

DREAMSAT-POSE: SPACECRAFT POSE ESTIMATION FROM SINGLE-VIEW 3D RECONSTRUCTIONS AND LEARNED 2D-3D FEATURE MATCHING

Josiane Uwumukiza¹, Jocelyn Zhao², Giovanni Lavezzi³, Giacomo Battaglia⁴, Paolo Panicucci⁵, Minduli C. Wijayatunga⁶, Victor Rodriguez-Fernandez⁷, and Richard Linares⁸

6-DoF pose estimation is a critical task in autonomous rendezvous and proximity operations. In the case of an unknown target, this task becomes challenging as it shall be paired with the reconstruction of the target shape model. In this article, we propose a novel framework for single-shot shape and pose estimation of unknown spacecraft objects. Given a single image, we first reconstruct a 3D shape model of the target, then estimate the relative six-degrees-of-freedom pose by learning dense 2D-3D correspondences. The image features are extracted using a frozen DINOv3 vision transformer, while the geometric features are computed from the reconstructed point cloud using a trainable dynamic graph convolutional neural network encoder. A dual-stream transformer matcher refines descriptors through alternating self- and cross-attention, producing soft correspondences that are passed to a Perspective- n -Point solver for pose recovery. We evaluate the method on the SPE3R dataset and consider FoundationPose as a representative baseline for current state-of-the-art capabilities. Results show reliable pose estimates achieving 0.157 degrees mean pointing error using only a single image and reconstructed geometry, demonstrating strong generalization to unseen spacecraft.

INTRODUCTION

6-degree-of-freedom (6-DoF) pose estimation is a critical capability for space robotic missions, including on-orbit docking, inspection, navigation, and debris removal. In these scenarios, spacecraft must often estimate the pose of an uncooperative target using extremely limited sensing. In particular, monocular cameras are frequently the only available sensors, providing a single RGB (Red-Green-Blue) image without reliable depth or scale information, making accurate pose estimation challenging. Classical Perspective- n -Point (PnP) pipelines [1] and recent deep learning-based methods, particularly Convolutional Neural Networks (CNNs) [2, 3], frequently struggle to generalize to unseen spacecraft [4, 5], limiting their operational utility.

To address these challenges, we leverage recent advances in representation learning. DINOv3 [6] provides viewpoint-consistent visual descriptors, while Dynamic Graph CNN (DGCNN) [7]

¹Undergraduate Student, Wellesley College, Wellesley, MA 02481, USA, ju100@wellesley.edu.

²Undergraduate Student, Massachusetts Institute of Technology, Cambridge, MA 02139, USA, jjz300@mit.edu.

³Research Scientist, Massachusetts Institute of Technology, Cambridge, MA 02139, USA, glavezzi@mit.edu.

⁴Ph.D. Student, Politecnico di Milano, Milan, Italy, giacomo.battaglia@polimi.it.

⁵Assistant Professor, Politecnico di Milano, Milan, Italy, paolo.panicucci@polimi.it.

⁶Assistant Professor, University of Illinois Urbana-Champaign, Urbana, IL 61801, USA, minduli@illinois.edu.

⁷Associate Professor, Universidad Politécnica de Madrid, Madrid, Spain, victor.rfernandez@upm.es.

⁸Associate Professor, Massachusetts Institute of Technology, Cambridge, MA 02139, USA, linaresr@mit.edu.

extracts expressive features from 3D point clouds by modeling local geometric relationships. While 2D matchers like SuperGlue [8] or registration tools like MinCD-PnP [9] have made progress, they remain limited by a lack of direct 2D-3D matching or a reliance on metrically scaled Computer Aided Design (CAD) models often obtained from depth sensors or structure-from-motion, which are not available in our setting. Alternatively, generative frameworks like DreamSat [10, 11] use diffusion priors to reconstruct 3D geometries from single observations.

In this work, we propose a single-view, correspondence-based framework for 6-DoF pose estimation that leverages reconstructed object geometry without requiring depth or multi-view supervision. Building on our previous work, DreamSat [10, 11], which enables single-view 3D reconstruction of spacecraft, we use the reconstructed structure as a geometric prior to guide pose estimation. This reformulates the problem from direct pose regression into a structured correspondence problem between 2D image features and 3D model points. We extract robust 2D features and keypoints using a frozen vision transformer (DINOv3) [6] from the image, train a graph-based network (DGCNN) [7] to encode the 3D CAD, and learn soft 2D-3D correspondences via a lightweight transformer with self- and cross-attention. These correspondences are then used within a classical PnP solver to recover the pose of the object. To evaluate the proposed framework, we utilize the SPE3R dataset [12, 13], including 64 unique spacecraft models. We also consider NVIDIA’s FoundationPose [14], a generalist 6D pose estimator capable of zero-shot generalization to unseen objects, as a representative baseline for current state-of-the-art capabilities. It must be noted that FoundationPose requires depth maps as inputs, whereas our methodology operates strictly on monocular RGB data; as such, FoundationPose serves as a best-case performance benchmark rather than a direct equivalent. We employ two primary experimental configurations: a ground-truth-based assessment to establish a performance upper bound, and a reconstruction-based assessment to evaluate the system’s utility for unknown targets. This analysis investigates whether generative 3D proxies can effectively substitute for a-priori CAD models in autonomous proximity operations. By combining modern representation learning with geometric constraints, our method achieves viable pose estimation from a single RGB image.

The paper is organized as follows. First, we introduce the single-view correspondence-based framework for 6-DoF pose estimation, which is called DreamSat-Pose. Then, we describe the mission scenario, dataset and evaluation metrics. Subsequently, we discuss the experimental design and results of the proposed pose estimation pipeline. Lastly, we provide final remarks to conclude the paper.

DREAMSAT-POSE

Our pipeline consists of four main components: (a) 3D reconstruction from a single RGB image of the space object, (b) feature extraction on the image and 3D reconstructed model, (c) feature matching to produce 2D-3D correspondences, and (d) pose solving to estimate the camera pose from matches. Figure 1 illustrates the flow of data through these stages. An RGB image is given as input to DreamSat to obtain a 3D reconstructed model. The same image is encoded using a frozen DINOv3 vision transformer to extract 2D features, while the reconstructed CAD model is processed by a DGCNN to obtain 3D descriptors. A transformer-based cross-attention module produces a soft correspondence matches. Top- k correspondences are used in an EPnP solver [15] to estimate pose.




Figure 1: Overview of the DreamSat-Pose pipeline

Input Image Processing Techniques

The input images from the SPE3R dataset are utilized at their native resolution 256×256 for the majority of the pipeline. However, we observed that when a spacecraft occupies only a small portion of the total image area the pose estimation performance degrades significantly. In these instances, the effective resolution of the target is insufficient for the DINOv3 encoder to extract discriminative visual descriptors, leading to unstable 2D-3D correspondences. To mitigate this, we employ a bounding-box-aware cropping and upscaling strategy. First, we utilize the segmentation mask to identify the tightest bounding box surrounding the detected spacecraft. This region is then cropped from the original frame and rescaled back to the 256×256 input size using bilinear interpolation. The corresponding segmentation mask is resized using nearest-neighbor interpolation to preserve binary foreground/background labels. This process maximizes the pixel density of the satellite’s structure within the input window and provides the feature matcher with a more consistent representation, thereby improving the robustness of the PnP solver. Future work will focus on improving the quality of correspondences, which, despite the cropping and upscaling strategy, remain a challenge. We will also explore image enhancements such as super-resolution and denoising and higher resolution images to further improve feature matching and reconstruction fidelity.

Single-View 3D Reconstruction

DreamSat integrates and benchmarks multiple state-of-the-art single-view generative 3D reconstruction models as described in our previous work [10, 11]. In this work, we adopt Hunyuan-3D-2.0 [16] due to its strong performance across both 2D perceptual and 3D geometric metrics on custom spacecraft and asteroid datasets, making it well-suited for downstream pose estimation tasks. As a result, DreamSat-Pose leverages Hunyuan-3D-2.0 reconstructions as the primary geometric representation in the proposed pose estimation pipeline. By reconstructing spacecraft geometry from a single observation, the proposed framework enables pose estimation in scenarios where traditional

model-based approaches are not feasible and no additional observations and depth information are available.

Image Feature Extraction

To enable generalization to unseen targets, a frozen DINOv3 ViT-B/16 [6] is used as the 2D feature extractor and keypoint selector. Given a 256×256 RGB image, the model produces a dense grid of patch descriptors (size 16×16 with patch size 16), where each descriptor is 768-dimensional. These descriptors are projected to 256 dimensions via a linear layer to reduce memory and align with the 3D feature space produced by the DGCNN encoder. In addition to descriptors, a lightweight score head produces a scalar saliency for each patch, highlighting informative regions such as edges and textured areas while remaining robust to harsh orbital lighting. Keypoints are selected by thresholding this score, yielding a set of candidate 2D feature points. A foreground mask is applied to ensure that the selected key points are on the object.

3D Feature Extraction

After 3D reconstruction, 1024 points are uniformly sampled from the surface mesh and normalized through centering and scaling to ensure a geometry-independent representation, serving as our candidate 3D keypoints. These points are fed into the DGCNN network, which constructs a k -NN graph on the points ($k = 20$) neighbors and applies a series of EdgeConv layers. EdgeConv is a graph convolution operator that learns local geometric features by computing edge representations between each point and its neighbors, allowing the network to capture both the properties of individual points and their local spatial relationships. In an EdgeConv, for each point, a small neural network processes the concatenation of the point’s feature and the relative feature of a neighbor, thereby encoding geometric relationships. The descriptors from four successive layers are concatenated to aggregate local surface curvature with global object context, resulting in a 256-dimensional embedding. After each EdgeConv, a max-pooling over the neighbors produces an updated feature for each point. We initialize the DGCNN randomly and allow it to train. This helped the 3D features to adapt to the distribution of our object models. The DGCNN is optimized using the matcher’s alignment and reprojection losses to co-adapt 3D geometric descriptors into a shared latent space with the 2D image features.

Cross-Modal Matching Network

The 2D-3D matching module takes as input the set of 2D keypoint descriptors $\{f_i^{2D}\}$ and 3D point descriptors $\{f_j^{3D}\}$ (of size N and M respectively), and outputs a soft correspondence matrix $P \in \mathbb{R}^{N \times M}$, as shown in Fig. 2. The architecture follows a Transformer-based dual encoder-decoder design, inspired by SuperGlue [8]. Both descriptor sets are projected into a common 256-dimensional space via linear layers, then augmented with positional encodings: normalized 2D coordinates ($x, y \in [0, 1]$) and normalized 3D coordinates (XYZ) are each passed through small Multilayer Perceptrons (MLPs) and added to their respective descriptors. Separate self-attention layers perform intra-modality refinement, followed by interleaved cross-attention layers where queries from one modality attend to keys and values from the other, enabling correspondence reasoning across modalities. The cosine similarity matrix $S \in \mathbb{R}^{N \times M}$ is computed using the dot product of the L_2 -normalized features:

$$S_{ij} = \langle \hat{f}_i^{2D}, \hat{f}_j^{3D} \rangle \quad (1)$$

Sinkhorn normalization[17] is then applied to S , yielding a doubly-stochastic matrix P that enforces one-to-one matching and remains differentiable during training. At inference, the top match per 2D point is selected as $\arg \max_j P_{ij}$, with matches below a confidence threshold discarded.




Figure 2: Dual-stream transformer architecture for 2D–3D correspondence matching. 2D image features and 3D geometric features are refined through alternating self- and cross-attention layers, followed by Sinkhorn normalization to produce the soft correspondence matrix P .

Loss Functions

We introduce a supervised geometric objective that leverages known camera poses to generate correspondence targets. Rather than direct parameter regression, the ground-truth (GT) rotation R_{gt} , translation t_{gt} , and camera intrinsics K are used to project each sampled 3D point $X_j \in \mathbb{R}^3$ into the image plane as $\tilde{u}_j = \pi(K(R_{gt}X_j + t_{gt}))$, where $\pi(\cdot)$ denotes perspective projection. For each selected 2D keypoint $u_i \in \mathbb{R}^2$, a soft GT correspondence distribution P_{ij}^{gt} is constructed by applying a local softmax over Gaussian-weighted Euclidean distances $d_{ij} = \|u_i - \tilde{u}_j\|_2$ within a top- k neighborhood. This soft formulation accounts for the inherent ambiguity of the correspondence problem, where multiple 3D points may project near a single 2D keypoint due to sparse sampling or local symmetry, thereby avoiding brittle hard assignments [18]. The predicted soft assignment matrix P is then optimized via the soft correspondence loss:

$$\mathcal{L}_{\text{soft}} = -\frac{1}{|\mathcal{V}|} \sum_{i \in \mathcal{V}} \sum_{j=1}^M P_{ij}^{gt} \log P_{ij} \quad (2)$$

where $\mathcal{V}$ denotes the set of valid 2D keypoints and M is the total number of 3D points. To provide a strong geometric signal to the DGCNN, we employ an expected-XYZ loss [19] to align the predicted expected 3D point $\hat{X}_i = \sum_{j=1}^M P_{ij} X_j$ with the target 3D point $X_i^{gt} = \sum_{j=1}^M P_{ij}^{gt} X_j$:

$$\mathcal{L}_{\text{xyz}} = \frac{1}{|\mathcal{V}|} \sum_{i \in \mathcal{V}} \left\| \hat{X}_i - X_i^{gt} \right\|_2^2 \quad (3)$$

Direct pixel-space consistency is enforced by the reprojection loss [2, 20], which minimizes the distance between the observed 2D keypoint u_i and the reprojection of the predicted soft match $\hat{X}_i$:

$$\mathcal{L}_{\text{reproj}} = \frac{1}{|\mathcal{V}|} \sum_{i \in \mathcal{V}} \left\| \pi(K(R_{gt}\hat{X}_i + t_{gt})) - u_i \right\|_2 \quad (4)$$

The assignment matrix is further regularized by a mutual nearest-neighbor loss $\mathcal{L}_{\text{mnn}}$ to ensure bidirectional matching coherence [8, 21]. The total supervised objective is a weighted combination of these terms:

$$\mathcal{L}_{\text{sup}} = \lambda_{\text{soft}}\mathcal{L}_{\text{soft}} + \lambda_{\text{xyz}}\mathcal{L}_{\text{xyz}} + \lambda_{\text{mnn}}\mathcal{L}_{\text{mnn}} + \lambda_{\text{reproj}}\mathcal{L}_{\text{reproj}} \quad (5)$$

using weights $\lambda_{\text{soft}} = 1.0$, $\lambda_{\text{xyz}} = 1.0$, $\lambda_{\text{mnn}} = 0.5$, and $\lambda_{\text{reproj}} = 0.01$. This formulation allows the matcher and DGCNN to co-adapt into a shared latent space, stabilized by an annealing curriculum that gradually sharpens the Gaussian bandwidth σ and neighborhood size k over time.

Pose Calculation

To translate learned correspondences into a 6-DoF pose (R, t) , we employ a robust geometric pipeline. Discrete matches are first filtered by the confidence scores in P , requiring a minimum of six points to avoid underdetermined solutions. To handle outliers arising from geometric symmetries or reconstruction noise, an initial pose is estimated using the Efficient PnP (EPnP) algorithm [15] within a RANSAC (RANDOM SAmple Consensus) framework of 200 iterations. This estimate is subsequently refined through iterative nonlinear Levenberg-Marquardt optimization to minimize the total reprojection error of the identified inlier set. This hybrid approach combines the generalization of learned features with the interpretability and precision of classical geometric solvers. Figure 3 shows the qualitative results of the proposed DreamSat-Pose pipeline.

EXPERIMENTAL SETUP

Mission Scenario and Assumptions


The proposed methodology aims to estimate the 6-DoF pose of an unknown spacecraft from a single RGB image without access to GT geometry, GT CAD models, or multi-view observations. We assume monocular sensing, unknown spacecraft 3D model, and no depth measurements. We assume non-cooperative targets, with the absence of fiducial markers or predefined landmarks. These assumptions define a challenging regime for pose estimation. Under these constraints, the system must rely on learned geometric reasoning. These constraints also influence design decisions throughout the pipeline. For example, reconstruction is performed from a single image rather than multi-view geometry. Similarly, learned descriptors are used instead of handcrafted features to improve robustness under varying conditions. With these constraints, this formulation closely reflects realistic operational conditions in space autonomy systems.

Dataset


The Spacecraft Pose Estimation and 3D Reconstruction (SPE3R) dataset [12] is an open-source satellite image dataset developed at Stanford’s Space Rendezvous Laboratory (SLAB) [13]. SPE3R comprises 64 unique satellite models, and each satellite model is characterized by 1,000 synthetic images and their corresponding binary segmentation masks, GT pose labels (orientation and position) and 3D point clouds. The synthetic images in SPE3R are of size 256×256 pixels. These images depict a satellite against either a black or Earth background, with added illumination features such as harsh lighting and shadows.

Data Splitting


For comparison, we adopt the same dataset split as used in the literature [2]. The dataset is divided into 57 training spacecraft and 7 unseen test spacecraft. Each spacecraft contains 1000 images with




(a) Input Image



(b) DINO Keypoints



(c) 3D Features (DGCNN)



(d) Pose Alignment

Figure 3: Qualitative results of the proposed DreamSat-Pose pipeline. (a) Input RGB image. (b) Keypoints selected by DINOv3, which are distributed across geometrically informative regions. (c) Learned 3D feature representation of the reconstructed object using DGCNN. (d) Estimated pose alignment in 3D space, showing agreement between the reconstructed model (blue) and ground truth (orange).

unique poses. This results in 7000 images. The train set images per object are split 80:20 for train and validation such that the validation set is a group of unseen images of a known spacecraft. Both training and inference are performed on a workstation equipped with NVIDIA RTX 4090 hardware.

Reconstruction Metrics

To evaluate the fidelity of the 3D reconstructions, we employ 3D geometric metrics [11]. Specifically, the structural accuracy is assessed by comparing the reconstructed meshes directly against GT CAD models using three metrics: the 3D Volumetric Intersection over Union (IoU), which measures global spatial overlap; the Chamfer Distance (CD), which quantifies the average proximity between predicted and true point clouds; and the Hausdorff Distance (HD), which captures the maximum geometric deviation between the reconstructed and GT surfaces. To ensure accurate evaluation, each reconstructed model is normalized and rotated to align its coordinate frame with the canonical SPE3R GT frame. Together, these metrics provide complementary measures of reconstruction quality and ensure that the generated models are geometrically reliable for downstream pose estimation tasks.

Pose Estimation Metrics

From the Satellite Pose Estimation Challenge [4], we consider the orientation error e_R , which is computed using the quaternion representation to avoid singularities and ensure numerical stability. The error is defined as the angular distance between the estimated quaternion q_{est} and the ground truth q_{gt} :

$$e_R = 2 \cos^{-1}(|q_{gt} \cdot q_{est}|) \quad (6)$$

We evaluate the pointing accuracy of the translation vector, known as the Line-of-Sight (LOS) error. This metric decouples the directional accuracy from the absolute range estimation. The pointing error e_p is calculated as the angular separation between the estimated and ground truth translation vectors:

$$e_p = \cos^{-1} \left(\frac{\mathbf{t}_{est} \cdot \mathbf{t}_{gt}}{\|\mathbf{t}_{est}\| \|\mathbf{t}_{gt}\|} \right) \quad (7)$$

The inclusion of e_p is related to the inherent limitations of monocular vision systems, i.e. the infeasibility to retrieve the absolute scale of a reconstructed object from a single monocular image without prior knowledge of the target’s physical dimensions. Absolute distance can generally only be recovered through triangulation, which requires multiple views and a known change in the chaser’s position and velocity, or by assuming a fixed model scale during the 2D-3D matching process. Consequently, the translation error is often dominated by range inaccuracies, while e_p provides a more direct measure of the geometric alignment performance.

To evaluate the geometric consistency of the estimated pose, we compute the Mean Reprojection Error (MRE) using the set of 2D-3D correspondences that survive the PnP RANSAC filtering. The MRE is defined as the average Euclidean distance, in pixels, between the observed 2D keypoints and the 3D points projected back onto the image plane using the estimated rotation $\mathbf{R}$, translation $\mathbf{t}$, and camera intrinsics $\mathbf{K}$. This metric quantifies the alignment accuracy of the PnP solver and serves as a measure of how well the reconstructed 3D model represents the visual features observed in the 2D image.

Frame ambiguity

One of the challenging aspect of evaluating pose for reconstructed spacecraft models is that the object-fixed reference frame is not uniquely defined by the geometry alone. In monocular spacecraft shape-and-pose estimation, the estimated pose especially orientation is meaningful only relative to the body frame assigned to the reconstructed shape; however, when that frame is generated from structural rules such as aligning axes with the principal or largest dimensions of the object, the rule may constrain the axis ordering without uniquely determining the positive and negative directions of those axes. Dimension-based body-axis alignment cannot enforce consistent positive and negative axis directions across different spacecraft models, so a prediction may be penalized even when the estimated shape and pose are geometrically consistent with the image, but expressed in a different valid body-frame convention [2]. This issue is further complicated by the fact that the reconstructed object is not a perfect replica of the GT model. Small geometric distortions, missing structures, smoothed components, or reconstruction artifacts can change the apparent principal dimensions of the shape, making frame extraction from the reconstruction less stable than frame extraction from the GT mesh. Thus, even if the same axis-assignment rule is applied to both models, the reconstructed frame may not align exactly with the GT frame before the sign ambiguity is considered. For a right-handed coordinate system, the valid sign-preserving frame conventions are the original frame and the three 180° rotations about the body axes, (x, y, z) , $(x, -y, -z)$, $(-x, y, -z)$, and $(-x, -y, z)$, corresponding to the permutation set used in the paper’s ambiguity-aware rotation loss [2, Eq. (8)-(9)]. This is related to, but distinct from, object symmetry: even when a spacecraft is not treated as physically symmetric, the constructed coordinate frame can still have an arbitrary sign convention. Consequently, a direct comparison between the estimated rotation and the GT rotation can report a large orientation error simply because the reconstructed body frame is flipped relative to the GT convention. To avoid treating this arbitrary convention mismatch as a physical pose error, we evaluate the rotation under each valid right-handed frame permutation and retain the convention that minimizes the pose error [2, Eq. (9)].

FoundationPose

We adopt NVIDIA’s FoundationPose [14] as a benchmark for our pipeline, since FoundationPose is a generalist 6D pose estimator and tracker trained on large-scale synthetic data from diverse 3D object repositories with domain randomization, demonstrating strong zero-shot generalization to unseen objects without fine-tuning. The method utilizes a transformer-based architecture to establish an iterative pose refinement process, where a neural network predicts the relative transformation required to align a 3D model with the observed image features.

FoundationPose is a model-based 6D object pose estimation method that, for each pose, takes as input an RGB image, an object mask, a depth map, a normalized 3D mesh (.obj), and camera intrinsics. In our setup, all required inputs are provided by SP3ER except for the depth map. Since absolute depth information is not directly available in the normalized dataset, we generate it by rendering the GT 3D model from the known camera pose. Although this provides references with idealized depth information, it establishes a baseline to compare the underlying pose estimation logic. However, it must be noted that FoundationPose’s dependency on depth maps limits its feasibility for typical monocular space missions where active depth sensors are unavailable.

EXPERIMENTAL DESIGN AND RESULTS

To quantify the performance and robustness of the DreamSat-Pose framework, we conduct a series of comparative simulations designed to isolate the impact of geometric fidelity on pose accuracy. Our experimental campaign is divided into two primary configurations. The first experiment is GT-based pose estimation. In this configuration, we provide both our DreamSat-Pose and FoundationPose with GT meshes provided by the SPE3R dataset, to establish an approximate upper bound for pose estimation performance. This allows us to determine the maximum achievable accuracy of the matching and tracking algorithms when structural uncertainty is removed. The second experiment is reconstruction-based pose estimation. To evaluate the pipeline’s utility for uncooperative targets, we perform pose estimation using the 3D meshes reconstructed directly from single-view SPE3R images. First, we consider only one reconstruction per test spacecraft. Then, due to reconstruction processing times, we select a subset of 100 images per spacecraft, split equally between black and Earth backgrounds (50 images each). These experiments serve as a full pipeline test, determining if a machine-learned geometric proxy can effectively substitute for a pre-existing CAD model in proximity operations. While the results presented are based on GT masks, preliminary experiments using the Segment Anything Model (SAM) [22] to automate foreground segmentation achieved comparable performance. This highlights the pipeline’s potential for deployment in scenarios where manual mask annotations are unavailable. Table 1 summarizes the pose estimation performance for DreamSat-Pose and FoundationPose which are discussed in the following sections.

Model	Test	N	Orientation Error (°)				Pointing Error (°)			
			Mean	Median	Min	Max	Mean	Median	Min	Max
FoundationPose	GT Test	7000	29.64	2.31	0.05	180.00	0.095	0.071	0.001	6.188
FoundationPose	Recon Test	7000	123.49	140.51	0.23	179.99	0.919	0.492	0.003	5.594
FoundationPose	Recon 100per	700	114.36	133.74	0.89	179.98	0.809	0.445	0.000	6.626
DreamSat-Pose	Validation	11395	29.32	12.77	0.27	118.08	0.092	0.050	0.000	3.547
DreamSat-Pose	GT Test	7000	48.90	47.65	0.57	110.88	0.133	0.094	0.001	2.637
DreamSat-Pose	Recon Test	7000	58.73	61.89	1.02	118.09	0.157	0.118	0.001	2.349
DreamSat-Pose	Recon 100per	700	57.84	59.30	2.91	114.93	0.179	0.123	0.004	2.139

Table 1: Pose estimation performance for FoundationPose and DreamSat-Pose. Orientation and pointing errors are reported in degrees.

DreamSat-Pose

In this section, the results of the experimental campaign for DreamSat-Pose are reported. Figures 4 and 5 illustrate the error distributions for both LOS direction and orientation for the validation set, the test set using GT CAD meshes (GT-Test), and the full pipeline using single-view reconstructed meshes (Recon-Test). The directional error demonstrates remarkable precision across all categories, with mean errors remaining well below 1° . This indicates that the pipeline is highly robust in identifying the target’s relative bearing, regardless of whether a GT or a reconstructed mesh is utilized. The low variance in the GT-Test and Recon-Test configurations suggests that the 2D-3D correspondence matching successfully captures the geometric centroid of the spacecraft. In contrast, the orientation error exhibits higher sensitivity to the fidelity of the 3D model. While the validation set achieves a mean orientation error of approximately 29.32° , the performance on the test set degrades when moving from GT meshes (mean $\approx 48.90^\circ$) to reconstructed meshes (mean $\approx 58.73^\circ$). The wider error distribution in the reconstruction-based configuration reflects structural uncertainties in the geometric proxies, particularly as no image enhancement is utilized in

this study. Future integration of techniques such as super-resolution or denoising could improve the fidelity of the single-view 3D reconstructions and bridge the performance gap with GT CAD models. Nevertheless, these results confirm that machine-learned reconstructions serve as a functional proxy for pose estimation in uncooperative scenarios where no prior model is available.




Figure 4: Pointing error distributions for DreamSat-Pose.




Figure 5: Orientation error distributions for DreamSat-Pose.

To assess the fidelity of the single-view 3D reconstructions in the Recon-Test results, we evaluate the reconstruction metrics. As shown in Fig. 6, Hunyuan-3D-2.0 achieves performance comparable to that reported in DreamSat [11], despite a significant increase in data complexity. While the original study uses screenshots of isolated models, i.e. lacking complex illumination or background contexts, the current work employs realistic synthetic imagery from the SPE3R dataset. Specifically, across 700 reconstructions, the distributions of CD, HD, and 3D IoU indicate that the generated point clouds preserve overall spacecraft geometry with moderate structural fidelity despite the challenges of single-view reconstruction. However, not all outputs conform to a satisfactory reconstruction, suggesting the need to investigate the use of image processing techniques to enhance the visual quality of the inputs provided to the reconstruction model. Notably, no clear correlation is observed between reconstruction quality metrics and downstream pose estimation performance, as samples with lower reconstruction error do not consistently produce lower pointing or orientation errors. The MRE of the RANSAC inlier points depends on the chosen threshold, averaging 0.9, 2.4, and 3.7 pixels for the 2, 5, and 8 pixel Ransac thresholds, respectively.




Figure 6: Distribution of 3D reconstruction performance metrics across the evaluation dataset. Lower values indicate better performance for CD and HD, whereas values closer to one are better for IoU.

Alternative Approach: FoundationPose

In this section, the results of the experimental campaign for FoundationPose are reported. Similarly, FoundationPose is evaluated under two settings: a GT test setting, in which the method is supplied with the original SPE3R CAD meshes, and a reconstruction-based setting, in which the method uses meshes reconstructed from single-view images.




Figure 7: Pointing error distributions for FoundationPose across GT test and reconstructed test settings.




Figure 8: Orientation error distributions for FoundationPose across GT test and reconstructed test settings.

The pointing and orientation error distributions for FoundationPose are shown in Fig. 7 and Fig. 8. On the GT CAD models, FoundationPose achieves strong pose estimation performance,

with a mean pointing error of 0.095° and a median pointing error of 0.071° . The narrow pointing error distribution indicates that the method consistently recovers the correct viewing direction of the spacecraft. Orientation estimation is also highly accurate for many samples, with a median error of 2.31° in the GT test setting. However, the mean orientation error is substantially higher at 29.64° , indicating the presence of a subset of cases with large angular deviations. When evaluated using reconstructed meshes rather than GT CAD models, FoundationPose exhibits a noticeable reduction in pose accuracy. The mean pointing error increases from 0.095° in the GT setting to 0.919° in the reconstructed setting. This decrease is expected, as reconstructed meshes contain geometric imperfections, surface smoothing artifacts, and structural inconsistencies that are absent from the GT models. Despite these limitations, the pointing estimates remain relatively accurate, with a mean pointing error below 1° . These results suggest that reconstructed geometry can still provide useful structural information for downstream pose estimation. The degradation is more pronounced for orientation estimation. In the reconstructed setting, the mean orientation error increases to 123.49° , with a median error of 140.51° . This indicates that FoundationPose is more sensitive to reconstruction artifacts when estimating full rotational alignment than when recovering the spacecraft line of sight. Although FoundationPose achieves very low orientation errors on the GT models, these results are not immune to body-frame ambiguity. The reconstructed and GT body frames may differ by a valid right-handed frame convention even when the estimated pose is geometrically correct. When the ambiguity-aware evaluation is applied, the mean orientation error is reduced from 123.49° to 53.35° for 1 reconstruction per 1000 images and from 114.36° to 49.84° for 100 reconstructions per object, with the updated per-satellite orientation distributions shown in Figure 9. The ambiguity-aware results are highlighted in red.




Figure 9: Orientation error distributions for FoundationPose on the GT and reconstructed test sets, with frame ambiguity handled by evaluating valid body-frame flips.

This ambiguity primarily affects quantitative evaluation against a fixed GT reference frame. In practical deployment, however, the reconstructed body frame defines the reference frame used by the downstream pose estimator. Once that frame convention is established, the ambiguity disappears because all subsequent pose estimates are expressed consistently in the reconstructed coordinate system rather than being compared against the GT frame. This result is consistent with the fact that orientation estimation depends strongly on fine-grained geometric details, while pointing estimation can remain accurate even when the reconstructed mesh is imperfect. It is important to note that FoundationPose was supplied with a depth map during inference. In this experimental setup, the depth map is generated from the known spacecraft geometry and camera pose, which provides idealized geometric information that would generally not be available in practical monocular on-orbit scenarios. In real missions, high-quality depth sensing may be unavailable or degraded due to sensor noise, sparse measurements, illumination constraints, or the absence of active depth sensors.

Consequently, the FoundationPose results should be interpreted as an approximate upper bound under favorable sensing assumptions.




Figure 10: Per-satellite orientation error distributions for FoundationPose in the GT-Test setting.

FoundationPose: Effect of Spacecraft Symmetry

Although the overall pointing performance of FoundationPose remains strong, the orientation error distributions reveal a clear object-dependent behavior. In particular, Fig. 8 shows that several samples exhibit large orientation errors despite accurate pointing estimates. This behavior is consistent with the inherent ambiguity of monocular pose estimation for geometrically symmetric spacecraft, where multiple rotations can produce nearly identical image projections. Figure 10 shows that this ambiguity is highly object-dependent. FoundationPose achieves orientation errors below approximately 5° for *aquarius_ud*, *galileo_actual_05*, and *herschel*, whose distinctive geometric structures uniquely constrain their orientation. In contrast, *cygnss_solo_39*, *eo-1_final*, *loral-1300com-main*, and *smart_1* exhibit substantially larger orientation errors despite maintaining accurate pointing estimates. This trend is further illustrated in Fig. 11, which compares representative spacecraft meshes in a common view. Spacecraft with distinctive, asymmetric structures provide stronger geometric constraints for rotation recovery. Conversely, spacecraft with repeated components or approximate planes of symmetry can admit multiple visually plausible orientations. In these cases, FoundationPose may predict an orientation that is visually similar to the ground truth but differs by a rotation about a symmetry axis. Therefore, the reported orientation error may overestimate the practical localization error for highly symmetric spacecraft. The consistently

low pointing errors nevertheless indicate that FoundationPose recovers the spacecraft line of sight accurately, even when orientation remains ambiguous.




Figure 11: Representative spacecraft meshes projected in a common view: EO-1, Galileo, Herschel, Aquarius, CYGNSS, Smart-1, and Loral-1300. The vertical line indicates an approximate plane of symmetry where present.

Runtime and Memory Performance

Once initialized, the DreamSat-Pose runtime is about 0.034 seconds per image and includes 2D feature extraction using DINOv3, 3D feature computation using the DGCNN representation, transformer-based correspondence matching, and final pose estimation using EPnP. The maximum system memory consumption during end-to-end inference is approximately 1.62 GB of RAM. The 3D reconstruction takes about 35-40 seconds when using the Hunyuan-3D-2.0 model. However, a faster option is available, Hunyuan-3D-2.0-mini-turbo-flash, which takes about 4 seconds per image; its usage and a trade off between reconstruction quality and speed will be explored in future work. On the other hand, FoundationPose runtime is approximately 1 second per image.

CONCLUSION

This paper proposes DreamSat-Pose, a framework for 6-DoF spacecraft pose estimation that operates under the minimal assumptions of a single RGB image without requiring depth sensing or a-priori CAD models. By reformulating pose estimation as a correspondence problem between 2D image features and generative 3D reconstructions, the pipeline provides a scalable solution for navigating around uncooperative or unknown targets. The results from our experiments demonstrate that high-quality 3D reconstructions can effectively serve as proxies for GT geometry. Moreover, the system demonstrates strong generalization to unseen spacecraft, suggesting that the learned representations capture fundamental geometric structures rather than object-specific appearance. At the same time, the results show that orientation remains the most challenging component of the pose estimation task. While the pipeline achieves consistently accurate pointing estimates, the orientation error remains more sensitive to reconstruction artifacts, object symmetries, repeated structural patterns, and body-frame ambiguities. The strong performance observed when using GT meshes indicates that the learned 2D-3D correspondence and PnP-based pose estimation pipeline can recover accurate pose information when reliable geometry is available. Therefore, as single-view reconstruction outputs continue to improve, the reconstruction-based results suggest a clear possibility for using generated 3D models as practical geometric proxies in monocular spacecraft pose estimation. The significance of this work to the future of autonomous space operations is twofold. First, it enables autonomous missions to unprepared and unknown targets by bypassing the current bottleneck where navigation systems rely on GT CAD models or cooperative markers; data that is often lost or outdated for defunct objects and debris. Second, it offers a path toward a significant reduction in sensor size, weight, and power (SWaP) by demonstrating that accurate 6-DoF estimation can be achieved through a monocular pipeline rather than relying on heavy, power-hungry LiDAR or stereo-camera rigs. Despite these promising outcomes, several limitations remain for operational

deployment. The current end-to-end inference latency can be suitable for situational awareness but might exceed the requirements for real-time closed-loop control and autonomous docking. Furthermore, orientation estimation remains a challenge due to the inherent geometric symmetries and repeated structural patterns of many spacecraft. Another limitation is the resolution and framing of the input imagery. The SPE3R images are 256×256 pixels, and the spacecraft does not always occupy the full image. Instead, the object of interest is often embedded within the frame, and for distant targets, the visible satellite can occupy only a small number of pixels. This limits the amount of visual detail available to the feature extractor and makes it harder to obtain meaningful image embeddings and reliable 2D-3D correspondences. Also, the reliance on photorealistic renderings leaves a realism gap, as the model has yet to be challenged by the sensor noise, compression artifacts, and extreme specular reflections characteristic of actual flight imagery. Future work will focus on investigating model distillation and weight quantization to reduce the computational footprint and memory usage, multi-hypothesis pose estimation and uncertainty modeling to better capture symmetric configurations, and validating the pipeline against realistic sensor noise and motion blur. Future work should also explore higher-resolution inputs, improved cropping and upscaling strategies, and image enhancement techniques such as super-resolution and denoising to improve feature extraction when the spacecraft occupies only a small region of the image. Ultimately, DreamSat-Pose aims to provide a robust framework for autonomous interaction with the complex and unpredictable modern orbital environment.

CODE AVAILABILITY

The code for this work can be accessed on GitHub (<https://dreamsatpose.github.io/>).

ACKNOWLEDGEMENTS

From the MIT side, the research was sponsored by the Department of the Air Force Artificial Intelligence Accelerator and was accomplished under Cooperative Agreement Number FA8750-19-2-1000. The views and conclusions contained in this document are those of the authors and should not be interpreted as representing the official policies, either expressed or implied, of the Department of the Air Force or the U.S. Government. The U.S. Government is authorized to reproduce and distribute reprints for Government purposes notwithstanding any copyright notation herein. Paolo Panicucci is supported by the Progetto Rocca Fellowship.

REFERENCES

- [1] L. Pauly, W. Rharbaoui, C. Shneider, A. Rathinam, V. Gaudillière, and D. Aouada, “A survey on deep learning-based monocular spacecraft pose estimation: Current state, limitations and prospects,” in *Acta Astronautica*, vol. 212, 2023, pp. 339–360. [Online]. Available: <https://www.sciencedirect.com/science/article/pii/S0094576523003995>
- [2] E. Bates and S. D’Amico, “Removing ambiguities in concurrent monocular single-shot spacecraft shape and pose estimation using a deep neural network,” in *AAS/AIAA Astrodynamics Specialist Conference*, 2025, paper AAS 25-146.
- [3] B. Candan and S. Servadio, “Relative pose estimation of an uncooperative target with camera marker detection,” *Aerospace*, vol. 12, no. 5, 2025. [Online]. Available: <https://www.mdpi.com/2226-4310/12/5/425>
- [4] B. Chen, J. Cao, A. Parra, and T.-J. Chin, “Satellite pose estimation with deep landmark regression and nonlinear pose refinement,” in *ICCVW*, 2019.
- [5] J. I. B. Pérez-Villar, García-Martín, J. Bescós, and M. Escudero-Viñolo, “Spacecraft pose estimation: Robust 2-d and 3-d structural losses and unsupervised domain adaptation by intermodel consensus,” *IEEE Transactions on Aerospace and Electronic Systems*, vol. 60, no. 3, pp. 2515–2525, 2024.

- [6] O. Siméoni, H. V. Vo, M. Seitzer, F. Baldassarre, M. Oquab, C. Jose, V. Khalidov, M. Szafraniec, S. Yi, M. Ramamonjisoa, F. Massa, D. Haziza, L. Wehrstedt, J. Wang, T. Darcet, T. Moutakanni, L. Sentana, C. Roberts, A. Vedaldi, J. Tolan, J. Brandt, C. Couprie, J. Mairal, H. Jégou, P. Labatut, and P. Bojanowski, “DINOv3,” 2025. [Online]. Available: <https://arxiv.org/abs/2508.10104>
- [7] Y. Wang, Y. Sun, Z. Liu, S. E. Sarma, M. M. Bronstein, and J. M. Solomon, “Dynamic graph cnn for learning on point clouds,” *ACM Transactions on Graphics (TOG)*, 2019.
- [8] P.-E. Sarlin, D. DeTone, T. Malisiewicz, and A. Rabinovich, “SuperGlue: Learning feature matching with graph neural networks,” in *CVPR*, 2020. [Online]. Available: <https://arxiv.org/abs/1911.11763>
- [9] P. An, J. Yang, M. Peng, Y. Yang, Q. Liu, X. Wu, and L. Nan, “Mincd-pnp: Learning 2d-3d correspondences with approximate blind pnp,” in *Proceedings of the IEEE/CVF International Conference on Computer Vision (ICCV)*, 2025, pp. 26 519–26 528.
- [10] N. Mathihalli, A. Wei, G. Lavezzi, P. Mun Siew, V. Rodriguez-Fernandez, H. Urrutxua, and R. Linares, “Dreamsat: Towards a general 3d model for novel view synthesis of space objects,” in *75th International Astronautical Congress 2024*. Milan, Italy: International Astronautical Federation, 10 2024.
- [11] S. Diaz, X. Hu, J. Uwumukiza, G. Lavezzi, V. Rodriguez-Fernandez, and R. Linares, “Dreamsat-2.0: Towards a general single-view asteroid 3d reconstruction,” in *2025 AAS/AIAA Astrodynamics Specialist Conference*, Boston, MA, USA, 08 2025.
- [12] Stanford University, “SPE3R: Synthetic Dataset for Satellite Pose Estimation and 3D Reconstruction,” <https://purl.stanford.edu/pk719hm4806>.
- [13] T. H. Park and S. D’Amico, “Rapid abstraction of spacecraft 3d structure from single 2d image,” in *AIAA SCITECH 2024 Forum*, 2024.
- [14] B. Wen, W. Yang, J. Kautz, and S. Birchfield, “Foundationpose: Unified 6d pose estimation and tracking of novel objects,” in *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*, 2024.
- [15] V. Lepetit, F. Moreno-Noguer, and P. Fua, “Ep n p: An accurate o (n) solution to the p n p problem,” *International journal of computer vision*, vol. 81, no. 2, pp. 155–166, 2009.
- [16] T. H. Team, “Hunyuan3d 2.0: Scaling diffusion models for high resolution textured 3d assets generation,” 2025.
- [17] M. Cuturi, “Sinkhorn distances: Lightspeed computation of optimal transport,” in *Proceedings of the 26th International Conference on Advances in Neural Information Processing Systems (NIPS 26)*, 2013, pp. 2292–2300.
- [18] S. Li, K. Han, K.-Y. K. Wong, Y. Zheng, and M. Niessner, “Correspondence networks with adaptive neighbourhood consensus,” in *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*, 2020, pp. 10 196–10 205.
- [19] E. Brachmann and S. Gumhold, “Learning less is more - 6d object pose estimation from synthetic rgb images,” in *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*, 2018, pp. 4654–4663.
- [20] T. H. Park, S. Sharma, and S. D’Amico, “Towards robust learning-based pose estimation of non-cooperative spacecraft,” *arXiv preprint arXiv:1910.07139*, 2019, primary reference for the original SPEED dataset and the 2019 challenge.
- [21] M. Tyszkiewicz, P. Fua, and E. Trulls, “Disk: Learning local features with reinforcement learning,” in *Advances in Neural Information Processing Systems (NeurIPS)*, vol. 33, 2020, pp. 10 254–10 265.
- [22] A. Kirillov, E. Mintun, N. Ravi, H. Mao, C. Rolland, L. Gustafson, T. Xiao, S. Whitehead, A. C. Berg, W.-Y. Lo, P. Dollár, and R. Girshick, “Segment anything,” 2023. [Online]. Available: <https://arxiv.org/abs/2304.02643>