jianyuan Wang, Minghao Chen, Nikita Karaev, Andrea Vedaldi, Christian Rupprecht, and David Novotny. VGGT: Visual geometry grounded transformer. In *CVPR*, pages 5294–5306, 2025c.
- Qianqian Wang, Vickie Ye, Hang Gao, Weijia Zeng, Jake Austin, Zhengqi Li, and Angjoo Kanazawa. Shape of motion: 4D reconstruction from a single video. In *ICCV*, pages 9660–9672, 2025d.
- Yiming Wang, Qihang Zhang, Shengqu Cai, Tong Wu, Jan Ackermann, Zhengfei Kuang, Yang Zheng, Frano Rajič, Siyu Tang, and Gordon Wetzstein. BulletTime: Decoupled control of time and camera pose for video generation. *arXiv preprint arXiv:2512.05076*, 2025e.
- Zhaoqing Wang, Xiaobo Xia, Zhuolin Bie, Jinlin Liu, Dongdong Yu, Jia-Wang Bian, and Changhu Wang. Taming camera-controlled video generation with verifiable geometry reward. *arXiv preprint arXiv:2512.02870*, 2025f.
- Zhou Wang, Alan C. Bovik, Hamid R. Sheikh, and Eero P. Simoncelli. Image quality assessment: From error visibility to structural similarity. *IEEE Transactions on Image Processing*, 13(4):600–612, 2004.

- Zhouxia Wang, Ziyang Yuan, Xintao Wang, Yaowei Li, Tianshui Chen, Menghan Xia, Ping Luo, and Ying Shan. MotionCtrl: A unified and flexible motion controller for video generation. In *ACM SIGGRAPH Conference Papers*, pages 1–11, 2024.
- Yulun Wu, Ruyi Zha, Wei Cao, Yingying Li, Yuanhao Cai, and Yaoyao Liu. Active view selection with perturbed gaussian ensemble for tomographic reconstruction. *arXiv preprint arXiv:2603.06852*, 2026.
- Hongwei Xue, Tiankai Hang, Yanhong Zeng, Yuchong Sun, Bei Liu, Huan Yang, Jianlong Fu, and Baining Guo. Advancing high-resolution video-language representation with large-scale video transcriptions. In *CVPR*, pages 5036–5045, 2022.
- Lihe Yang, Bingyi Kang, Zilong Huang, Zhen Zhao, Xiaogang Xu, Jiashi Feng, and Hengshuang Zhao. Depth anything v2. *NeurIPS*, 37:21875–21911, 2024a.
- Zhuoyi Yang, Jiayan Teng, Wendi Zheng, Ming Ding, Shiyu Huang, Jiazheng Xu, Yuanming Yang, Wenyi Hong, Xiaohan Zhang, Guanyu Feng, et al. CogVideoX: Text-to-video diffusion models with an expert transformer. *arXiv preprint arXiv:2408.06072*, 2024b.
- Chun-Han Yao, Yiming Xie, Vikram Voleti, Huaizu Jiang, and Varun Jampani. SV4D 2.0: Enhancing spatio-temporal consistency in multi-view video diffusion for high-quality 4D generation. In *ICCV*, pages 13248–13258, 2025.
- Meng You, Zhiyu Zhu, Hui Liu, and Junhui Hou. NVS-Solver: Video diffusion model as zero-shot novel view synthesizer. *arXiv preprint arXiv:2405.15364*, 2024.
- Mark Yu, Wenbo Hu, Jinbo Xing, and Ying Shan. TrajectoryCrafter: Redirecting camera trajectory for monocular videos via diffusion models. In *ICCV*, pages 100–111, 2025.
- 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.
- David Junhao Zhang, Roni Paiss, Shiran Zada, Nikhil Karnad, David E. Jacobs, Yael Pritch, Inbar Mosseri, Mike Zheng Shou, Neal Wadhwa, and Nataniel Ruiz. ReCapture: Generative video camera controls for user-provided videos using masked video fine-tuning. In *CVPR*, pages 2050–2062, 2025a.
- Hao Zhang, Chun-Han Yao, Simon Donné, Narendra Ahuja, and Varun Jampani. Stable part diffusion 4D: Multi-view RGB and kinematic parts video generation. *arXiv preprint arXiv:2509.10687*, 2025b.
- Junyi Zhang, Charles Herrmann, Junhwa Hur, Varun Jampani, Trevor Darrell, Forrester Cole, Deqing Sun, and Ming-Hsuan Yang. MonST3R: A simple approach for estimating geometry in the presence of motion. *arXiv preprint arXiv:2410.03825*, 2024.
- Richard Zhang, Phillip Isola, Alexei A. Efros, Eli Shechtman, and Oliver Wang. The unreasonable effectiveness of deep features as a perceptual metric. In *CVPR*, pages 586–595, 2018.
- Brian Nlong Zhao, Jiajun Wu, and Shangzhe Wu. Web-scale collection of video data for 4D animal reconstruction. *arXiv preprint arXiv:2511.01169*, 2025.
- Sixiao Zheng, Minghao Yin, Wenbo Hu, Xiaoyu Li, Ying Shan, and Yanwei Fu. VerseCrafter: Dynamic realistic video world model with 4D geometric control. *arXiv preprint arXiv:2601.05138*, 2026.
- Hongkuan Zhou, Wei Cao, Aifen Sui, and Zhenshan Bing. What matters to enhance traffic rule compliance of imitation learning for end-to-end autonomous driving. *arXiv preprint arXiv:2309.07808*, 2023.
- Hongkuan Zhou, Stefan Schmid, Yicong Li, Lavdim Halilaj, Xiangtong Yao, and Wei Cao. Predicting the road ahead: A knowledge graph based foundation model for scene understanding in autonomous driving. In *European Semantic Web Conference*, pages 116–132. Springer, 2025a.
- Kaichen Zhou, Yuhan Wang, Grace Chen, Xinhai Chang, Gaspard Beaudouin, Fangneng Zhan, Paul Pu Liang, and Mengyu Wang. PAGE-4D: Disentangled pose and geometry estimation for VGGT-4D perception. *arXiv preprint arXiv:2510.17568*, 2025b.
- Tinghui Zhou, Richard Tucker, John Flynn, Graham Fyfe, and Noah Snavely. Stereo magnification: learning view synthesis using multiplane images. *ACM Transactions on Graphics*, 37(4):65, 2018.

Appendix

A Dataset Statistics

Figure 6 summarizes the retained cases: yaw values concentrate between $\pm 90^\circ$ and $\pm 120^\circ$ with pitch mostly in $[-30^\circ, 0^\circ]$, and the 62 source clips split into 24 robot, 23 human, 14 animal, and 1 other.




Figure 6 Dataset statistics. (a) Yaw and pitch parameters for the retained target trajectories. Each point is one redirection case; color denotes the subject family (robot, human, animal, other) and marker size denotes the radius scale (0.7, 1.0, or 1.4). Yaw values cluster between $\pm 90^\circ$ and $\pm 120^\circ$, with most pitch values in $[-30^\circ, 0^\circ]$. (b) Source-clip counts across subject families in the current release: robot 24 (39%), human 23 (37%), animal 14 (23%), other 1 ($\sim 2\%$), totaling the 62 source clips that yield 83 redirection cases.

B Safety Filtering and Release Policy

Safety filtering. Of the 103 target trajectories produced by the construction pipeline, we exclude 20 cases that involve minors as the foreground subject, leaving the 83 retained cases on which all quantitative results in the paper are computed.

Release policy. The benchmark is released on Hugging Face; links to the dataset and code are available from the project page (<https://redirect4d-bench.cvlgroup.web.illinois.edu/>). The public dataset contains derived assets (4D point clouds, target trajectories, rendered geometry, subject masks, prompts, and metadata) plus a small public sample for inspection. Original source RGB is not redistributed; following Zhao et al. (2025), the released code provides a recovery script using YouTube identifiers in the metadata.

C Qwen3-VL Subject-Recognition Prompt

The subject-recognition component in Sec. 4.4 uses Qwen3-VL-32B (Bai et al., 2025b) as a binary judge: given a cropped generated frame and a subject class {cls}, the judge decides whether the foreground subject is catastrophically broken (e.g., extra or missing limbs, fragmented body, fused parts). The prompt template is shown in Fig. 7; separate defect lists are used for biological/articulated and mechanical subjects, and the model’s textual response is parsed into a binary recognition decision.

D Compute Resources

All experiments run on a single internal node with $8 \times$ NVIDIA A100-SXM4-80GB GPUs.

```

[Shared Instruction]
You are shown a single frame cropped from an AI-generated video. The crop contains a {cls}. The
camera angle may be unusual.

Question: is the {cls} catastrophically broken, i.e., so badly deformed that it is clearly not a
real {cls}?

[Biological / Articulated Subjects]
Only answer "yes" if ONE of these is undeniably visible:
  a. DUPLICATION: two heads, two tails, two torsos, or two complete instances of a body part
  that should appear only once.
  b. WRONG LIMB COUNT: a clearly visible extra limb, or a main limb entirely missing where
  anatomy demands it.
  c. MELTED / BLOBBED: the {cls}'s body or limbs have dissolved into a continuous amorphous blob
  with no recognizable structure.
  d. REVERSED JOINT: a knee, elbow, hock, wrist, or ankle bends in the opposite direction to the
  species' natural anatomy.
If the {cls} looks like a plausible albeit imperfect rendering of a real {cls}, answer "no".
Imperfect details, blur, odd poses, long or S-curved necks, humps, trunks, stripes, fur / skin
texture, dance poses, unusual camera angles, cropped parts, self-occlusion, and natural
class-specific features are NEVER defects.

[Mechanical Subjects]
Only answer "yes" if ONE of these is undeniably visible:
  a. DUPLICATION: two of a main component that should appear only once (e.g., two gear hubs, two
  robot torsos / heads, or two complete structural pieces).
  b. MELTED / BLOBBED: structural parts have dissolved into a continuous amorphous blob with no
  recognizable mechanical form.
  c. IMPOSSIBLE STRUCTURE: a rigid component fluidly bent like rubber, a hinge twisted beyond
  its physical range, or a snapped structural piece floating apart.
  d. MIS-ATTACHED PART: a main component sprouts from a clearly impossible location, such as a
  robot limb attached at the head.
If the {cls} looks like a plausible albeit imperfect rendering of a real {cls}, answer "no".
Unusual camera angles, cropped parts, stylized aesthetics, visible internal mechanisms, unusual
colors, minor jitter, and class-specific mechanical features are NEVER defects.

[Output Rule]
When in doubt, always answer "no". Flag "yes" only for defects that would stop someone scrolling
past on social media to say "wait, that's wrong."

Respond with exactly one line. Start with "no" if the {cls} is plausible; otherwise start with
"yes" followed by a brief description of which part is catastrophically broken.

```

Figure 7 Qwen3-VL recognition prompt. The prompt receives a crop from a generated frame and a subject class placeholder {cls}. Separate defect lists are used for biological/articulated and mechanical subjects. The response is parsed as a binary recognition decision: catastrophic defects make the frame invalid, while plausible but imperfect renderings remain valid.

4D-reconstruction pipeline. The construction pipeline (Sec. 3) is a one-time cost per dataset version. End-to-end on the 83 retained cases it consumes roughly 40 A100-hours, dominated by foreground completion (SV4D 2.0 (Yao et al., 2025) pseudo multi-view synthesis and VGGT-1B (Wang et al., 2025c) point-map prediction), background reconstruction (ViPE (Huang et al., 2025a) lyra configuration), and target-view rendering with Wan2.2 VACE (Wang et al., 2025a). Per-stage wall times on 8 GPUs: 5 min for source-RGB caching, 5 min for SAM source-mask sharding, and 5 h for the reconstruction-render-refine loop.

Baseline inference. Each baseline is run once per case on a single A100-80GB. Per-case averages over the 83 retained cases are: ReCamMaster (Bai et al., 2025a) ~5.7 min/case (~8 A100-hours total), TrajectoryCrafter (Yu et al., 2025) ~7 min/case (~10 A100-hours total), Gen3C (Ren et al., 2025) ~17 min/case (~23 A100-hours total), and FreeOrbit4D (Cao et al., 2026a) ~10 min/case (~14 A100-hours total). Sharded over 8 GPUs the wall-clock times are approximately 1 h, 1.5 h, 2 h 50 min, and 1.7 h respectively.

Metric evaluation. Per evaluated method, the camera-accuracy protocol re-runs ViPE (Huang et al., 2025a) on the generated video; this takes ~3.6 A100-hours total (~27 min wall on 8 GPUs). Subject-mask extraction with SAM (Carion et al., 2025) costs under 1 A100-hour per method (~5 min wall per shard on 8 GPUs). Subject recognition with Qwen3-VL-32B (Bai et al., 2025b) in 4-bit quantization costs ~10 A100-hours per method (~75 min wall on 8 GPUs).

Aggregate budget. Producing all numbers reported in the paper (the four-baseline benchmark plus all camera-accuracy and object-fidelity metrics on 83 cases) takes roughly 200 A100-80GB GPU-hours on top of the one-time ~ 40 -hour construction cost. This excludes preliminary runs and failed configurations that are not reported in the paper, which together used a comparable amount of compute, primarily for tuning baseline trajectory encoding and validating the metric pipeline.

E Additional Per-Case Comparisons

We provide seven additional per-case comparisons (Figs. 8 to 14) using the same layout as Fig. 5. Each figure displays, on the left, the input source video and the four baselines’ generated target-view videos, with a visual-check mark per baseline; on the right, panel (a) shows the foreground-complete 4D point cloud with the target trajectory arc, and panel (b) overlays the target trajectory with each baseline’s ViPE-recovered trajectory in 3D. The accompanying per-case table reports the same columns as Table 1 except that aggregate distribution metrics (FID/FVD) are omitted; cIoU abbreviates cMaskIoU, and **bold** marks the best score per column. Each case is chosen to expose a specific discrepancy between traditional video metrics (CLIP (Radford et al., 2021), VBench (Huang et al., 2024)) and the visual check: traditional scores frequently reward outputs that are visibly broken in ways the Redirect4D-Bench metrics surface directly.




Figure 8 Bear case. ReCamMaster (Bai et al., 2025a) scores best on six of the ten CLIP/VBench columns by emitting frames close to the source view: the bear stays intact ($R = 1.000$), but the camera barely moves ($\text{RotErr} = 31.9^\circ$) and the bear is left in the wrong target-view location ($R\text{-cMaskIoU} = 0.402$). TrajectoryCrafter (Yu et al., 2025) follows the arc ($\text{RotErr} = 4.12^\circ$) but twists the bear’s body ($R = 0.444$). Gen3C (Ren et al., 2025) fails on both axes: its camera overshoots the arc ($\text{RotErr} = 24.0^\circ$) and the bear is visibly deformed in the later frames. Only FreeOrbit4D (Cao et al., 2026a) follows the arc to the rear view ($\text{RotErr} = 3.53^\circ$) and keeps the bear intact and correctly placed ($R\text{-cMaskIoU} = 0.981$).




Figure 9 Dancer case 1. This trajectory combines yaw, pitch, and roll, and every method has RotErr above 12° . ReCamMaster (Bai et al., 2025a) scores best on three VBench columns and keeps the dancer intact ($R = 1.000$), but its camera drifts 24.9° off the arc and the dancer slides to the frame border, far from her target-view location ($R\text{-cMaskIoU} = 0.043$). TrajectoryCrafter (Yu et al., 2025) passes the visual check, yet its camera overshoots the arc (largest TransErr = 1.157). Gen3C (Ren et al., 2025) distorts the background and the dancer in the later frames ($R = 0.800$). FreeOrbit4D (Cao et al., 2026a) has the lowest pose errors (RotErr = 12.1°) and places the dancer best ($R\text{-cMaskIoU} = 0.886$ vs. at most 0.308 for the other methods).




Figure 10 Dancer case 2. ReCamMaster (Bai et al., 2025a) keeps the ballerina intact and passes the visual check, but its camera barely orbits (RotErr = 24.5° , $R\text{-cMaskIoU} = 0.408$): a trajectory failure that plausible-looking frames can hide and that RotErr exposes. TrajectoryCrafter (Yu et al., 2025) (RotErr = 1.95°) and Gen3C (Ren et al., 2025) (RotErr = 2.78°) track the trajectory closely but break the ballerina’s limbs ($R = 0.422$ and 0.222). Only FreeOrbit4D (Cao et al., 2026a) does well on both axes (RotErr = 1.65° , $R\text{-cMaskIoU} = 0.905$).




Figure 11 Dancer case 3. Gen3C (Ren et al., 2025) wins both camera-pose columns (RotErr = 1.98° , TransErr = 0.237) yet renders the dancer as fragments with smeared limbs ($R = 0.311$, $R\text{-cMaskIoU} = 0.074$). TrajectoryCrafter (Yu et al., 2025) follows the path less closely (RotErr = 6.47°) and stretches her body ($R = 0.578$). ReCamMaster (Bai et al., 2025a) keeps the dancer intact ($R = 1.000$) but barely orbits (RotErr = 24.4°), so she appears away from her target-view location ($R\text{-cMaskIoU} = 0.235$). FreeOrbit4D (Cao et al., 2026a) follows the arc (RotErr = 4.82°) and is the only method that keeps the dancer intact and correctly placed ($R\text{-cMaskIoU} = 0.899$).




Figure 12 Deer case. ReCamMaster (Bai et al., 2025a) wins several VBench columns, but its camera diverges 32.9° from the arc and pushes the deer to the frame border, where only part of it remains visible ($R\text{-cMaskIoU} = 0.072$). TrajectoryCrafter (Yu et al., 2025) (RotErr = 3.05°) and Gen3C (Ren et al., 2025) (RotErr = 10.17°) follow the path more closely but distort the deer’s body and legs ($R = 0.733$ and 0.689). FreeOrbit4D (Cao et al., 2026a) follows the arc (RotErr = 4.20°) and keeps the deer intact and correctly placed ($R\text{-cMaskIoU} = 0.886$).




Figure 13 Robot case 1. ReCamMaster (Bai et al., 2025a) wins TransErr (0.327), but its camera barely orbits (RotErr = 28.4°) and the robot drifts toward the frame border, away from its target-view location (R -cMaskIoU = 0.281). Gen3C (Ren et al., 2025) tracks the path closely (RotErr = 2.41°) and TrajectoryCrafter (Yu et al., 2025) less so (RotErr = 10.9°), but both distort the robot’s head and torso, so recognition fails on almost every frame (R = 0.067 and 0.089). FreeOrbit4D (Cao et al., 2026a) (RotErr = 8.78°) keeps the robot intact from the new viewpoint and places it best (R -cMaskIoU = 0.842).




Figure 14 Robot case 2. TrajectoryCrafter (Yu et al., 2025), Gen3C (Ren et al., 2025), and FreeOrbit4D (Cao et al., 2026a) all reach low rotation error (under 2°), yet only FreeOrbit4D keeps the robot intact and correctly placed (R -cMaskIoU = 0.930). TrajectoryCrafter fails recognition on most frames (R = 0.244); Gen3C keeps a plausible robot on about 60% of frames (R = 0.622) but misplaces it (cMaskIoU = 0.127). ReCamMaster (Bai et al., 2025a) keeps the robot recognizable (R = 0.978) but under-rotates (RotErr = 17.6°) and mislocates it (R -cMaskIoU = 0.252).