Equivariant Neural Inertial Odometry for Microgravity Space Robotics

Sejun Ahn[1], Yuheng Qiu[2], Chanhee Park[1], Brian Coltin[3], Ryan Soussan[3], and Pyojin Kim[1,*]

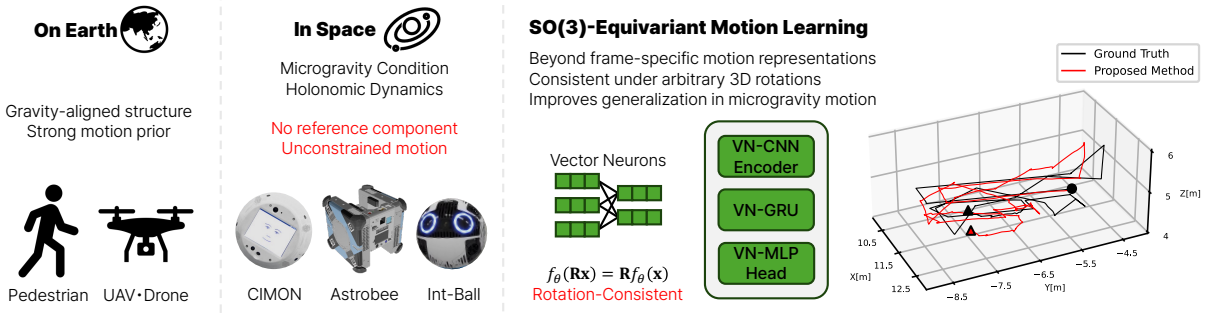

Fig. 1. Comparison between terrestrial and microgravity motion conditions and overview of the proposed SO(3)-equivariant neural inertial odometry for space free-flying robots [1], [2], [3]. On Earth, gravity provides a dominant directional prior and constrains motion, whereas in microgravity, motion is fully unconstrained with arbitrary 3D rotations. The proposed model leverages SO(3)-equivariant motion learning with vector neurons to ensure rotation-consistent predictions in the architectural-level and improved generalization in microgravity motions.

***Abstract*— We propose an SO(3)-equivariant neural inertial odometry framework for microgravity free-flying robots, where the absence of gravity removes dominant directional priors and leads to unconstrained 3D motion. Existing neural inertial odometry methods are largely developed for gravity-dominant terrestrial settings and do not transfer naturally to microgravity environments, where no prior method has been specifically designed. To address these limitations, we encode rotational symmetry directly in the motion network using vector neurons and prepend a learning-based IMU correction module to suppress axis-dependent sensor artifacts that violate the equivariance assumption. Our model predicts displacement and covariance in a rotation-consistent manner for translation estimation under known orientation, and improves both accuracy and generalization on the Astrobee simulation dataset.**

I. Introduction

Astrobee [1] is a free-flying robot that operates inside the International Space Station (ISS), where reliable onboard navigation remains challenging because of visual degradation, including lighting variation, clutter, and repetitive structure [4]. Although IMUs are unaffected by these visual conditions, classical inertial navigation rapidly drifts due to sensor noise and bias.

Recent neural inertial odometry (NIO) methods learn motion priors directly from IMU data, but their success has been demonstrated primarily in terrestrial settings with strong motion regularities. Astrobee, in contrast, operates in microgravity, where gravity no longer provides a dominant directional cue for inertial learning. Its holonomic and unconstrained 3D motion weakens the physical priors that terrestrial NIO methods typically exploit, making generalization to unseen motions substantially more difficult.

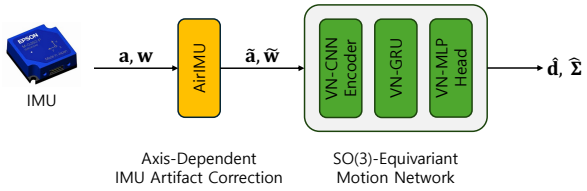

Fig. 2. Overview of the proposed SO(3)-equivariant NIO pipeline.

The main challenge, therefore, is to model IMU motion under full 3D rotational symmetry. Existing strategies such as coordinate normalization, yaw-only augmentation [5], [6], and soft equivariance [7], [8] are tailored to gravity-dominant settings (Fig. 1), and no prior NIO method directly addresses microgravity.

To address this limitation, we present, to the best of our knowledge, the first NIO framework designed for microgravity, based on vector neurons [9]. To better satisfy the symmetry assumption in the presence of real IMU artifacts, we prepend the motion network with a learning-based IMU correction module [10]. The resulting framework predicts displacement and uncertainty in a rotation-consistent manner. We focus on translation estimation under known orientation and evaluate whether architectural SO(3)-equivariance provides a better inductive bias for microgravity IMU-only learning than frame-dependent formulations developed for terrestrial motion.

II. Methodology

Fig. 2 summarizes the proposed two-stage pipeline: an axis-dependent IMU correction module that suppresses non-equivariant artifacts, followed by an SO(3)-equivariant motion network that predicts rotation-consistent displacement and uncertainty. By encoding rotational symmetry directly in the architecture, the proposed model promotes stronger generalization to unseen input orientations.

*Corresponding author.

[1]Sejun Ahn, Chanhee Park, and Pyojin Kim are with the Department of Mechanical and Robotics Engineering, Gwangju Institute of Science and Technology (GIST), Gwangju, Republic of Korea. {ahnsejun, chanipark}@gm.gist.ac.kr, pjinkim@gist.ac.kr

[2]Yuheng Qiu is with the Robotics Institute, Carnegie Mellon University, Pittsburgh, PA 15213, USA. yuhengq@andrew.cmu.edu

[3]Brian Coltin and Ryan Soussan are with the Intelligent Robotics Group (KBR Inc.), NASA Ames Research Center, Moffett Field, CA 94035, USA. {brian.coltin, ryan.soussan}@nasa.gov

A. IMU Correction

We first remove non-equivariant artifacts from the raw IMU measurements. In practice, raw IMU signals contain axis-dependent bias and noise, which break the rotational consistency assumed by an SO(3)-equivariant model. To mitigate this issue, we adopt the learning-based IMU correction strategy of AirIMU [10] before the motion network.

Given raw IMU measurements ${}^{B}\mathbf{a}$ and ${}^{B}\boldsymbol{\omega}$, AirIMU predicts correction terms ${}^{B}\sigma_{acc}$ and ${}^{B}\sigma_{gyro}$. The corrected measurements are defined as

$$ {}^{B}\tilde{\mathbf{a}} = {}^{B}\mathbf{a} + {}^{B}\sigma_{acc}, \qquad {}^{B}\tilde{\boldsymbol{\omega}} = {}^{B}\boldsymbol{\omega} + {}^{B}\sigma_{gyro}. \tag{1} $$

B. SO(3)-Equivariant Motion Network

We map the corrected IMU window to displacement and uncertainty using an SO(3)-equivariant motion network built on vector neurons (VNs) [9]. VNs represent latent activations as 3D vectors, allowing rotations applied to the input to propagate directly through the latent space. As a result, linear and nonlinear operations can be constructed to preserve SO(3)-equivariance throughout the network.

The proposed motion network consists of a VN-CNN encoder that extracts local motion patterns, followed by stacked VN-GRU layers that model temporal dependencies. Two VN-MLP decoder heads predict the displacement $\hat{\mathbf{d}} \in \mathbb{R}^3$ and its corresponding uncertainty $\hat{\boldsymbol{\Sigma}} \in \mathbb{R}^{3\times3}$. Because all latent operations are built on VNs, the network preserves SO(3)-equivariance and therefore produces rotation-consistent displacement and covariance estimates.

Formally, for an arbitrary rotation $R \in \mathrm{SO}(3)$ and corrected IMU input $\mathcal{X}$, the predictions satisfy

$$ \hat{\mathbf{d}}(R\mathcal{X}) = R\hat{\mathbf{d}}(\mathcal{X}), \tag{2} $$

$$ \hat{\boldsymbol{\Sigma}}(R\mathcal{X}) = R\hat{\boldsymbol{\Sigma}}(\mathcal{X})R^{\top}. \tag{3} $$

C. Loss Function

We train the motion network with a displacement regression loss and an uncertainty-aware probabilistic loss:

$$ \mathcal{L} = \mathcal{L}_{\text{Huber}} + \lambda\mathcal{L}_{\text{NLL}}. \tag{4} $$

$$ \mathcal{L}_{\text{Huber}} = \begin{cases} \frac{1}{2}\|\mathbf{d}-\hat{\mathbf{d}}\|^2, & \text{if } \|\mathbf{d}-\hat{\mathbf{d}}\| < \delta, \\ \delta\left(\|\mathbf{d}-\hat{\mathbf{d}}\| - \frac{1}{2}\delta\right), & \text{otherwise}, \end{cases} \tag{5} $$

$$ \mathcal{L}_{\text{NLL}} = \frac{1}{2}(\mathbf{d}-\hat{\mathbf{d}})^{\top}\hat{\boldsymbol{\Sigma}}^{-1}(\mathbf{d}-\hat{\mathbf{d}}) + \frac{1}{2}\log\left|\hat{\boldsymbol{\Sigma}}\right|. \tag{6} $$

Here, $\mathbf{d}$ and $\hat{\mathbf{d}}$ denote the ground-truth and predicted displacements, respectively. The final objective is a weighted sum of the two terms.

III. Evaluation

Dataset. We evaluate the proposed method on a simulation dataset generated using the Astrobee Robot Software [11] in Gazebo. The dataset contains 41 sequences with a total duration of 315 minutes, recorded at 62.5 Hz for ground-truth states and 125 Hz for IMU measurements. We partition the sequences into seen and unseen splits following [12].

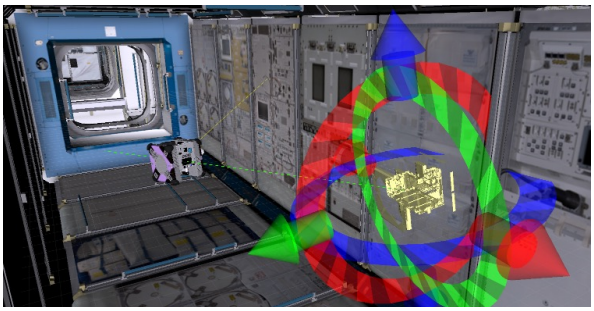

Fig. 3. Astrobee in the Astrobee Robot Software simulator. We collect sequences using user commands and scripted flight plans.

TABLE I

Summary of compared neural inertial odometry baselines.

Method	Input	Target	Representation
RoNIN	$\mathbf{a}$, $\boldsymbol{\omega}$	$\mathbf{v}$	World
TLIO	$\mathbf{a}$, $\boldsymbol{\omega}$	$\mathbf{d}$, Σ	World
IMO	$\boldsymbol{\omega}$, $\mathbf{T}_{\text{thrust}}$	$\mathbf{d}$, Σ	World
AirIO	$\mathbf{a}$, $\boldsymbol{\omega}$, $\boldsymbol{\xi}$	$\mathbf{v}$, Σ	Body
Ours	$\mathbf{a}$, $\boldsymbol{\omega}$	$\mathbf{d}$, Σ	SO(3)-equivariant (Body)

Baselines. We compare against representative NIO baselines summarized in Table I. RoNIN [5], TLIO [6], and IMO [13] are adapted to the ISS world frame; for IMO, the thrust input is replaced with the commanded 3D force from the Astrobee simulator. AirIO [12] retains its body-frame formulation with additional attitude encoding. We evaluate all methods in a network-only setting with ground-truth orientation and no downstream filtering.

Evaluation Metrics. We report RMSE ATE and RMSE RTE following [12], where ATE evaluates the full-trajectory translation error and RTE evaluates the relative translation error over fixed windows. The RTE interval is set to 5 seconds.

Training Details. We train the model using Adam with an initial learning rate of 0.001. Each training sample uses a window of 200 frames, corresponding to 2 seconds of resampled IMU input at 100 Hz. In our setup, we use a single NVIDIA RTX 4090 GPU with a batch size of 256, apply early stopping with a patience of 10 epochs.

Results. Table II shows that the proposed method achieves the lowest ATE and RTE on both the seen and unseen splits, with the largest margin on the unseen split. These results show that the SO(3)-equivariant formulation improves not only accuracy on seen trajectories but also generalization to unseen microgravity motion patterns.

Figs. 4 and 5 further support this result on representative unseen sequences with distinct motion characteristics, including long-range planar translation and repeated out-of-plane directional changes. In both cases, the proposed method remains more stable than the baselines, indicating stronger robustness across different unseen motion regimes.

TABLE II

ATE AND RTE (M) ON THE ASTROBEE SIMULATION DATASET.

Sequence	Metric	RoNIN	TLIO	AirIO	IMO	Ours
Seen	ATE	2.3639	2.2911	3.0095	2.7928	**0.8335**
	RTE	0.2560	0.2424	0.2756	0.3012	**0.0876**
Unseen	ATE	2.3859	2.5291	3.2572	2.9146	**1.1413**
	RTE	0.2177	0.2139	0.2590	0.2687	**0.1204**

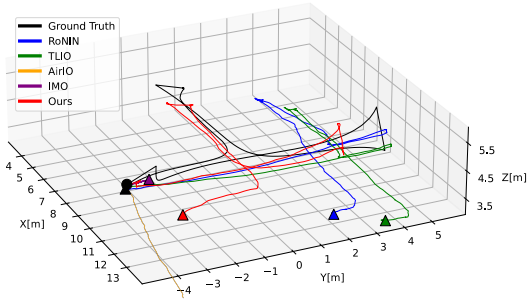

Fig. 4. Qualitative comparison on an unseen sequence with long-range translational motion primarily in the XY plane. The results reveal drift accumulation and trajectory stability under a long-range XY-planar motion pattern.

IV. ABLATION STUDY

This ablation isolates representation-dependent generalization failure in existing NIO baselines. In particular, methods trained with conventional world-frame or body-frame targets can fail under specific unseen motion conditions, depending on how the test sequence shifts relative to the training distribution in each representation. We therefore construct controlled cases to expose these failure modes and to test whether the proposed SO(3)-equivariant model remains robust when frame-specific baselines break down.

We define three controlled cases to analyze representation-dependent generalization. Case A is unseen in both world and body frames; Case B isolates world-frame distribution shift while keeping body-frame motion relatively consistent due to face-forwarding; and Case C keeps the world-frame trajectory similar while varying the body-frame motion.

Segment-wise Direction Error. To evaluate directional consistency of the model-predicted translation, we define the segment-wise direction error (SDE) as

$$e_{i,\Delta t} = \arccos\left(\left[\frac{(\mathbf{p}_{i+\Delta t} - \mathbf{p}_i)\cdot(\hat{\mathbf{p}}_{i+\Delta t} - \hat{\mathbf{p}}_i)}{\|\mathbf{p}_{i+\Delta t} - \mathbf{p}_i\|\,\|\hat{\mathbf{p}}_{i+\Delta t} - \hat{\mathbf{p}}_i\|}\right]_{-1}^{1}\right) \quad (7)$$

where only segments satisfying $\|\mathbf{p}_{i+\Delta t} - \mathbf{p}_i\| > \epsilon$ and $\|\hat{\mathbf{p}}_{i+\Delta t} - \hat{\mathbf{p}}_i\| > \epsilon$ are considered. The final SDE is obtained by averaging $e_{i,\Delta t}$ over all valid segments.

Results. Table III and Fig. 6 show clear representation-dependent failure modes in conventional baselines. In contrast, the proposed method shows a more consistent trend across A–C without the catastrophic failures observed in conventional baselines, with the largest gains on unseen sequences where frame-specific distribution shifts are most pronounced. These results indicate that the SO(3)-equivariant model is less sensitive to representation-specific distribution mismatch than conventional world-frame or body-frame formulations.

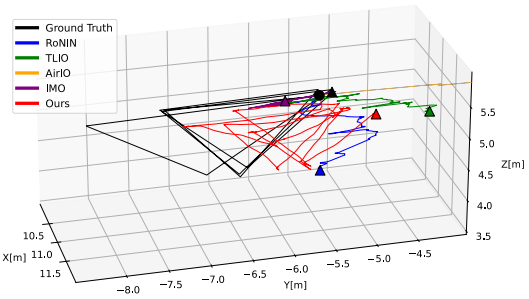

Fig. 5. Qualitative comparison on an unseen sequence with repeated triangular motion in the YZ plane. The results demonstrate the model's ability to generalize to an out-of-distribution trajectory dominated by motion along the z-axis.

TABLE III

SEGMENT-WISE DIRECTION ERROR (SDE, RAD) ON CASES A, B, AND C

Case	Sequence		RoNIN	TLIO	AirIO	IMO	Ours
A	**Seen**	**X**	0.0360	**0.0340**	0.5572	0.1389	0.1357
	Unseen	**Y**	1.2500	1.0525	1.6080	1.5757	**0.3525**
		Z	1.5633	1.5492	1.5584	1.5723	**1.0315**
B	**Seen**	**X**	0.0382	**0.0321**	0.2079	0.1346	0.0692
	Unseen	**Y**	0.2266	0.1738	**0.1002**	1.6263	0.2371
		Z	1.5656	1.5732	1.0631	1.5713	**0.2630**
C	**Seen**	**X**	0.0337	**0.0209**	0.2166	0.1063	0.0511
	Unseen	**Y**	0.5012	**0.2991**	1.5629	0.4854	0.3869
		Z	0.7013	0.5725	1.5690	2.7302	**0.4506**

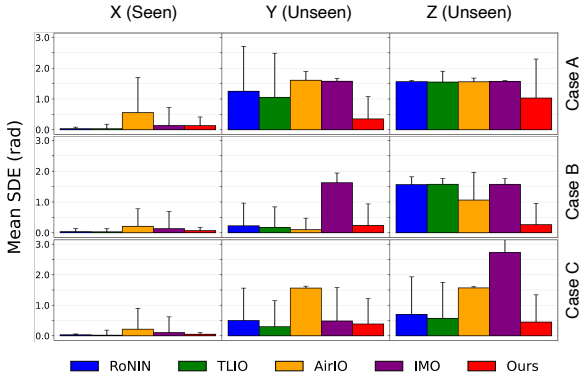

Fig. 6. Bar plot of mean segment-wise direction error (SDE) for the ablation cases, showing stronger directional consistency of the proposed method across controlled representation shifts.

V. CONCLUSION

We present an SO(3)-equivariant NIO framework for microgravity free-flying robots, where gravity-aligned assumptions from terrestrial inertial learning no longer hold. The proposed method combines a learning-based IMU correction module with a vector neuron motion network to preserve rotational consistency in translation estimation. Experiments on Astrobee simulation data demonstrate improved accuracy and stronger robustness to unseen motion conditions than existing baselines, supporting architectural equivariance as an effective inductive bias for inertial learning under unconstrained 3D microgravity motions.

VI. Limitation and Future Work

The evaluation is conducted exclusively on a simulation dataset. While this enables controlled analysis of SO(3)-equivariant generalization under microgravity motion, validation on real-world data [4] remains necessary. Addressing the sim-to-real gap through domain adaptation strategies to bridge discrepancies between simulated and physical IMU characteristics is left for future work. In addition, this work focuses specifically on translation estimation under known orientation, with the primary goal of investigating whether architectural SO(3)-equivariance provides a better inductive bias for displacement learning in microgravity. Orientation estimation under microgravity is a distinct problem in the learning-based inertial odometry literature and is left for future work. Combining displacement estimation with a microgravity orientation module within a unified framework remains another avenue for future work.

Acknowledgement

This work was supported by the National Research Foundation of Korea (NRF) grant funded by the Korea government (MSIT) (No.RS-2024-00358374) and the GIST Future-Leading Specialized Research Project grant funded by the GIST in 2026.

References

[1] T. Smith, O. Alexandrov, J. Barlow, J. Benavides, M. Bualat, R. Carlino, B. Coltin, J. Cortez, E. Daley, J. Feller, L. Flückiger, T. Fong, J. Fusco, R. G. Ruiz, K. Hamilton, S. Kanis, A. Katterhagen, Y. Kim, J. F. Love, M. McIntyre, B. McLachlan, A. M. Vargas, Z. Moratto, M. Moreira, T. Morse, H. Orosco, I.-W. Park, C. Provencher, H. Sanchez, K. Sharif, E. Smith, R. Soussan, A. Symington, R. O. Talavera, V. To, D. Wheeler, and J. Yoo, "Astrobee: Free-flying robots for the international space station," *IEEE Transactions on Field Robotics*, pp. 1–1, 2026.

[2] S. P. Yamaguchi, A. M. Vargas, T. Eisenberg, C. Rogon, T. Yamamoto, S. Inoue, C. Kössl, B. Coltin, T. Smith, and J. V. Benavides, "Free-flying crew cooperative robots on the iss: A joint review of astrobee, cimon, and int-ball operations," in *2025 International Conference on Space Robotics (iSpaRo)*, 2025, pp. 402–409.

[3] S. Mitani, M. Goto, R. Konomura, Y. Shoji, K. Hagiwara, S. Shigeto, and N. Tanishima, "Int-ball: Crew-supportive autonomous mobile camera robot on iss/jem," in *2019 IEEE Aerospace Conference*. IEEE, 2019, pp. 1–15.

[4] S. Kang, R. Soussan, D. Lee, B. Coltin, A. M. Vargas, M. Moreira, K. Browne, R. Garcia, M. Bualat, T. Smith, J. Barlow, J. Benavides, E. Jeong, and P. Kim, "Astrobee iss free-flyer datasets for space intravehicular robot navigation research," *IEEE Robotics and Automation Letters*, vol. 9, no. 4, pp. 3307–3314, 2024.

[5] S. Herath, H. Yan, and Y. Furukawa, "Ronin: Robust neural inertial navigation in the wild: Benchmark, evaluations, and new methods," in *2020 IEEE International Conference on Robotics and Automation (ICRA)*, 2020, pp. 3146–3152.

[6] W. Liu, D. Caruso, E. Ilg, J. Dong, A. I. Mourikis, K. Daniilidis, V. Kumar, and J. Engel, "Tlio: Tight learned inertial odometry," *IEEE Robotics and Automation Letters*, vol. 5, no. 4, pp. 5653–5660, 2020.

[7] X. Cao, C. Zhou, D. Zeng, and Y. Wang, "Rio: Rotation-equivariance supervised learning of robust inertial odometry," in *2022 IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*, 2022, pp. 6604–6613.

[8] R. K. Jayanth*, Y. Xu*, Z. Wang, E. Chatzipantazis, K. Daniilidis, and D. Gehrig, "EqNIO: Subequivariant neural inertial odometry," in *The Thirteenth International Conference on Learning Representations*, 2025.

[9] C. Deng, O. Litany, Y. Duan, A. Poulenard, A. Tagliasacchi, and L. Guibas, "Vector neurons: A general framework for so(3)-equivariant networks," in *2021 IEEE/CVF International Conference on Computer Vision (ICCV)*, 2021, pp. 12 180–12 189.

[10] Y. Qiu, C. Wang, C. Xu, Y. Chen, X. Zhou, Y. Xia, and S. Scherer, "Airimu: Learning uncertainty propagation for inertial odometry," *arXiv preprint arXiv:2310.04874*, 2023.

[11] L. Fluckiger and B. Coltin, "Astrobee robot software: Enabling mobile autonomy on the iss," Tech. Rep., 2019.

[12] Y. Qiu, C. Xu, Y. Chen, S. Zhao, J. Geng, and S. Scherer, "Airio: Learning inertial odometry with enhanced imu feature observability," *IEEE Robotics and Automation Letters*, vol. 10, no. 9, pp. 9368–9375, 2025.

[13] G. Cioffi, L. Bauersfeld, E. Kaufmann, and D. Scaramuzza, "Learned inertial odometry for autonomous drone racing," *IEEE Robotics and Automation Letters*, vol. 8, no. 5, pp. 2684–2691, 2023.