

Vision Foundation Models with Synthetic-Only Training for Monocular Spacecraft Pose Estimation

John Church* Vazghen Nikolian*

Abstract—We present an improvement on previous spacecraft pose estimation architectures that results in the lowest published mean rotation errors we know of on the SPEED+ lightbox and sunlamp test sets for a known, non-cooperative spacecraft. By using a previously established heatmap-based pose estimation architecture and adapting a large self-supervised ViT foundation model (DINOv3) in place of the smaller convolutional and ViT encoders of previous work, we show that pose estimation accuracy improves from 300M to 840M parameters with no saturation yet observed. We also evaluate our 840M model on a Jetson Orin NX 16GB, measuring single-pass network inference at 133.8 ms per crop with a board draw of 32.0 W. These measurements demonstrate embedded inference feasibility on a processor family with orbital flight heritage. Our resulting model outperforms previous models across lightbox and sunlamp domains while training only on synthetic data. Our best model, using DINOv3 840M adapted with LoRA as the encoder (rank 64, three-seed ensemble with four-rotation test-time augmentation), results in 1.56° mean rotation error on sunlamp and 1.17° on lightbox, compared to the previous best mean rotation errors we know of on these test sets, 2.66° and 1.75° by EagerNet.

I. INTRODUCTION

In the coming years the number of missions that require a rendezvous between two craft in orbit will likely increase considerably [1], [2], [3]. In the best cases, the target craft in question will be equipped with modern docking features that reduce the risk and complexity of the maneuver. Features like IR dots around docking ports and shared telemetry over stable connections between craft make pose estimation, the problem of an approaching craft knowing the target’s 6D pose (relative position and orientation), straightforward. However, many missions will need to be conducted that are not so easy. Targets that are legacy, natural (such as asteroids), broken, or otherwise uncooperative will nonetheless need to be approached and therefore their 6D pose will need to be accurately ascertained without pre-planned cooperative systems assisting.¹

An additional layer of difficulty is added when considering the constraints imposed by the particular environment the missions will take place in and the material limits of the craft that will be engaging in them. Imagery taken of the target craft is subject to harsh sun lighting that, without the aid of atmospheric diffusion, casts pitch black shadows and overexposes metallic surfaces in direct light. Earth albedo



Fig. 1. Example SPEED+ imagery: synthetic training render (left), lightbox HIL test image (center), sunlamp HIL test image (right). Imagery from SPEED+ [5], CC BY-NC-SA 4.0.

provides some soft fill illumination but cannot be relied upon. Active sensors that would alleviate these visual difficulties somewhat, such as LiDAR, add mass, power draw, and system complexity (the RVS 3000-3D rendezvous LiDAR used for the MEV-1 docking, for example, is specified at 12.4–15.3 kg and roughly 71 W nominal power [4]). While monocular pose estimation also incurs computational and energy costs, most methods can run on an embedded GPU without requiring a custom processor designed specifically for pose estimation.

For these reasons, pose estimation using greyscale monocular imagery has been an object of research interest [5]. Given the unique difficulty of obtaining real training imagery, the problem of solving pose estimation has gone hand in hand with the problem of bridging the gap between simulated and real training data. SPEED+ was developed as a benchmark by Park et al. [5], a collaboration between Stanford’s Space Rendezvous Laboratory (SLAB) and ESA’s Advanced Concepts Team, and followed by a competition in 2021–2022, the Satellite Pose Estimation Competition (SPEC2021), hosted on ESA’s Kelvins platform [6]. The SPEED+ dataset includes both simulated images of a satellite ‘Tango’ along with real images taken of a physical Tango model staged to simulate harsh direct sun (‘sunlamp’) as well as earth albedo (‘lightbox’). The real imagery comprises 9,531 hardware-in-the-loop (HIL) test images of the mockup in Stanford’s TRON facility alongside roughly 60,000 synthetic images (training and validation), examples of which are shown in Fig. 1.

Significant progress has been made in recent years in solving both the sim to real gap in this domain as well as the central problem of pose estimation itself. The competition was won with mean rotation errors of 2.73° on sunlamp (team lava1302) and 3.19° on lightbox (team TangoUnchained) [6], though the top entrants all used some of the real imagery during training. The competition postmortem further improved on the competition results with entrants such as FA-VAE [7]

*Both authors contributed equally to this work; author order is alphabetical.

¹Here, non-cooperative refers to the absence of assistance from the target spacecraft; our method still assumes that its geometry and the 3D locations of the selected keypoints are known.

and the semi-supervised domain adaptation of Chen et al. [8]. In more recent years, further refinements have pushed the edge further while showing that state of the art results are possible without ever using non-simulated data in training. Impressive results include those of Ulmer et al.’s EagerNet [9] whose technique also markedly improves on many previous methods by producing state of the art results while requiring only an approximate 3D model of the target. Park and D’Amico, the team behind SPNv3 [10], showed that it was possible to produce competitive results on a model of only 22.7M parameters (compared to EagerNet’s approximately 88.6M-parameter backbone). Concurrently with our research, Bergia et al. [11] also implemented a DINOv3 backbone with a heatmap head, obtaining their best result by self-training on unlabeled HIL imagery with DINOv3-L at 224×224 input. Our contribution uses low-rank adapters at 448×448 input, scales to the 840M encoder, and includes controls that isolate the source of the gains to encoder scale and adaptation. On the same DINOv3-L encoder, our single-pass synthetic-only model reaches $1.90^\circ / 2.26^\circ$ (lightbox / sunlamp) against their $3.27^\circ / 4.06^\circ$, a 42% / 44% reduction in mean rotation error, and our best configuration reaches $1.17^\circ / 1.56^\circ$ against their self-trained $2.57^\circ / 3.25^\circ$. Table I summarizes these prior results. Recent direct-regression approaches trained on SPEED+ synthetic imagery, such as FastPose-ViT [12], PAID-ViT [13], and Cov2Pose [14], report mean rotation errors above 11° on the HIL domains and are not included in Table I. Recent keypoint-based approaches include mask-guided attention and adaptive keypoint suppression by Peng et al. [15] and a depth-aware keypoint loss by Yang et al. [16]; neither study reports separate results on the full SPEED+ lightbox and sunlamp test sets.

Previous research has generally used small models compared to the large ViT foundation models increasingly used in other applications. Some models, such as SPNv3, were specifically designed around operating on limited compute. In the coming years, though, it is likely that onboard compute resources will become more common and less cost prohibitive [18]. We set out to answer whether scaling improves pose estimation and to measure the computational cost of the resulting model on embedded hardware. Our research shows that there is indeed significant headroom and that previous architectures, when adapted to use cutting edge SSL vision foundation models, produce state of the art mean rotation errors on the SPEED+ lightbox and sunlamp test sets that improve on all previously published results we know of for those domains.

II. METHOD

We began by implementing the SPNv3 architecture described by Park and D’Amico in “Bridging the Domain Gap for Flight-Ready Spaceborne Vision” [10]. SPNv3 is a three stage architecture involving:

- 1) A bounding box detection algorithm that crops the high resolution input image to produce a smaller image of

the target object suitable to be fed into the subsequent vision model.

- 2) Detection of K predefined keypoints on the image of the target
- 3) A pose solution solver that takes the 2D keypoint locations and produces a 6D pose solution using the well-established Perspective-n-Point (PnP) family of solvers [19].

For our purposes, we can more or less treat steps 1 and 3 of this architecture as fixed. Following SPNv3, we use ground-truth-derived crops in place of a bounding box detection algorithm for stage 1 so that our comparison isolates the keypoint detection task central to stage 2.² The pose solution solving math in step 3 is already well established.³ Remaining error in pose solutions for this architecture stem from inaccuracies in placing keypoints in stage 2.

In the case of SPNv3, stage 2 uses a neural net feeding heads that output K heatmaps for the keypoints ($K = 11$ for Tango). SPNv3 trained ViTPose, originally proposed by Xu et al. [20], for this purpose. For our implementation, we decided to tune DINOv3 [21], a self-supervised vision foundation model by Meta AI, in place of ViTPose’s ViT-S backbone. We trained the 300M variant using both LoRA [22] and DoRA [23], and the 840M variant using LoRA, across only synthetic imagery. As with previous implementations, we found that random seeding multiple models and predicting via consensus between an ensemble of 3 produced superior results for the 300M model. For the rank-64 840M model, the three independently trained models with four-rotation test-time augmentation (the crop rotated by 0° , 90° , 180° , and 270° , with the back-rotated heatmaps averaged using each view’s peak confidence as its weight) gave mean rotation errors of 1.19 – 1.20° on lightbox and 1.49 – 1.79° on sunlamp. Combining the three models with the same augmentation gave 1.17° and 1.56° , respectively. We report the single-model and ensemble configurations separately because they require four and twelve network passes per frame.

Our training, as mentioned, was purely on the official synthetic training split of SPEED+ (approximately 48,000 renders), with the 11,994-image synthetic validation split held out. Across all rank-16 training runs our configuration was held constant with the exception of the encoder itself and per-encoder input normalization matched to the encoder’s specific pre-training statistics. The rank-64 runs raise the adapter rank from 16 to 64 (with LoRA scaling α/r of 2 instead of 1), and the rank-64 840M runs additionally lower the learning rate to 3×10^{-4} for training stability. Network inputs were always the same 448×448 crops. The heatmap head was trained from scratch in every run (two transposed-convolution upsampling blocks and a 1×1 convolution producing 112×112 heatmaps,

²Our SPEED+ results therefore evaluate the keypoint-and-PnP stage given a ground-truth-derived crop, the same protocol as SPNv3. Comparisons with methods that include their own detector, such as EagerNet and the SPEC2021 entries, are not under an identical protocol.

³During training, crops are randomly enlarged and shifted to build robustness against imperfect bounding boxes, as in SPNv3.

TABLE I

SELECTED PRIOR PUBLISHED RESULTS ON SPEED+ HIL. * = USED REAL HIL IMAGERY DURING TRAINING; † = 3-SEED ENSEMBLE. VALUES ARE REPORTED UNDER EACH SOURCE’S OWN TRAINING AND CROP/DETECTION PROTOCOL AND ARE NOT A CONTROLLED COMPARISON UNDER ONE COMMON PIPELINE. BOLD DENOTES THE BEST DISPLAYED VALUE IN EACH METRIC COLUMN, INCLUDING ROUNDED TIES. OUR ROTATION-ERROR VALUES IN TABLE II ARE UNTHRESHOLDED MEANS; THE PRIOR VALUES IN THIS TABLE ARE REPRODUCED AS REPORTED AND MAY APPLY THE HIL PRECISION FLOOR.

Method	Year	Params	lightbox			sunlamp		
			$\bar{E}_t$	$E_q[^\circ]$	E_{pose}	$\bar{E}_t$	$E_q[^\circ]$	E_{pose}
TangoUnchained* [6]	2022	—	0.018	3.19	0.073	0.015	4.30	0.090
lava1302* [6]	2022	—	0.048	6.66	0.165	0.011	2.73	0.059
SPNv2 + ODR* [17]	2023	52.5M	0.025	5.58	0.122	0.026	9.79	0.197
Chen et al.* [8]	2024	—	0.018	2.87	0.068	0.014	2.75	0.062
EagerNet [9]	2023	88.6M	0.009	1.75	0.039	0.013	2.66	0.059
SPNv3-S [10]	2024	22.7M	0.016	3.04	0.069	0.020	4.37	0.097
SPNv3-S† [10]	2024	22.7M	0.017	2.69	0.064	0.020	3.88	0.088
SPNv3-B† [10]	2024	86.3M	0.012	2.03	0.047	0.015	3.40	0.074
Bergia et al. [11] (synthetic only, 40% holdout)	2026	303M	0.012	3.27	0.069	0.016	4.06	0.087
Bergia et al.* [11] (+ self-training, 40% holdout)	2026	303M	0.012	2.57	0.057	0.015	3.25	0.072
Bergia et al.* [11] (+ self-training, full set)	2026	303M	0.013	2.87	0.063	0.015	3.27	0.072

trained with a mean-squared-error loss against Gaussian targets). The original optimization schedule used AdamW with learning rate 10^{-3} , weight decay 10^{-4} , 30 epochs at batch size 32, and one warmup epoch followed by cosine decay; the rank-64 840M runs retain this schedule with the lower learning rate stated above. Thanks to modern fine-tuning methods such as LoRA and DoRA we were able to achieve our results while adding trainable parameters equivalent to roughly 0.5–3% of total encoder weights in each tuning run, with the original encoder weights remaining frozen; adapters are placed on the query/key/value and output projections of every attention block. Since SPEED+ does not provide a real-image validation split, we follow the standards set in prior work and use the lightbox and sunlamp sets both to compare settings and to choose inference-time settings. No real images were used for training. As an additional check on our results, we used only the even-numbered frames of each real set to choose the inference-time settings and then scored the odd-numbered frames with the configuration selected that way, a single model with four-rotation test-time augmentation. This resulted in mean rotation errors of 1.22° on lightbox and 1.52° on sunlamp, consistent with the full-set results.

III. CONTROLS

We conducted a number of control runs in an effort to confirm the gains we achieved were due to tuning a larger encoder and not a difference in our architecture or improvements generally in base encoders (Table III):

- 1) Since we followed SPNv3 closely architecturally, we conducted our own training run using a similar baseline encoder (22M DeiT3 [24] ViT-S) and obtained results in the same range.
- 2) To test whether modern ViT encoders can support pose estimation without fine tuning, we attempted to train only the head while using DINOv3 and SigLIP2-L [25]

as encoders in their own respective runs. Each failed to produce usable results when the encoder was frozen for training (DINOv3: 18% / 29% catastrophic flips on lightbox / sunlamp; SigLIP2-L: 54% / 70%).

- 3) EUPE-S [26], a modern 22M parameter encoder, was trained with LoRA to test whether modern encoders still produced similar results to similarly sized encoders used in previous models like SPNv3. EUPE-S with LoRA actually ended up producing significantly worse results compared to our fully trained 22M DeiT3 [24] ViT-S baseline, showing that using modern encoders and fine-tuning techniques alone does not improve on previous results.
- 4) We trained SigLIP2-L using LoRA to confirm that other modern large pretrained ViT models show similar improvements with tuning and that our results are not unique to DINOv3’s architecture. DINOv3 300M does seem to be a better model for this application, as SigLIP2-L (316M) remained 18–23% behind DINOv3-L (303M) with LoRA (1.90° / 2.26° vs 2.47° / 2.74°).

IV. HARDWARE TESTING

While not a primary focus of this paper, we did want to benchmark the performance of our system on hardware that has already flown real missions as a preliminary test to see if computation of this size is viable. This is especially important since many past approaches, including the SPNv3 approach on which we based our model, cite computation constraints as a reason for opting for a smaller model.

Bergia et al. [11] also evaluated DINOv3-L on a Jetson Orin NX at 224×224 input using TensorRT FP16. Our measurements characterize the larger 840M model at 448×448 input, including measured board power.

We chose to test on an NVIDIA Jetson Orin NX 16GB, a 70×45 mm system-on-module with a 1,792-core Ampere

TABLE II

OUR RESULTS. SELECT STATE OF THE ART RESULTS ARE RENDERED IN GREY ITALIC ROWS FOR COMPARISON. FLIP% = FRACTION OF FRAMES WITH ROTATION ERROR ABOVE 30° , I.E. A CATASTROPHIC ORIENTATION ERROR; ON THIS NEARLY SYMMETRIC TARGET THESE ARE FREQUENTLY THE MIRRORED POSE. SINGLE-MODEL ROWS ARE SEED 2026, THE FIRST TRAINED; THE THREE-SEED SUMMARIES AVERAGE SEEDS 2026, 2027, AND 2028 FOR THE INDICATED ONE-PASS OR FOUR-PASS CONFIGURATION, WITH SAMPLE STANDARD DEVIATION REPORTED FOR ROTATION ERROR. TTA = CONFIDENCE-WEIGHTED AVERAGE OF 4 ROTATED PASSES. $\ddagger$ = LIGHTBOX TRANSLATION MEANS ARE DOMINATED BY A SINGLE PHYSICALLY INVALID PNP SOLUTION WITH A TRANSLATION ERROR OF 500 M OR MORE. E_{pose} FOLLOWS THE REVISED SPEED+ DEFINITION USED IN SPNV3.

Model	Params	Passes	lightbox				sunlamp			
			$\bar{E}_t$	$E_q[^\circ]$	flip%	E_{pose}	$\bar{E}_t$	$E_q[^\circ]$	flip%	E_{pose}
<i>EagerNet [9] (reference)</i>	<i>88.6M</i>	<i>4</i>	<i>0.009</i>	<i>1.75</i>	—	<i>0.039</i>	<i>0.013</i>	<i>2.66</i>	—	<i>0.059</i>
<i>SPNV3-B$\ddagger$ [10] (reference)</i>	<i>86.3M</i>	<i>3</i>	<i>0.012</i>	<i>2.03</i>	—	<i>0.047</i>	<i>0.015</i>	<i>3.40</i>	—	<i>0.074</i>
DINOv3-L, LoRA r16	303M	1	0.0099	1.90	0.43	0.0428	0.0117	2.26	0.47	0.0510
L r16 + 3-seed ensemble	303M	3	0.0096	1.69	0.39	0.0387	0.0115	2.17	0.36	0.0492
L r16 + 4-rot TTA	303M	4	0.0124	1.72	0.49	0.0421	0.0139	2.01	0.32	0.0488
L r16 + ensemble + TTA	303M	12	$\ddagger$	1.56	0.31	$\ddagger$	0.0122	2.00	0.32	0.0470
DINOv3-L, DoRA r16	303M	1	$\ddagger$	1.79	0.33	$\ddagger$	0.0136	2.35	0.36	0.0544
L DoRA r16 + 4-rot TTA	303M	4	0.0097	1.59	0.30	0.0371	0.0124	2.15	0.39	0.0497
DINOv3-H+, LoRA r16	840M	1	0.0082	1.34	0.13	0.0312	0.0108	1.84	0.21	0.0428
H+ r16, 3-seed mean $\pm$ SD	840M	1	0.0082	1.41 $\pm$ 0.07	0.25	0.0326	0.0109	1.87 $\pm$ 0.06	0.27	0.0434
H+ r16 + 3-seed ensemble	840M	3	0.0085	1.31	0.19	0.0309	0.0102	1.72	0.21	0.0400
H+ r16 + 4-rot TTA	840M	4	0.0076	1.20	0.12	0.0281	0.0112	1.67	0.18	0.0401
H+ r16 + ensemble + TTA	840M	12	0.0075	1.21	0.13	0.0283	0.0096	1.68	0.21	0.0388
DINOv3-L, LoRA r64	303M	1	$\ddagger$	1.46	0.19	$\ddagger$	0.0117	1.94	0.32	0.0454
L r64 + 4-rot TTA	303M	4	0.0098	1.35	0.18	0.0330	0.0121	1.71	0.21	0.0418
DINOv3-H+, LoRA r64	840M	1	0.0085	1.41	0.19	0.0327	0.0099	1.77	0.29	0.0407
H+ r64, 3-seed mean $\pm$ SD	840M	1	0.0084	1.37 $\pm$ 0.04	0.21	0.0319	0.0099	1.82 $\pm$ 0.06	0.29	0.0414
H+ r64 + 3-seed ensemble	840M	3	0.0076	1.22	0.10	0.0286	0.0095	1.74	0.29	0.0398
H+ r64 + 4-rot TTA (seed 2026)	840M	4	0.0077	1.20	0.09	0.0283	0.0105	1.49	0.14	0.0363
H+ r64 + TTA, 3-seed mean $\pm$ SD	840M	4	0.0075	1.197 $\pm$ 0.004	0.11	0.0281	0.0117	1.67 $\pm$ 0.16	0.24	0.0407
H+ r64 + ensemble + TTA	840M	12	0.0073	1.17	0.10	0.0273	0.0113	1.56	0.21	0.0383

GPU, 16 GB of shared memory, and a configurable 10–40 W power envelope [27] from a module family flown in 2026 on Aethero’s Phobos [28] and EDGX’s STERNA [29] payloads. Our preliminary findings are positive, showing that our 840M model is able to run at 133.8 ms per single network pass at 32.0 W measured board draw, and 419.8 ms at 10.8 W; four-view batches take 501.9 ms and 1.63 s at those respective operating points.

It should be noted that the Orin, while small enough to be suitable for some missions, is a heavier module, in power and thermal terms, than is often needed or flown in many cases (its rated envelope is 10–40 W, and we measured 10.8–32 W of total board draw). Additionally, we have not tested the capabilities of the chip under radiation or in a flight-representative thermal environment. Given these limitations, these results are informative for assessing viability in cases where a mission is already being flown with sufficient compute on board. More testing is needed for true flight-ready assessment and for viability on smaller, more energy-constrained craft.




Fig. 2. Example frames which fail for two training seeds of the 300M DINOv3-L + LoRA model and for the 22M baseline: near-black lightbox frames (left) and glare-saturated sunlamp frames (right). Imagery from SPEED+ [5], CC BY-NC-SA 4.0.

V. RESULTS

Our rank-16 300M model in a single pass (1.90° lightbox / 2.26° sunlamp) achieves lower mean rotation errors than every previously published single-pass model result we know of on the full lightbox and sunlamp test sets (best prior single model, SPNV3-B: 2.18° / 3.61°), and with 4-rotation test-time augmentation (1.72° / 2.01°) it also exceeds the next best result, EagerNet, with four network passes, matching EagerNet’s rotation count (1.75° / 2.66°). Our rank-16 840M model in a single pass (1.34° / 1.84°) exceeds every rank-16 multi-pass configuration of the 300M model and all previously published mean rotation-error results we know of on the full lightbox and sunlamp test sets, with or without ensemble or

TABLE III

CONTROLS RESULTS. mean° = MEAN ROTATION ERROR OVER ALL FRAMES; bulk° = MEAN OVER FRAMES WITH ERROR AT OR BELOW 30° , I.E. THE CONDITIONAL MEAN ON NON-FLIPPED FRAMES; $\text{flip}\%$ AS IN TABLE II. ALL CONTROL ROWS USE SEED 2026 AND ONE NETWORK PASS PER CROP; THE LoRA CONTROLS USE RANK 16.

Model	Params	Trained	lightbox			sunlamp		
			mean°	bulk°	$\text{flip}\%$	mean°	bulk°	$\text{flip}\%$
Baseline (SPNv3-style)	22M	all	3.60	2.20	1.5	5.19	3.21	2.1
DINOv3-L frozen	303M	head	15.47	3.65	18.1	24.49	5.01	29.2
SigLIP2-L frozen	316M	head	44.37	7.35	53.9	58.92	10.18	70.4
EUPE-S + LoRA	22M	head+LoRA	7.99	3.32	5.9	8.93	3.83	6.4
SigLIP2-L + LoRA	316M	head+LoRA	2.47	1.83	0.7	2.74	2.19	0.5

TABLE IV

EMBEDDED NETWORK BENCHMARK ON JETSON ORIN NX 16GB, DINOv3-H+ WITH FOLDED RANK-64 LoRA, SEED 2026, TENSORRT BF16, 448×448 INPUTS. LATENCIES ARE MEDIAN GPU COMPUTE TIMES OVER 200 TIMED EXECUTIONS AFTER WARM-UP. BATCH 1 IS ONE NETWORK PASS; BATCH 4 MEASURES A BATCH OF FOUR VIEWS, THE NETWORK WORKLOAD FOR SINGLE-MODEL ROTATION TTA. BATCH-4 LATENCY IS PER BATCH, NOT PER VIEW. THESE TIMINGS EXCLUDE ACQUISITION, CROP PREPARATION, TRANSFERS, TTA AGGREGATION, AND POSE SOLVING. TIMING RUNS USE RANDOM INPUT TENSORS; AGREEMENT WITH THE PYTORCH EVALUATION WAS CHECKED SEPARATELY ON REAL CROPS. POWER IS MEASURED TOTAL BOARD INPUT DRAW, DISTINCT FROM THE NAMED POWER MODE; RAM IS OBSERVED PEAK SYSTEM MEMORY, INCLUDING RUNTIME OVERHEAD.

Power mode	latency [ms]		board draw [W]		peak RAM [GB]
	Batch 1	Batch 4	idle	load	
MAXN_SUPER	133.8	501.9	7.8	32.0	5.4
40 W	133.1	502.3	6.7	32.4	6.2
25 W	360.6	1343	6.8	13.4	6.6
15 W	410.5	1597	6.7	12.7	5.3
10 W	419.8	1630	4.9	10.8	6.0




Fig. 3. Percent of the retained real frames the model incorrectly flipped, at various rates of least-confident frame abstention, for the 300M DINOv3-L + LoRA model.

test-time augmentation (Table II).

With rank-64 adaptation, the 840M three-seed ensemble and four-rotation test-time augmentation gives 1.17° mean rotation error on lightbox and 1.56° on sunlamp, with catastrophic-flip rates of 0.10% and 0.21% and pose scores of 0.0273 and 0.0383, respectively. These results are computed over

the full test sets without abstention. This configuration uses twelve network passes per frame; the Jetson measurements separately characterize one-pass and four-pass inference for a single model.

For the rank-64 840M ensemble with four-rotation test-time augmentation, median translation error is 2.5 cm on lightbox and 3.4 cm on sunlamp, at target distances of 2.5–9.6 m. The ensemble fails on 13 of the 9,531 real frames, and one frame in each domain fails for every training seed. These are often near-black or glare-saturated (Fig. 2). Many of these failures are flips that result in an extreme incorrect pose estimation in excess of 30 degrees. These flips often happen when the model mirrors the locations of the keypoints on the largely symmetric Tango satellite. A range check on the solved distance would remove the $\ddagger$ cases in Table II, but it catches only one flip per domain while the other flips solve to a plausible distance.

For the original rank-16 single-pass models, we observed that the model’s confidence was low in most cases when catastrophic flips occurred. The 300M model showed a median per-frame confidence (the mean of its K per-keypoint heatmap peak values) of 0.90 on correct frames vs 0.39 on flipped frames in the lightbox domain and 0.88 vs 0.41 in the sunlamp domain; the 840M model gives 0.92 vs 0.31 and 0.91 vs 0.39 on the respective domains. When broken down by degrees of error with and without flips included, our 300M model has a mean rotation error of $1.90^\circ / 2.26^\circ$ (lightbox / sunlamp) over all frames but $1.51^\circ / 1.82^\circ$ over the non-flipped frames alone; the 840M model $1.34^\circ / 1.84^\circ$ versus $1.23^\circ / 1.52^\circ$. As can be seen in Fig. 3, the model generally gives much lower confidence to flipped predictions. By retrospectively withholding the 2% least-confident frames, the 300M model’s flip rate falls from 0.43% / 0.47% to 0.09% / 0.07% and its mean rotation error to $1.46^\circ / 1.81^\circ$. It is possible to remove every flip in both real domains by withholding 10% of frames which would lower the mean error to $1.24^\circ / 1.62^\circ$; the 840M model needs only 2% withheld to reach zero flips in both domains (from 0.13% / 0.21%, against 1.5% / 2.1% for our fully trained 22M baseline), lowering its mean error from $1.34^\circ / 1.84^\circ$ to $1.16^\circ / 1.47^\circ$, at the cost of introducing occasional blank responses between prediction frames. Our results tables do not employ this technique in order to remain comparable to previous methods, which report over the full

test sets.

VI. CONCLUSION

This paper shows that adapting large self-supervised ViT models such as DINOv3 improves on the results of previous methods that use relatively smaller encoders while holding other architecture elements constant. We find that these large encoders do not perform well when frozen but that modern fine tuning methods like LoRA and DoRA allow for adaptation of a small percent of encoder weights to achieve state of the art capabilities in this domain. This result suggests that future missions that require onboard pose estimation of a target craft, and that have available compute for a model of this size (on the Jetson Orin NX measured here, one network pass in 134–420 ms at 11–32 W of board draw), have a more accurate option than the previous state of the art. Additionally, our rank-16 results show that scaling from 300M to 840M cuts mean rotation error by roughly 15–30% and roughly halves catastrophic flips, suggesting there may still be headroom to improve results further with larger foundation models. Our model was only trained and tested on a single target satellite with a well-defined 3D model and predefined keypoints, a stronger geometric prior than approximate-model methods such as EagerNet require, and limited to predictions from single frames.

REFERENCES

- [1] J. Mulvaney, D. Arney, C. Williams *et al.*, “In-space servicing, assembly, and manufacturing (ISAM) state of play – 2025 edition,” NASA Langley Research Center, Tech. Rep., Oct. 2025, nTRS 20250008988, <https://ntrs.nasa.gov/citations/20250008988>.
- [2] Secure World Foundation, “Commercial and civil robotic rendezvous and proximity operations,” Secure World Foundation, Fact Sheet FS25/3, Jun. 2025. [Online]. Available: https://cdn.prod.website-files.com/66dccc6872f6ed23bce1db235/684ad842a9a1efe9b1bbb5dd_Fact%20Sheet%202025_Commercial_Civil%20RPO.pdf
- [3] Novaspace, “In-orbit services market to generate \$3 billion in revenue over the next decade,” Press release, Aug. 2026. [Online]. Available: <https://nova.space/press-release/in-orbit-services-market-to-generate-3-billion-in-revenue-over-the-next-decade/>
- [4] Jena-Optronik GmbH, “RVS 3000 product family data sheet,” https://www.jena-optronik.de/products/rendezvous-sensors/rvs-3000.html?file=tl_files%2Fpdf%2FData+Sheet+RVS+3000+Product+Family.pdf, accessed September 13, 2026.
- [5] 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)*, Big Sky, MT, USA, 2022, pp. 1–15.
- [6] T. H. Park, M. Märtens, M. Jawaid *et al.*, “Satellite pose estimation competition 2021: Results and analyses,” *Acta Astronautica*, vol. 204, pp. 640–665, 2023.
- [7] Y. Liu, R. Zhou, D. Du, S. Cao, and N. Qi, “Feature-aided pose estimation approach based on variational auto-encoder structure for spacecrafts,” *Chinese Journal of Aeronautics*, vol. 37, no. 8, pp. 329–341, 2024.
- [8] S. Chen, W. Yang, W. Wang, J. Mai, J. Liang, and X. Zhang, “Spacecraft homography pose estimation with single-stage deep convolutional neural network,” *Sensors*, vol. 24, no. 6, p. 1828, 2024.
- [9] M. Ulmer, M. Durner, M. Sundermeyer, M. Stoiber, and R. Triebel, “6D object pose estimation from approximate 3D models for orbital robotics,” in *2023 IEEE/RSJ International Conference on Intelligent Robots and Systems (IROS)*, Detroit, MI, USA, 2023, pp. 10 749–10 756.
- [10] T. H. Park and S. D’Amico, “Bridging the domain gap for flight-ready spaceborne vision,” *Journal of Spacecraft and Rockets*, vol. 63, no. 1, pp. 289–307, 2026.
- [11] S. Bergia, I. Caracciolo, F. Stesina, and M. Fantino, “Self-supervised vision transformers for spacecraft pose estimation across domain shift,” in *2026 IEEE 13th International Workshop on Metrology for AeroSpace (MetroAeroSpace)*, 2026, pp. 70–75.
- [12] P. Ancey, A. Price, S. Javed, and M. Salzmann, “FastPose-ViT: A vision transformer for real-time spacecraft pose estimation,” in *IEEE/CVF Winter Conference on Applications of Computer Vision (WACV)*, 2026, arXiv:2512.09792.
- [13] Z. Xie, Y. Yang, Y. Zhang, G. Xie, Y. Liu, and S. Zhang, “Precision-aware illumination-disentangled vision transformer for spacecraft 6D pose estimation,” *arXiv preprint arXiv:2606.11619*, 2026.
- [14] N. Ali Ousalah, P. Rostami, V. Gaudillière *et al.*, “Cov2Pose: Leveraging spatial covariance for direct manifold-aware 6-DoF object pose estimation,” in *IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*, 2026, arXiv:2603.19961.
- [15] L. Peng, R. Jin, S. Gao *et al.*, “Precision-aware adaptive suboptimality suppression for pose recovery of non-cooperative on-orbit targets,” *Advances in Space Research*, vol. 78, no. 2, pp. 1627–1643, 2026.
- [16] H. Yang, J. Huang, H. Ren, H. Sun, and R. Tian, “Six degrees of freedom pose estimation of non-cooperative space object based on approximate three dimensional keypoint similarity loss,” *Engineering Applications of Artificial Intelligence*, vol. 178, p. 115074, 2026.
- [17] T. H. Park and S. D’Amico, “Robust multi-task learning and online refinement for spacecraft pose estimation across domain gap,” *Advances in Space Research*, vol. 73, no. 11, pp. 5726–5740, 2024.
- [18] NASA, “High-performance spaceflight computing (HPSC),” <https://www.nasa.gov/game-changing-development-projects/high-performance-spaceflight-computing-hpsc/>.
- [19] V. Lepetit, F. Moreno-Noguer, and P. Fua, “EPnP: An accurate O(n) solution to the PnP problem,” *International Journal of Computer Vision*, vol. 81, no. 2, pp. 155–166, 2009.
- [20] Y. Xu, J. Zhang, Q. Zhang, and D. Tao, “ViTPose: Simple vision transformer baselines for human pose estimation,” in *Advances in Neural Information Processing Systems (NeurIPS)*, vol. 35, 2022, pp. 38 571–38 584.
- [21] O. Siméoni, H. V. Vo, M. Seitzer *et al.*, “DINOv3,” *arXiv preprint arXiv:2508.10104*, 2025.
- [22] E. J. Hu, Y. Shen, P. Wallis *et al.*, “LoRA: Low-rank adaptation of large language models,” in *International Conference on Learning Representations (ICLR)*, 2022.
- [23] S.-Y. Liu, C.-Y. Wang, H. Yin *et al.*, “DoRA: Weight-decomposed low-rank adaptation,” in *Proceedings of the 41st International Conference on Machine Learning (ICML)*, ser. Proceedings of Machine Learning Research, vol. 235. PMLR, 2024, pp. 32 100–32 121.
- [24] H. Touvron, M. Cord, and H. Jégou, “DeiT III: Revenge of the ViT,” in *Computer Vision – ECCV 2022*. Springer, 2022, pp. 516–533.
- [25] M. Tschannen, A. Gritsenko, X. Wang *et al.*, “SigLIP 2: Multilingual vision-language encoders with improved semantic understanding, localization, and dense features,” *arXiv preprint arXiv:2502.14786*, 2025.
- [26] C. Zhu, S. Suri, C. Jose *et al.*, “Efficient universal perception encoder,” *arXiv preprint arXiv:2603.22387*, 2026.
- [27] NVIDIA Corporation, “Jetson Orin modules: Technical specifications,” <https://www.nvidia.com/en-us/autonomous-machines/embedded-systems/jetson-orin/>, accessed September 13, 2026.
- [28] Aethero, “Phobos satellite launch (press release),” <https://www.aethero.com/news/press-release-phobos-satellite-launch/>, Mar. 2026, accessed September 13, 2026.
- [29] NVIDIA Corporation, “Space computing: On-orbit AI and accelerated computing,” <https://www.nvidia.com/en-us/edge-computing/space-computing/>, accessed September 13, 2026.

Author information

John Church

johnrobertchurch@protonmail.com
<https://orcid.org/0000-0002-9472-8101>

Vazghen Nikolian

nikolainvazghen@gmail.com
<https://orcid.org/0009-0003-6337-3496>