

Reconstructing Satellites in 3D from Amateur Telescope Images

Zhiming Chang, Boyang Liu, Yifei Xia, Youming Guo, Boxin Shi, and He Sun

Abstract—Monitoring space objects is crucial for space situational awareness, yet reconstructing 3D satellite models from ground-based telescope images is super challenging due to atmospheric turbulence, long observation distances, limited viewpoints, and low signal-to-noise ratios. In this paper, we propose a novel computational imaging framework that overcomes these obstacles by integrating a hybrid image pre-processing pipeline with a joint pose estimation and 3D reconstruction module based on controlled Gaussian Splatting (GS) and Branch-and-Bound (BnB) search. We validate our approach on both synthetic satellite datasets and on-sky observations of China’s Tiangong Space Station and the International Space Station, achieving robust 3D reconstructions of low-Earth orbit satellites from ground-based data. Quantitative evaluations using SSIM, PSNR, LPIPS, and Chamfer Distance demonstrate that our method outperforms state-of-the-art NeRF-based approaches, and ablation studies confirm the critical role of each component. Our framework enables high-fidelity 3D satellite monitoring from Earth, offering a cost-effective alternative for space situational awareness. Project page: <https://ai4scientificimaging.org/3DSatellites>.

Index Terms—Computational Imaging, Astrophotography, 3D Reconstruction, Gaussian Splatting, Imaging through Turbulence

1 INTRODUCTION

THE monitoring of man-made objects in space is an essential endeavor in space domain awareness (SDA), particularly critical for ensuring the safe operation and on-orbit servicing of civilian satellites in the era of burgeoning large-scale satellite constellations, such as SpaceX’s Starlink [1]. Within the broad spectrum of satellite monitoring activities, including detection [2], categorization [3], and tracking [4], 3D reconstruction stands out as it uniquely provides unparalleled insights into the real-time operational status of satellites.

To date, 3D reconstruction has only been investigated in the space-based space surveillance missions, i.e. leveraging the videos captured by a satellite to retrieve the 3D information of its neighboring space objects. 3D reconstruction from ground-based observations, despite its advantages in terms of cost-effectiveness and versatility, remains under-explored because of a series of unique challenges compared with its space-based counterparts, including but not limited to: 1) atmospheric turbulence significantly distorts images captured by ground telescopes; 2) the long observation distance complicates the accurate estimation of a satellite’s pose; and 3) the rapid motion of satellites necessitates short exposure times, resulting in low signal-to-noise (SNR) ratios




Fig. 1: Problem definition. Our objective is to reconstruct a 3D model of a satellite in low-Earth orbit using limited, blurry, and noisy images captured by an amateur ground-based telescope, which are significantly affected by atmospheric turbulence and light pollution.

and constrained observational perspectives. As a result, classical 3D reconstruction approaches, no matter geometry or learning-based, fail in this challenging scenario.

In this paper, we propose an innovative computational imaging approach for reconstructing satellites in 3D using pure ground-based telescope images. Building on recent advances in imaging through turbulence and differentiable

- Zhiming Chang and He Sun are with the College of Future Technology and the National Biomedical Imaging Center, Peking University. Zhiming Chang is also with the College of Engineering, Peking University, Beijing, China (email: changzhiming@stu.pku.edu.cn, hesun@pku.edu.cn)
- Yifei Xia, Boxin Shi are with the State Key Laboratory of Multimedia Information Processing and National Engineering Research Center of Visual Technology, School of Computer Science, Peking University, Beijing, China (email: {yfxia, shiboxin}@pku.edu.cn)
- Boyang Liu is with the Beijing Jingdun Technology Co., Ltd. (email: pkulby@foxmail.com)
- Youming Guo is with the National Laboratory on Adaptive Optics and the Institute of Optics and Electronics, Chinese Academy of Sciences, Chengdu, China (email: guoyouming@ioe.ac.cn)
- Corresponding author: He Sun

volumetric rendering, we designed a cutting-edge 3D astrophotography pipeline that integrates image restoration foundation models for pre-processing and an innovative 3D Gaussian Splatting approach for joint satellite pose and shape reconstruction. This novel pipeline for the first time enables imaging 3D space objects from limited, distorted, and noisy ground telescope observations. We validate our method using both simulated and experimental datasets. Notably, we achieve the successful reconstruction of 3D models of the International Space Station (ISS) and China’s Space Station (CSS, or Tiangong Station), satellites in low Earth orbit, using solely images from an amateur 35cm-aperture telescope.

2 RELATED WORK

2.1 Spaceborne Satellite Monitoring

In space missions involving surveillance and servicing, accurate 3D reconstruction and pose estimation are vital for tasks like satellite maintenance, debris removal [5], and autonomous navigation. Due to limitations in size and power for onboard instruments, these tasks often rely on monocular camera observations from companion satellites. However, traditional computer vision algorithms like SIFT and SURF [6], [7] often struggle with the unique characteristics of space images, which typically exhibit high noise levels, simple textures, and symmetric features.

Recent advances have leveraged deep learning to overcome these challenges. Synthetic datasets like SPEED [8], SPEED+ [9], and SHIRT [10] have enabled the training of convolutional neural network (CNN)-based approaches that, when combined with predefined 3D object models (e.g., point clouds [11] or wireframes [12], [13]), yield robust pose estimation. Furthermore, emerging research is beginning to integrate pose estimation with simultaneous 3D reconstruction to monitor unknown objects—such as space debris—from limited 2D images [14], [15], [16]. Nevertheless, these models have primarily been tailored for the high-quality observations obtained by spaceborne telescopes. No established methods exist for reconstructing 3D satellite models purely from corrupted ground-based observations. Addressing this gap remains a critical challenge for future work in space situational awareness.

2.2 Imaging through Turbulence

Ground-based observations of space objects encounter significant image quality degradation due to atmospheric turbulence. Random fluctuations in the atmosphere’s refractive index introduce distortions and blurring in captured images. Adaptive optics (AO) [17] mitigates turbulence effects by actively correcting wavefront aberrations in real time using deformable mirrors [18], [19], but these systems are typically expensive and complex—limiting their use to large observatories.

An attractive alternative is computational correction [20], [21], [22]. For example, lucky imaging [23] captures high-frame-rate video and selects brief moments when turbulence is minimal—statistically, some frames will be less affected by distortions—and then combines these “lucky” frames to produce a higher-quality image. Recent progress

in turbulence simulation [24], [25], [26] has also facilitated the development of deep learning methods for turbulence mitigation. These methods train neural networks on pairs of degraded and ground-truth images to predict and correct turbulence effects [27], [28], [29], [30], often surpassing classical techniques in reconstruction quality and robustness.

2.3 3D Reconstruction and View Synthesis

3D reconstruction is a foundational technique in computer vision and graphics that enables the modeling and visualization of complex objects. Classical approaches typically leverage geometric cues from multiple images and rely on explicit representations. Techniques such as structure-from-motion (SfM) and multi-view stereo (MVS) build explicit models [31]—often in the form of point clouds or meshes—by leveraging epipolar geometry and bundle adjustment [32], [33]. While effective in many scenarios, these techniques can struggle with texture-less surfaces, occlusions, sparse viewpoints, and complex lighting conditions, challenges often encountered in 3D satellite reconstruction.

Neural Radiance Fields (NeRF) [34], [35] have revolutionized 3D reconstruction and view synthesis by modeling a scene’s volumetric density and color using deep networks. This implicit representation enables photorealistic renderings that capture intricate lighting and transparency from sparse viewpoints [36], [37]. Variants such as NeRF— and BARF further incorporate motion information inspired by SfM, allowing simultaneous 3D reconstruction and pose estimation from unposed images. However, despite their ability to capture fine geometric and appearance details, NeRF-based methods typically require significant training time and offer limited ease for scene editing or post-processing.

More recent advancements have introduced 3D Gaussian splatting [38], a technique that represents a scene as a collection of anisotropic Gaussians rather than an implicit function. This approach yields an efficient, interpretable representation that supports real-time volumetric rendering. By “splatting” each Gaussian onto the image plane and blending overlapping contributions, the method achieves high-fidelity reconstructions at interactive frame rates. The 3D Gaussian splatting family not only delivers competitive reconstruction quality but also significantly reduces rendering time, making it particularly promising for applications that demand both accuracy and speed, as well as ease of editing [39], [40], [41]. Multiple pose-free 3D Gaussian splatting reconstruction methods, such as SelfSplat [42], NoPoSplat [43], and FreeSplat [44], have also been introduced recently. However, these techniques typically require images with rich texture information, or sometimes rely on pre-trained depth estimators (effective only for specific scenes) to provide additional regularization.

2.4 Equatorial Mount and Meridian Flip

In astronomical imaging, telescopes are commonly mounted on equatorial mounts [45] to compensate for the Earth’s rotation and the apparent motion of astronomical objects. These mounts have two orthogonal axes to enable smooth tracking: the right ascension (RA) axis, aligned with Earth’s rotational axis and pointing toward the celestial north pole, and the declination (Dec) axis, which adjusts the telescope’s




Fig. 2: Illustration of equatorial mount tracking and meridian flip. (a) The telescope tracks a target across the celestial sphere. (b) The target reaches the local meridian. (c) Without adjustment, the telescope risks mechanical collision with the mount. (d) A meridian flip is performed, rotating the RA axis by 180° to safely continue tracking.

elevation (see Fig. 2(a)). During tracking, the RA axis rotates at either the sidereal rate or the satellite’s orbital rate to keep the target centered, while the Dec axis typically requires only minor adjustments.

However, as the target crosses the local meridian—a vertical plane passing through the zenith—the RA axis must continue rotating (Fig. 2(b)). This motion can bring the telescope dangerously close to the mount or tripod, risking mechanical collision (Fig. 2(c)). To prevent this, most equatorial mounts execute a meridian flip [46], [47], rotating the RA axis by 180° to reposition the telescope on the opposite side (Fig. 2(d)). This operation introduces a brief interruption and a mirrored camera orientation, typically causing a pause of around 30 seconds.

While such delays are negligible for slowly drifting targets like stars or nebulae, they become problematic for fast-moving satellites, which remain visible for only 2–3 minutes per pass. Compounding this issue, the meridian region often offers the best imaging conditions due to minimal atmospheric distortion. To address these challenges in ground-based satellite tracking, we should reconfigure the mount alignment to avoid meridian flips—an approach we detail in Sec. 4.2.1, referred to as the “Fake Polar Axis” method.

3 METHOD

Our method addresses the challenging problem of 3D reconstruction of space objects from unposed, monocular images captured by ground-based telescopes. In contrast to conventional 3D reconstruction tasks, our approach must contend with several unique challenges: atmospheric turbulence, extremely long observation distance, limited viewpoints, and

sparse, low-quality features. As shown in Fig. 3, our method overcomes these issues through two key innovations: 1) a hybrid image pre-processing pipeline combining classical and deep learning techniques, and 2) a joint pose-refinement and 3D reconstruction framework using Gaussian Splatting (GS) with controlled, iterative optimization.

3.1 Hybrid Telescope Image Pre-processing Pipeline

Ground-based telescope images are often compromised by atmospheric turbulence, noise, and light pollution, all of which distort and blur critical features. To address these issues, we synergize classical astronomical techniques with modern foundation models in a three-stage pipeline, as detailed in Fig. 3 (a).

3.1.1 Turbulence Mitigation via Lucky Imaging

We begin by applying lucky imaging [48] to high-frame-rate, short-exposure image sequences, statistically selecting the top 12% of frames based on sharpness and contrast metrics. These selected frames then undergo centroid-based alignment to correct for bulk satellite motion. Each frame is subsequently divided into multiple patches, and intensity-normalized cross-correlation is applied per patch to correct residual sub-pixel jitter induced by local turbulence. The aligned frames $\{\tilde{I}_t\}_{t \in S}$ are stacked to suppress noise and geometric distortions. The stacked image is computed as:

$$I_{\text{stacked}}(x, y) = \frac{1}{|S|} \sum_{t \in S} \tilde{I}_t(x, y),$$

where S denotes the subset of selected sharp frames. This averaging reduces zero-mean random noise by a factor of $1/\sqrt{|S|}$ and improves the signal-to-noise ratio (SNR).

3.1.2 Multi-Scale Sharpening with Wavelet Decomposition

While stacking reduces noise, it may also blur high-frequency details and suppress low-contrast structures due to short exposure. To restore such details, we apply multi-scale sharpening based on Gaussian wavelet decomposition, which separates the image into different frequency bands and selectively enhances the high-frequency components. Let $W_j(I)$ denote the wavelet decomposition of image I at scale j ($j = 1, \dots, 6$, from fine, 256×256 , to coarse, 8×8). We apply a multiplicative gain factor α_j to each level:

$$\tilde{W}_j(I) = \alpha_j \cdot W_j(I),$$

with

$$(\alpha_1, \alpha_2, \alpha_3, \alpha_4, \alpha_5, \alpha_6) = (1.1, 1.1, 1.1, 1.1, 1.1, 5.0).$$

The substantially larger gain at level 6 compensates for global contrast loss and enhances the visibility of extended, low-contrast structures (e.g., solar wings).

3.1.3 Artifact Suppression via Video Restoration Transformer

Classical sharpening techniques sometimes introduce artifacts or over-enhance noise. To mitigate these issues, we employ the Video Restoration Transformer (VRT) [49], a foundation model pre-trained on large-scale natural images, to further refine the pre-processed images. Leveraging self-attention mechanisms and temporal consistency, VRT denoises the pre-processed images, preserving essential structural details for accurate 3D reconstruction.




Fig. 3: Overview of the 3D satellite reconstruction method. (a) **Image pre-processing:** A novel pipeline mitigates atmospheric turbulence and noise by combining classical signal processing with a Video Restoration Transformer (VRT). Stacking (lucky imaging) boosts the video’s SNR, wavelet sharpening enhances high-frequency features, and VRT removes non-Gaussian noise and complex artifacts, significantly improving raw telescope images. (b) **Joint pose Estimation and 3D reconstruction:** Our method begins with a modified SfM for coarse pose initialization. In this approach, we employ an orthographic projection model—appropriate given the minimal perspective effects in long-distance observations—and sequentially register adjacent frames to regularize motion. Using the resulting sparse points, we initialize a 3D Gaussian point cloud and iteratively refine the Gaussian parameters and the camera poses. Pose estimation is optimized through branch-and-bound (BnB) search to minimize projection errors. Notably, we regulate the growth rate of the Gaussians to ensure they first capture the satellite’s overall geometry (low-frequency information) before learning finer details (high-frequency information). This strategy is particularly effective when initial pose estimates are inaccurate. After several iterations, we halt pose updates while continuing to train the Gaussians. (c) **Post-processing:** Toward the end of training, KNN filtering removes isolated or noisy points from the Gaussian point cloud. (d) **Training schedule:** An overview of the complete training process is presented.

3.2 Joint Pose Estimation and 3D Gaussian Splatting Reconstruction

Once high-quality images are obtained through pre-processing, we tackle the dual problem of reconstructing the object’s 3D shape and estimating its pose. This step is particularly challenging given the long observation distance, limited viewpoint, and scarce features present in our satellite images. Our solution comprises three important stages:

3.2.1 Initialization Using Modified SfM

We begin by establishing an initial estimate of the object’s pose and rough 3D structure using a modified SfM approach. In astrophotography, the satellite moves rapidly in orbit while the telescope remains fixed on Earth. To better leverage existing computer vision tools, we reframe the problem by treating the satellite as stationary and the telescope as if it “moves” around it.

Our approach first leverages pre-trained deep learning models—SuperPoint for feature extraction and SuperGlue

for feature matching—to detect and pair keypoints robustly despite distortions, blur, and sparse textures inherent in telescope images. To mitigate the challenges posed by the long observation distance, we adopt an orthographic projection model [50] (instead of a perspective model) in the initial step, which is more robust when perspective effects are minimal. Applying this model to the first three frames yields an initial point cloud along with corresponding camera poses.

Next, we revert to the full perspective camera model and sequentially integrate additional frames, as shown in the “Modified SfM” module in Fig. 3 (b). For each new image, we first apply the Random Sample Consensus (RANSAC) algorithm to filter out incorrect feature correspondences between adjacent images, ensuring more accurate feature matches and better associations with the 3D point cloud. Based on these refined correspondences, we then estimate the initial camera pose using the Perspective-n-Point (PnP) algorithm. Additionally, to prevent error accumulation, we perform a global bundle adjustment after each new frame

is added, refining all camera poses simultaneously. Additionally, every five frames we perform triangulation and integrate new features (e.g., from the 5th and 6th frames, then the 10th and 11th, etc.) into the point cloud. This iterative process continues until all images have been processed, resulting in a complete initial 3D model and the satellite’s pose estimates.

3.2.2 Controlled Gaussian Splatting with BnB Pose Refinement

The core of our reconstruction pipeline is a novel integration of 3D Gaussian splatting with iterative pose refinement [51]. We begin by initializing Gaussian points from the SfM-derived point cloud, then iteratively optimize the GS parameters while holding camera poses fixed, and vice versa. This alternating optimization enables simultaneous recovery of the satellite’s 3D shape and accurate camera poses despite challenges such as long observation distance, limited viewpoints, and sparse features.

During the early GS training phase, we fix the number of Gaussians and optimize only their parameters—position, covariance, opacity, and spherical harmonic coefficients (up to third order). By limiting the number of Gaussians, we force each one to cover a larger region, effectively capturing the satellite’s coarse geometry while preventing overfitting to high-frequency artifacts when pose estimates are imprecise. After a set number of epochs, we automatically clone and split the Gaussians to increase their count, then resume training on the expanded set with the count held fixed. This controlled growth strategy enables the model to first learn coarse geometry and gradually capture finer details.

In the first 10,000 epochs of training, we refine the camera poses through a Branch-and-Bound (BnB) search performed every 1,000 epochs. Candidate poses are generated by applying controlled perturbations to the current estimate. Specifically, for rotation, we sample a random rotation angle, $\Delta\theta$, and a random axis, $\vec{n}$, then compute the rotation perturbation using Rodrigues’ formula:

$$R_{\text{noise}} = I + \sin(\Delta\theta)[\vec{n}]_{\times} + (1 - \cos(\Delta\theta))[\vec{n}]_{\times}^2, \quad (1)$$

where $[\vec{n}]_{\times}$ is the skew-symmetric matrix of $\vec{n}$. Translation perturbations are drawn from a zero-mean Gaussian distribution:

$$T_{\text{noise}} \sim \mathcal{N}(0, \sigma_T^2), \quad (2)$$

Each candidate pose is then given by:

$$R_i = R \cdot R_{\text{noise}}, \quad T_i = T + T_{\text{noise}}, \quad (3)$$

At each BnB stage, we sample 200 pose candidates (example camera poses are shown in red in Fig. 3(b)). We evaluate these candidates using a photometric consistency loss and select the pose $(R_{\text{best}}, T_{\text{best}})$ that minimizes this loss. After each stage, we shrink the translation range to one-fourth and the rotation range to one-half of the previous interval, centering both on the current camera pose estimate. This refinement schedule ensures stable convergence within ten rounds. The best pose is then passed to the next GS training phase. By alternating between Gaussian splatting with adaptive point growth and BnB-based pose refinement, our method jointly converges to accurate 3D structure and camera poses.

Our implementation builds on the original 3D Gaussian Splatting framework [38], using the same hyperparameter set across all experiments—synthetic data, CSS, and ISS—without scene-specific tuning, and keeping all other settings unchanged except for the modifications described above.

3.2.3 Explicit Post-processing of the Learned Point Cloud

Because the reconstructed 3D Gaussians may include stray points from background noise, we further apply a statistical filtering step to improve model accuracy. As illustrated in Fig. 3(c), we examine each Gaussian’s spatial location (ignoring its appearance and shape attributes), compute the mean distance to its k nearest neighbors, and remove any point whose mean distance exceeds the global mean plus one standard deviation ($\mu + \sigma$). After filtering, we resume Gaussian splatting training on the refined 3D Gaussians for an additional 500 epochs to repair discontinuities and further improve reconstruction quality.

Fig. 3 (d) summarizes the complete training schedule of our 3D reconstruction pipeline—including initialization, GS training, pose refinement, and explicit filtering. Our integrated approach, enabled by careful algorithmic design, effectively overcomes multiple fundamental challenges of reconstructing space objects from degraded ground-based telescope images.

4 EXPERIMENTS

In this section, we evaluate our method using both simulated and real on-sky data. We first benchmark our pre-processing method against multiple image deblurring baselines, and then compare our 3D reconstruction performance with that of cutting-edge NeRF-based baselines [53], [54]. All pose-free 3D GS baselines are excluded from our discussion, as experiments indicate that none are applicable due to the sparse textures inherent in astrophotography. We quantitatively assess reconstruction accuracy using standard imaging and 3D metrics—SSIM, PSNR, LPIPS, and Chamfer Distance (CD)—and evaluate camera pose accuracy via the Absolute Trajectory Error (ATE). Our method demonstrates superior quality and robustness on synthetic ground-based telescope images as well as on real observations of China’s Space Station (CSS) and the International Space Station (ISS).

4.1 Simulation Results

4.1.1 Synthetic Ground Telescope Satellite Images

We generated three distinct 3D satellite models using Blender, each approximately 60 meters in size, and simulated ground telescope images for each. Placed in a low-Earth orbit at 650 km altitude, each satellite is visible from a stationary ground telescope in a 13° segment, which corresponds to an apparent 116° relative rotation between the satellite and the telescope (see Fig. 7 (a)). Our virtual telescope—with a 0.35-meter aperture and a 3.2-meter focal length—rendered 700 high-quality images per model. To simulate realistic observing conditions, we introduced a full degradation model incorporating both atmospheric turbulence and light pollution. Specifically, each clean image was degraded as:

$$I_{\text{noisy}} = (I_{\text{clean}} * \text{PSF}_{\text{turb}}) + B + n_{\text{ph}} + n_{\text{read}}.$$

TABLE 1: Quantitative comparison of image pre-processing quality between our method and various baselines on synthetic data. The best results in each column are in **bold**, and second-best are underlined. Our method achieves the highest PSNR and SSIM metrics. For a detailed explanation of each method, please refer to the Fig. 4’s caption.

Method	Simulation 1			Simulation 2			Simulation 3		
	PSNR [↑]	SSIM [↑]	LPIPS [↓]	PSNR [↑]	SSIM [↑]	LPIPS [↓]	PSNR [↑]	SSIM [↑]	LPIPS [↓]
Dirty Images	23.28	0.07	0.41	23.02	0.07	0.36	23.16	0.07	0.36
Restormer	24.83	0.18	0.09	24.76	0.18	0.10	25.06	0.18	0.10
Stacking	27.00	0.41	0.15	27.15	0.41	0.17	27.75	0.41	0.16
Stacking + SelfDeblur	23.89	0.46	<u>0.10</u>	24.39	0.45	0.10	24.39	0.45	0.10
Stacking + Restormer	26.78	0.39	<u>0.14</u>	26.93	0.38	0.16	27.49	0.39	0.15
Stacking + Sharpening	<u>27.60</u>	0.53	<u>0.10</u>	<u>27.92</u>	<u>0.53</u>	<u>0.11</u>	28.3	0.52	0.11
Ours (Stacking + Sharpening + VRT)	27.74	0.53	0.11	28.02	0.55	0.13	28.47	0.54	0.11




Fig. 4: Comparison of pre-processed images obtained using our method and various baselines on synthetic data. “Dirty images” denote the raw data; “Restormer” is a pre-trained deblurring network, selfDeblur [52] is a self-supervised blind deconvolution method, and “VRT” is a pre-trained multi-task video restoration network used for denoising. “Stacking” refers to the frame alignment and averaging process from lucky imaging, while “Sharpening” applies wavelet-based enhancement. Our method yields the most faithful reconstruction.

TABLE 2: Quantitative comparison of pose estimation accuracy (Absolute Trajectory Error, ATE) for various methods on synthetic data. “N/A” indicates failure to converge. The top half presents results for pose initialization, while the bottom half shows outcomes after joint 3D reconstruction and pose refinement. Notably, our pose initialization method is the only one that produces reasonable initial estimates, and our overall approach achieves the highest accuracy after joint optimization.

Method	Absolute Trajectory Error [↓]
Pose Initialization	
Colmap / Hloc / Detector-Free SfM	N/A
Dust3r	138.81
RayDiffusion	138.76
Ours	2.61
After pose Refinement	
After BnB	2.43
NeRF—	2.57
BARF (Constrain Sampling)	2.54



Fig. 5: Comparison of estimated poses generated by our method and various baselines. Methods that fail to converge, such as Colmap, Hloc, and Detector-free SfM, are omitted. Our approach achieves the highest pose initialization accuracy.

Here, PSF_{turb} is the optical point spread function degraded by atmospheric turbulence, simulated using the Phase-to-Space Transform (P2S) method [25]. P2S approximates Kolmogorov turbulence by sampling Zernike polynomials and PCA bases learned from real turbulence data. The

turbulence strength is controlled via the Fried parameter r_0 , which we sampled uniformly from the range [0.07m, 0.35m], equivalent to restricting the turbulence integration height to roughly the lowest 10 km of the atmosphere. Light pollution was modeled as a combination of additive background bias $B \sim \mathcal{U}(0.05 I_{\text{clean}}^{\max}, 0.07 I_{\text{clean}}^{\max})$, signal-dependent photon noise $n_{\text{ph}} \sim \mathcal{N}(0, \sigma_1^2(I_{\text{turb}} + B)^2)$, and signal-independent

TABLE 3: Quantitative comparison of 3D reconstruction accuracy between our method and NeRF-based approaches. Results are evaluated using three simulation experiments and two on-sky observations, China’s Space Station (CSS) and the International Space Station (ISS). PSNR, SSIM, and LPIPS metrics are computed by comparing novel views with Blender renderings for simulated data and with pre-processed reserved validation views for on-sky data. Best results in each column are shown in **bold**, and second-best values are underlined.

Method	Simulation 1				Simulation 2				Simulation 3				CSS			ISS		
	PSNR [↑]	SSIM [↑]	LPIPS [↓]	CD [↓]	PSNR [↑]	SSIM [↑]	LPIPS [↓]	CD [↓]	PSNR [↑]	SSIM [↑]	LPIPS [↓]	CD [↓]	PSNR [↑]	SSIM [↑]	LPIPS [↓]	PSNR [↑]	SSIM [↑]	LPIPS [↓]
Ours	25.32	0.90	0.16	0.07	27.53	0.91	0.14	0.08	28.10	0.91	0.11	0.07	34.68	0.96	0.03	36.87	0.98	0.02
NeRF—	16.09	0.53	0.37	0.34	18.06	<u>0.88</u>	0.29	<u>0.34</u>	16.24	<u>0.76</u>	0.36	0.42	15.09	0.67	0.34	29.05	<u>0.92</u>	0.10
BARF	17.64	<u>0.55</u>	<u>0.40</u>	0.38	17.28	0.47	0.35	0.35	<u>22.82</u>	0.53	0.18	0.43	23.52	<u>0.77</u>	0.15	30.85	0.88	0.12
BARF (Con.)	<u>24.20</u>	0.54	0.16	<u>0.34</u>	28.67	0.62	0.11	0.34	22.07	0.56	<u>0.14</u>	<u>0.37</u>	<u>31.04</u>	<u>0.87</u>	<u>0.06</u>	<u>31.53</u>	<u>0.92</u>	<u>0.06</u>




Fig. 6: Box plots comparing PSNR, SSIM, and LPIPS results for our method versus NeRF-based methods on Simulation 1. Our approach consistently delivers stable performance across both training and validation sets. In contrast, even with the constrained sampling strategy illustrated in Fig. 8, NeRF-based methods exhibit significant discrepancies between the training and validation sets, indicating lower robustness.




Fig. 7: Satellite tracking geometry and the Fake Polar method: (a) A 13° ground track corresponds to a 116° apparent angular sweep from the telescope. (b) Tilting the mount’s polar axis into the orbital plane avoids meridian flip and shifts tracking motion primarily to the Dec axis.

readout noise $n_{\text{read}} \sim \mathcal{N}(0, \sigma_2^2)$, where $I_{\text{turb}} = I_{\text{clean}} * \text{PSF}_{\text{turb}}$. This procedure resulted in a dataset of 14,000 degraded frames per model, each at a resolution of 512×512 pixels. All synthetic data and code are available on our project page (see Abstract).

4.1.2 Raw Telescope Image Pre-processing

We first validate our pre-processing pipeline on the simulated data. Following our pipeline—lucky imaging and stacking (every 100 images), wavelet-based multi-scale sharpening, and VRT denoising—we reduce the 14,000 distorted frames to 140 high-quality images per satellite model. We compared our method against other cutting-edge deblurring algorithms. Table 1 summarizes the evaluation metrics (SSIM, PSNR, LPIPS) across various preprocessing methods. The results indicate that our proposed pipeline (Stacking + Sharpening + VRT) consistently surpasses pure deep learning approaches (e.g., Restormer [55]), traditional




Fig. 8: Constrained sampling strategy for BARF. Dark blue triangles represent cameras, yellow rays denote light paths, light blue dots indicate the original uniform sampling region, and gray dots mark the constrained sampling region suggested by the initial camera-to-target distance estimate.

techniques (e.g., Stacking + Sharpening), self-supervised methods, and other hybrid strategies, specifically in PSNR and SSIM metrics.

4.1.3 Robust Pose Initialization

From the 140 processed images, we select every tenth frame (e.g., 1st, 10th, 20th) for the training set and reserve the intermediate frames for validating the reconstruction. We first compare the pose estimation accuracy of our method against existing techniques, including feature-based SfM methods—such as Colmap [31], hierarchical localization [56], and detector-free SfM [57]—as well as several feed-forward deep learning approaches, including DUS3R [58] and RayDiffusion [59]. Fig. 5 and Table 2 compare the estimated poses to the ground truth (GT) on the training set. Feature-based SfM methods accumulate noise, which leads

(a) 3D Point Cloud Reconstruction Results

Simulation 1

Simulation 2



(b) Novel View Synthesis Results




Fig. 9: Comparison of 3D reconstruction results on synthetic data using our method versus NeRF-based baselines. (a) Reconstructed 3D point clouds from our method compared with the ground truth. (b) Novel view synthesis results from our method, NeRF-based baselines, and the ground truth. Results from two simulation datasets (1 & 2) are presented. Our approach accurately recovers the satellite’s 3D geometry, as well as produces clearer structures and higher SNR in novel views. In contrast, NeRF-based methods suffer from overfitting, resulting in pronounced artifacts and noise that diminish the quality of the reconstructed views.

to inaccurate keypoint matching and causes bundle adjustment (BA) to fail to converge; consequently, no reasonable pose estimation can be reported. Similarly, deep learning methods struggle with extreme focal lengths, structural symmetry, and poor textures, resulting in incorrect depth estimation and subsequent pose errors. Other approaches, such as PoseDiffusion [60], MV-DUST3R+ [61], and Colmap-free 3DGS [62], also exhibit substantial pose estimation errors. In contrast, our pose initialization method achieves an ATE of only 2.61, with the estimated pose distribution closely aligning with the ground truth. As explained in Sec.3.2.1, our superior results stem from three key factors: (1) a robust orthographic projection model; (2) the use of RANSAC algorithms to eliminate erroneous matches; and (3) a strategy that leverages only the matching relationships between adjacent frames to relax the convergence conditions. These design choices enable our method to overcome the limitations of traditional SfM and deep learning approaches, significantly enhancing reliability for pose estimation of texture-poor space objects.

4.1.4 Joint 3D Reconstruction and Pose Refinement

We then compare our pipeline with two NeRF-based baseline methods, NeRF-- and BARF, that support joint 3D reconstruction and pose estimation. To ensure a fair comparison, we use our customized pose initialization for all baselines; otherwise, neither NeRF-- nor BARF exhibits generalization ability. In addition, traditional NeRF-based methods uniformly sample the space between the cameras and the object (illustrated by the light blue dots in Fig. 8). This approach fails in our sparse-view, long-distance observational scenario because of overfitting in the excessively large sampling space. Consequently, we explore a BARF


variant (named BARF (Con.)) that restricts the sampling region to the vicinity of the initial object distance predicted by our modified SfM (indicated by the gray dots in Fig. 8). This adaptive strategy significantly reduces overfitting and improves reconstruction accuracy.

Fig. 9 compares the reconstructed point clouds (appearance and shape attributes of 3D Gaussians neglected) and novel views from our method with those produced by baseline techniques. Although NeRF-based methods have been carefully tuned for this task, they still struggle to generalize to unseen viewpoints. In contrast, our GS-based approach employs controlled point growth to focus sampling within the object region, accurately capturing fine structural details. This advantage is confirmed by superior SSIM, PSNR, and LPIPS metrics (Table 3 and Fig. 6). In the meantime, our method delivers more accurate pose estimates after the joint optimization of poses and the 3D object (Table 2).


To further quantify the precision of our reconstructions, we measure the Chamfer distance (CD) between the reconstructed models and the ground-truth point clouds. As illustrated in Table 3, our method demonstrates significantly lower CD values, showcasing its ability to accurately capture 3D geometries. As visualized by the point clouds and the novel views in Fig. 9, our method accurately recovers the 3D structure of the satellite, while NeRF-based methods fail to extract any structural information.

Our method also significantly reduces training time. A 3D satellite model is reconstructed in approximately 4 minutes on a single NVIDIA RTX 4090 using our method (2 minutes for GS training and 2 minutes for BnB pose refinement), much faster than NeRF--’s 1.5 hours and BARF’s over 4 hours.

(a) Ours 3D Point Cloud Reconstruction Results
China's Space Station (CSS)



International Space Station (ISS)



(b) Novel View Synthesis Results




Fig. 10: Comparison of 3D reconstruction results on real observations using our method versus NeRF-based baselines. (a) Reconstructed 3D point clouds produced by our method from multiple viewpoints. (b) Novel view synthesis results from our method and NeRF-based baselines. Results from two on-sky observations, CSS and ISS, are presented. Our approach accurately recovers the satellite's 3D geometry, as well as produces clearer structures and higher SNR in novel views than NeRF-based methods, which use reserved pre-processed images as references.



Fig. 11: The Chinese Space Station (CSS) and the imaging instrument. On the left, the official schematic of the CSS, featuring four primary modules: 1 - Tianhe Core Module, 2 - Mengtian Laboratory Module, 3 - Wentian Laboratory Module, and 4 - Tianzhou-6 Cargo Spacecraft. In the lower right corner, the Celestron C14HD telescope paired with a QHY5III678M planetary camera.

4.2 On-Sky Results

We apply our algorithm to real observations of China's Tiangong Space Station (CSS) and the ISS. Since both low-Earth orbit satellites have well-documented geometric specifications, we are able to assess our method's accuracy by comparing the reconstructions with publicly available information.

4.2.1 Observations of CSS and ISS


The CSS was observed on September 15, 2023, from 20:48:30 to 20:51:17 UTC near Miyun Reservoir, Beijing, with a closest approach of 389.4 km and a peak elevation of 78.3°. The ISS was observed on August 21, 2024, from 20:20:06 to

TABLE 4: Metrology of the China's Space Station modules and their solar wings.

Category	Module	Measured	GT [63]
Length	Tianhe Core Module	16.11	16.6
	Mengtian Laboratory	17.30	17.88
	Wentian Laboratory	17.30	17.88
	Tianzhou-6	9.71	10.6
Diameter	Tianhe Core Module	5.65	4.2
	Mengtian Laboratory	5.50	4.2
	Wentian Laboratory	5.63	4.2
	Tianzhou-6	3.30	3.35
Solar wing length	Tianhe Core Module	25.44	×
	Mengtian Laboratory	55.62	55
	Wentian Laboratory	55.62	55
	Tianzhou-6	4.89	×
Solar wing angle	Tianhe Core Module	0.78	×
	Mengtian Laboratory	63.01	×
	Wentian Laboratory	68.71	×
	Tianzhou-6	3.22	×

20:22:09 UTC in Yanqing District, Beijing, with a closest approach of 442 km and a peak elevation of 86°. For both observations, we used a Celestron C14HD telescope (35cm aperture, f/11) with a QHY5III678M camera (shown in the lower right corner of Fig. 11), achieving a spatial sampling rate of 0.106"/pixel (0.2m/pixel at 400km). This sampling rate is smaller than the optical resolution: The diffraction limit of our 35 cm telescope is approximately 0.39", corresponding to 76cm resolution at 400km. Our observations employed the "Fake Polar Axis" method with a German equatorial mount (CGX-L), whose polar axis was tilted into the satellite's orbital plane rather than aligned with the celestial pole (see Fig. 7 (b)). This orientation eliminates the need for meridian flips and enables uninterrupted tracking during the satellite pass. As most of the satellite's apparent

(a) 3D Point Cloud Reconstruction Results



(b) Novel View Synthesis results




Fig. 12: Ablation studies of our 3D satellite reconstruction pipeline. Reconstructed point clouds (a) and novel views (b) of Simulation 2 data illustrate the effect of removing individual steps. (1) **Ours**: The complete pipeline produces accurate point cloud and novel views. (2) **Original GS**: Using the original GS algorithm—initialized with our modified SfM but lacking all subsequent steps—leads to overfitting and inaccurate pose estimation, resulting in chaotic point clouds and severe artifacts in novel views. (3) **No Preprocessing**: Omitting preprocessing yields noisy point clouds and overly-smoothed novel view renderings. (4) **No Regulation**: The absence of constraints on point cloud growth rate exacerbates overfitting, producing super noisy point clouds. (5) **No BnB**: Excluding BnB refinement introduces artifacts in both point clouds and novel view synthesis due to erroneous pose estimation. (6) **No Filtering**: Omitting point cloud filtering results in persistent background clutter and residual outliers. These studies underscore the critical role of each step in achieving stable, detailed 3D reconstructions and high-quality novel view synthesis.

TABLE 5: Ablation studies of our 3D satellite reconstruction pipeline on three simulation datasets. Quantitative metrics (PSNR, SSIM, LPIPS, CD) for novel view synthesis are reported. Best results in each column are highlighted in **bold**, and second-best are underlined.

Method	Simulation 1				Simulation 2				Simulation 3			
	PSNR [↑]	SSIM [↑]	LPIPS [↓]	CD [↓]	PSNR [↑]	SSIM [↑]	LPIPS [↓]	CD [↓]	PSNR [↑]	SSIM [↑]	LPIPS [↓]	CD [↓]
Ours	25.33	0.90	<u>0.16</u>	0.07	27.53	0.91	<u>0.14</u>	0.08	28.10	0.91	<u>0.11</u>	0.07
Original GS	20.96	0.35	0.18	0.45	23.81	0.36	0.17	0.15	22.12	0.34	0.20	0.17
No Preprocessing	20.87	<u>0.87</u>	0.22	0.38	25.76	0.37	0.21	0.15	22.24	<u>0.85</u>	0.22	0.13
No Regulation	22.98	<u>0.87</u>	0.14	<u>0.21</u>	26.23	0.87	0.13	0.16	27.72	0.65	0.10	0.12
No BnB	22.51	0.86	0.18	0.30	25.89	<u>0.90</u>	0.16	<u>0.14</u>	<u>28.01</u>	0.91	0.12	<u>0.09</u>
No Filtering	<u>24.98</u>	0.62	0.17	0.31	<u>26.89</u>	0.61	0.16	0.16	27.88	0.63	0.12	0.14

motion is projected onto the declination (Dec) axis, only minor compensation is required from the right ascension (RA) axis, reducing field rotation—especially near the center of the image sequence. We implemented this method using a custom tracking program. Images were captured at 5 ms exposures, averaging 88 fps for CSS and 55 fps for ISS, resulting in datasets of 13,800 images for CSS and 6,770 for ISS, respectively. The relative angular sweeps between the telescope and the satellites span approximately 99° for CSS and 82° for ISS.

4.2.2 3D Reconstruction and Metrology

We first apply the pre-processing pipeline to extract clean frames from the raw observations, selecting 138 clean frames for CSS and 116 for ISS. Following our simulation protocol,

we chose one image out of every ten for training—yielding 14 images for CSS and 12 for ISS—with the remainder reserved for validation. Our method is again benchmarked against three NeRF-based baselines, NeRF—, BARF and BARF’s variant (BARF with Constrained Sampling, abbreviated as BARF (Con.)).

The 3D reconstruction results for CSS and ISS (Fig. 10) closely mirror our simulation findings, while NeRF-based methods struggle to generalize to unseen viewpoints with limited training data, our approach achieves more accurate 3D reconstruction and generated more realistic new views. Quantitative metrics reported in Table 3 also demonstrates our method’s superior performance.

We further validate our reconstructions via detailed metrology. Specifically, we measure the lengths and diam-

eters of four CSS modules (Tianhe, Mengtian, Wentian, and Tianzhou-6) and determine the angles and dimensions of their solar arrays (see Fig. 11). Using our reconstructed 3D Gaussian splatting model, we densely sample 2D views to identify frames where each module appeared in orthogonal perspective at its maximum extent, then compute dimensions via pixel counts. The solar array angles are derived from their spatial relationships with the main modules. As summarized in Table 4, our metrology results align with the officially released CSS specifications.

4.3 Ablation Study

We finally conduct comprehensive ablation studies on our simulation data to evaluate the contribution of each component in our 3D reconstruction framework. Fig. 12 display reconstructed point clouds and representative novel view synthesis results for the naive Gaussian Splatting method and when individual steps are omitted.

Each ablation visibly degrades reconstruction performance, highlighting the importance of every step in our pipeline for accurate pose estimation, high-fidelity 3D reconstruction, and consistent novel view synthesis. The absence of any component results in errors like blurriness, ghosting, structural discontinuities, fringe artifacts, or noisy backgrounds, underscoring the necessity of our complete method.

Quantitative evaluation using PSNR, SSIM, LPIPS, and CD metrics (see Table 5) confirms that every step in our pipeline contributes to reconstruction quality. Notably, the modified SfM-based initialization, controlled Gaussian growth, and BnB-search-based pose refinement are the most critical innovations in our method.

5 CONCLUSION

In this paper, we present a novel 3D imaging framework for reconstructing satellites from ground-based telescope observations. Our approach integrates a hybrid telescope image pre-processing pipeline with a joint pose-refinement and 3D reconstruction framework that leverages controlled Gaussian Splatting (GS) and Branch-and-Bound (BnB) search. We validated our method on synthetic datasets and on-sky observations, achieving, for the first time, 3D imaging of China's Tiangong Space Station and the International Space Station using ground-based telescopes.

Due to hardware limitations, such as the telescope's aperture size, our current 3D imaging resolution remains restricted. Additionally, our current model, based on a simple photometric-consistency loss, does not fully capture strong specular reflections from metallic surfaces and solar panels, which may bias the fit under highly specular conditions. In future work, we plan to use larger, more stable telescopes to improve image resolution and signal-to-noise ratio, enabling the reconstruction of smaller satellites. We will also develop modular interfaces so that updated denoising and matching networks can be seamlessly integrated. Additionally, we plan to incorporate material properties into our 3D satellite models to better simulate sunlight reflection and occlusion effects. We will further extend our approach to capture specular reflections by using reflectance-aware Gaussian

splatting. Additionally, we will explore other applications of our method for 3D imaging at long observation distances and with low-quality images.

ACKNOWLEDGMENTS

This work was supported by National Natural Science Foundation of China (Grant No. 62371007, 62136001, 62088102, 12173041).

REFERENCES

- [1] SpaceX, "Starlink," <https://www.starlink.com>, accessed: [November, 2024].
- [2] J. Fletcher, I. McQuaid, P. Thomas, J. Sanders, and G. Martin, "Feature-based satellite detection using convolutional neural networks," in *Proceedings of the Advanced Maui Optical and Space Surveillance Technologies Conference*, vol. 11, 2019.
- [3] I. Sanad, Z. Vali, and D. G. Michelson, "Statistical classification of remote sensing satellite constellations," in *2020 IEEE Aerospace Conference*. IEEE, 2020, pp. 1–15.
- [4] G. Hawkins, D. Edwards, and J. McGeehan, "Tracking systems for satellite communications," in *IEE Proceedings F (Communications, Radar and Signal Processing)*, vol. 135, no. 5. IET, 1988, pp. 393–407.
- [5] C. P. Mark and S. Kamath, "Review of active space debris removal methods," *Space policy*, vol. 47, pp. 194–206, 2019.
- [6] D. G. Lowe, "Distinctive image features from scale-invariant keypoints," *International journal of computer vision*, vol. 60, pp. 91–110, 2004.
- [7] S. Sharma *et al.*, "Comparative assessment of techniques for initial pose estimation using monocular vision," *Acta Astronautica*, vol. 123, pp. 435–445, 2016.
- [8] S. Sharma and S. D'Amico, "Neural network-based pose estimation for noncooperative spacecraft rendezvous," *IEEE Transactions on Aerospace and Electronic Systems*, vol. 56, no. 6, pp. 4638–4658, 2020.
- [9] T. H. Park, M. Märtens, G. Lecuyer, D. Izzo, and S. D'Amico, "Speed+: Next-generation dataset for spacecraft pose estimation across domain gap," in *2022 IEEE Aerospace Conference (AERO)*. IEEE, 2022, pp. 1–15.
- [10] T. H. Park and S. D'Amico, "Adaptive neural-network-based unscented kalman filter for robust pose tracking of noncooperative spacecraft," *Journal of Guidance, Control, and Dynamics*, vol. 46, no. 9, pp. 1671–1688, 2023.
- [11] S. Qiao, H. Zhang, G. Meng, M. An, F. Xie, and Z. Jiang, "Deep-learning-based satellite relative pose estimation using monocular optical images and 3d structural information," *Aerospace*, vol. 9, no. 12, p. 768, 2022.
- [12] S. Sharma, J. Ventura, and S. D'Amico, "Robust model-based monocular pose initialization for noncooperative spacecraft rendezvous," *Journal of Spacecraft and Rockets*, vol. 55, no. 6, pp. 1414–1429, 2018.
- [13] S. D'Amico, M. Benn, and J. L. Jørgensen, "Pose estimation of an uncooperative spacecraft from actual space imagery," *International Journal of Space Science and Engineering* 5, vol. 2, no. 2, pp. 171–189, 2014.
- [14] H. Zhang, Q. Wei, and Z. Jiang, "3d reconstruction of space objects from multi-views by a visible sensor," *Sensors*, vol. 17, no. 7, p. 1689, 2017.
- [15] K. Dennison and S. D'Amico, "Vision-based 3d reconstruction for navigation and characterization of unknown, space-borne targets," *Austin, TX, Jan*, 2023.
- [16] T. H. Park and S. D'Amico, "Rapid abstraction of spacecraft 3d structure from single 2d image," in *AIAA SCITECH 2024 Forum*, 2024, p. 2768.
- [17] F. Roddier, "Adaptive optics in astronomy," 1999.
- [18] H. Sun, N. J. Kasdin, and R. Vanderbei, "Identification and adaptive control of a high-contrast focal plane wavefront correction system," *Journal of Astronomical Telescopes, Instruments, and Systems*, vol. 4, no. 4, pp. 049006–049006, 2018.
- [19] H. Sun, A. Goun, S. Redmond, M. Galvin, T. Groff, M. Rizzo, and N. Jeremy Kasdin, "High-contrast integral field spectrograph (hcifs): multi-spectral wavefront control and reduced-dimensional system identification," *Optics Express*, vol. 28, no. 15, pp. 22412–22423, 2020.

- [20] N. Anantrasirichai, A. Achim, and D. Bull, "Atmospheric turbulence mitigation for sequences with moving objects using recursive image fusion," in *IEEE international conference on image processing*, 2018, pp. 2895–2899.
- [21] N. Anantrasirichai, A. Achim, N. G. Kingsbury, and D. R. Bull, "Atmospheric turbulence mitigation using complex wavelet-based fusion," *IEEE Transactions on Image Processing*, vol. 22, no. 6, pp. 2398–2408, 2013.
- [22] Z. Mao, N. Chimitt, and S. H. Chan, "Image reconstruction of static and dynamic scenes through anisoplanatic turbulence," *IEEE Transactions on Computational Imaging*, vol. 6, pp. 1415–1428, 2020.
- [23] N. M. Law, C. D. Mackay, and J. E. Baldwin, "Lucky imaging: high angular resolution imaging in the visible from the ground," *Astronomy & Astrophysics*, vol. 446, no. 2, pp. 739–745, 2006.
- [24] A. Schwartzman, M. Alterman, R. Zamir, and Y. Y. Schechner, "Turbulence-induced 2d correlated image distortion," in *IEEE International Conference on Computational Photography*, 2017, pp. 1–13.
- [25] Z. Mao, N. Chimitt, and S. H. Chan, "Accelerating atmospheric turbulence simulation via learned phase-to-space transform," in *Proceedings of the IEEE/CVF International Conference on Computer Vision*, 2021, pp. 14759–14768.
- [26] N. Chimitt, X. Zhang, Z. Mao, and S. H. Chan, "Real-time dense field phase-to-space simulation of imaging through atmospheric turbulence," *IEEE Transactions on Computational Imaging*, vol. 8, pp. 1159–1169, 2022.
- [27] X. Zhang, N. Chimitt, Y. Chi, Z. Mao, and S. H. Chan, "Spatio-temporal turbulence mitigation: A translational perspective," *arXiv preprint arXiv:2401.04244*, 2024.
- [28] R. Yasarla and V. M. Patel, "Learning to restore images degraded by atmospheric turbulence using uncertainty," in *2021 IEEE International Conference on Image Processing*, 2021, pp. 1694–1698.
- [29] S. N. Rai and C. Jawahar, "Removing atmospheric turbulence via deep adversarial learning," *IEEE Transactions on Image Processing*, vol. 31, pp. 2633–2646, 2022.
- [30] K. Mei and V. M. Patel, "Ltt-gan: Looking through turbulence by inverting gans," *IEEE Journal of Selected Topics in Signal Processing*, 2023.
- [31] J. L. Schonberger and J.-M. Frahm, "Structure-from-motion revisited," in *Proceedings of the IEEE conference on computer vision and pattern recognition*, 2016, pp. 4104–4113.
- [32] M. A. Fischler and R. C. Bolles, "Random sample consensus: a paradigm for model fitting with applications to image analysis and automated cartography," *Communications of the ACM*, vol. 24, no. 6, pp. 381–395, 1981.
- [33] B. Triggs, P. F. McLauchlan, R. I. Hartley, and A. W. Fitzgibbon, "Bundle adjustment—a modern synthesis," in *Vision Algorithms: Theory and Practice: International Workshop on Vision Algorithms Corfu, Greece, September 21–22, 1999 Proceedings*. Springer, 2000, pp. 298–372.
- [34] B. Mildenhall, P. P. Srinivasan, M. Tancik, J. T. Barron, R. Ramamoorthi, and R. Ng, "Nerf: Representing scenes as neural radiance fields for view synthesis," *Communications of the ACM*, vol. 65, no. 1, pp. 99–106, 2021.
- [35] T. Müller, A. Evans, C. Schied, and A. Keller, "Instant neural graphics primitives with a multiresolution hash encoding," *ACM Trans. Graph.*, vol. 41, no. 4, pp. 102:1–102:15, Jul. 2022. [Online]. Available: <https://doi.org/10.1145/3528223.3530127>
- [36] S. Fridovich-Keil, G. Meanti, F. R. Warburg, B. Recht, and A. Kanazawa, "K-planes: Explicit radiance fields in space, time, and appearance," in *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, 2023, pp. 12479–12488.
- [37] A. Cao and J. Johnson, "Hexplane: A fast representation for dynamic scenes," in *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, 2023, pp. 130–141.
- [38] B. Kerbl, G. Kopanas, T. Leimkühler, and G. Drettakis, "3d gaussian splatting for real-time radiance field rendering," *ACM Transactions on Graphics*, vol. 42, no. 4, 2023.
- [39] Y. Chen, Z. Chen, C. Zhang, F. Wang, X. Yang, Y. Wang, Z. Cai, L. Yang, H. Liu, and G. Lin, "Gaussianeditor: Swift and controllable 3d editing with gaussian splatting," *arXiv preprint arXiv:2311.14521*, 2023.
- [40] J. Fang, J. Wang, X. Zhang, L. Xie, and Q. Tian, "Gaussianeditor: Editing 3d gaussians delicately with text instructions," *arXiv preprint arXiv:2311.16037*, 2023.
- [41] J. Huang and H. Yu, "Point'n move: Interactive scene object manipulation on gaussian splatting radiance fields," *arXiv preprint arXiv:2311.16737*, 2023.
- [42] G. Kang, J. Yoo, J. Park, S. Nam, H. Im, S. Shin, S. Kim, and E. Park, "Selfsplat: Pose-free and 3d prior-free generalizable 3d gaussian splatting," *arXiv preprint arXiv:2411.17190*, 2024.
- [43] B. Ye, S. Liu, H. Xu, X. Li, M. Pollefeys, M.-H. Yang, and S. Peng, "No pose, no problem: Surprisingly simple 3d gaussian splats from sparse unposed images," *arXiv preprint arXiv:2410.24207*, 2024.
- [44] Y. Wang, T. Huang, H. Chen, and G. H. Lee, "Freesplat: Generalizable 3d gaussian splatting towards free view synthesis of indoor scenes," *Advances in Neural Information Processing Systems*, vol. 37, pp. 107326–107349, 2024.
- [45] W. Contributors, "Equatorial coordinate system," 2023, accessed: 2025-06-25. [Online]. Available: https://en.wikipedia.org/wiki/Equatorial_coordinate_system
- [46] A. McConahay, "The meridian flip," in *Using Sequence Generator Pro and Friends*. Springer, 2019, pp. 233–240.
- [47] W. C. Observatory, "What is a meridian flip," <https://wolfcreek.space/index.php/2020/11/03/what-is-a-meridian-flip/>, 2020, accessed: 2025-06-25.
- [48] E. Kraaikamp, "Autostakkert!" Software available from AutoStakkert website, 2024, accessed: 2024-03-20. [Online]. Available: <https://www.astrokrai.nl/software/latest.php>
- [49] J. Liang, J. Cao, Y. Fan, K. Zhang, R. Ranjan, Y. Li, R. Timofte, and L. Van Gool, "Vrt: A video restoration transformer," *arXiv preprint arXiv:2201.12288*, 2022.
- [50] L. F. Julià, P. Monasse, and M. Pierrot-Deseilligny, "The orthographic projection model for pose calibration of long focal images," *Image Processing On Line*, vol. 9, pp. 231–250, 2019.
- [51] E. D. Zhong, A. Lerer, J. H. Davis, and B. Berger, "Cryodrgn2: Ab initio neural reconstruction of 3d protein structures from real cryo-em images," in *Proceedings of the IEEE/CVF International Conference on Computer Vision*, 2021, pp. 4066–4075.
- [52] D. Ren, K. Zhang, Q. Wang, Q. Hu, and W. Zuo, "Neural blind deconvolution using deep priors," in *Proceedings of the IEEE/CVF conference on computer vision and pattern recognition*, 2020, pp. 3341–3350.
- [53] Z. Wang, S. Wu, W. Xie, M. Chen, and V. A. Prisacariu, "Nerf-: Neural radiance fields without known camera parameters," *arXiv preprint arXiv:2102.07064*, 2021.
- [54] C.-H. Lin, W.-C. Ma, A. Torralba, and S. Lucey, "Barf: Bundle-adjusting neural radiance fields," in *Proceedings of the IEEE/CVF international conference on computer vision*, 2021, pp. 5741–5751.
- [55] S. W. Zamir, A. Arora, S. Khan, M. Hayat, F. S. Khan, and M.-H. Yang, "Restormer: Efficient transformer for high-resolution image restoration," in *Proceedings of the IEEE/CVF conference on computer vision and pattern recognition*, 2022, pp. 5728–5739.
- [56] P.-E. Sarlin, C. Cadena, R. Siegwart, and M. Dymczyk, "From coarse to fine: Robust hierarchical localization at large scale," in *CVPR*, 2019.
- [57] X. He, J. Sun, Y. Wang, S. Peng, Q. Huang, H. Bao, and X. Zhou, "Detector-free structure from motion," in *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, 2024, pp. 21594–21603.
- [58] S. Wang, V. Leroy, Y. Cabon, B. Chidlovskii, and J. Revaud, "Dust3r: Geometric 3d vision made easy," in *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, 2024, pp. 20697–20709.
- [59] J. Y. Zhang, A. Lin, M. Kumar, T.-H. Yang, D. Ramanan, and S. Tulsiani, "Cameras as rays: Pose estimation via ray diffusion," *arXiv preprint arXiv:2402.14817*, 2024.
- [60] J. Wang, C. Rupprecht, and D. Novotny, "Posediffusion: Solving pose estimation via diffusion-aided bundle adjustment," in *Proceedings of the IEEE/CVF International Conference on Computer Vision*, 2023, pp. 9773–9783.
- [61] Z. Tang, Y. Fan, D. Wang, H. Xu, R. Ranjan, A. Schwing, and Z. Yan, "Mv-dust3r+: Single-stage scene reconstruction from sparse views in 2 seconds," *arXiv preprint arXiv:2412.06974*, 2024.
- [62] Y. Fu, S. Liu, A. Kulkarni, J. Kautz, A. A. Efros, and X. Wang, "Colmap-free 3d gaussian splatting," in *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, 2024, pp. 20796–20805.
- [63] Wikipedia, "Tiangong space station," 2024, [Online; accessed 28-March-2024]. [Online]. Available: https://en.wikipedia.org/wiki/Tiangong_space_station

Zhiming Chang received the M.S. degree from the College of Engineering, Peking University, Beijing, China, in 2025. His research interests include computational imaging, 3D reconstruction, and autonomous driving.

Boyang Liu received his Ph.D. degree from the International Centre for Radio Astronomy Research (ICRAR) at the University of Western Australia in 2021, and his bachelor's degree from Peking University in 2012. He is currently involved in the large optical telescope project EAST, initiated by Peking University. His interests also include public outreach in astronomy and the observational study of artificial objects in low Earth orbit.

Yifei Xia is an M.S. student in the School of Computer Science at Peking University. He received the B.S. from Peking University in 2023. His research interests include computational photography and video generation.

Youming Guo received the B.S. degree in automation from the University of Science and Technology of China, Anhui, China, in 2009 and the Ph.D. degree in signal and information processing from the Institute of Optics and Electronics, Chinese Academy of Sciences, Sichuan, China, in 2014. In 2014, he joined the Institute of Optics and Electronics, Chinese Academy of Sciences, where he is currently a Researcher and a Ph.D. Supervisor. He has authored or coauthored more than 70 papers in high-impact domestic and international journals. He has been granted more than 20 patents and was the principal investigator for projects including the National Natural Science Foundation of China (NSFC), high-tech initiatives, and innovation special zones. His primary research areas include adaptive optics, intelligent control, and computational imaging. Dr. Guo is a member of the Youth Innovation Promotion Association of the Chinese Academy of Sciences (CAS), a Western Young Scholar of CAS, and a council member of the Sichuan Youth Science and Technology Association.

Boxin Shi received the BE degree from the Beijing University of Posts and Telecommunications, the ME degree from Peking University, and the PhD degree from the University of Tokyo, in 2007, 2010, and 2013. He is currently a Boya Young Fellow Associate Professor (with tenure) and Research Professor at Peking University, where he leads the Camera Intelligence Lab. Before joining PKU, he did research with MIT Media Lab, Singapore University of Technology and Design, Nanyang Technological University, National Institute of Advanced Industrial Science and Technology, from 2013 to 2017. His papers were awarded as Best Paper, Runners-Up at CVPR 2024, ICCP 2015, and selected as Best Paper candidate at ICCV 2015. He is an associate editor of TPAMI/IJCV and an area chair of CVPR/ICCV/ECCV. He is a senior member of IEEE.

He Sun is an Assistant Professor at the National Biomedical Imaging Center, Peking University, China. Prior to joining Peking University, he was a Postdoctoral Researcher in the Department of Computing and Mathematical Sciences at the California Institute of Technology. He received the Ph.D. degree in Mechanical and Aerospace Engineering from Princeton University in 2019 and the bachelor's degree in Engineering Mechanics and Economics from Peking University in 2014. His research primarily focuses on computational imaging, which tightly integrates optics, control, signal processing and machine learning to push the boundary of scientific imaging. His past work has contributed to multiple real science missions, including the Event Horizon Telescope for black hole interferometric imaging and space telescope missions for exoplanet detection.