Vision-Based Relative State Estimation for Non-Cooperative Spacecraft Using Deep Learning and Adaptive Kalman Filtering

Safa Mesut Bostancı
Aerospace Engineering
Middle East Technical University
Ankara, Türkiye
mesut.bostanci@metu.edu.tr

Halil Ersin Söken
Aerospace Engineering
Middle East Technical University
Ankara, Türkiye
esoken@metu.edu.tr

***Abstract*—Autonomous proximity operations for on-orbit servicing and debris removal require accurate estimation of a target's translational and rotational motion. While deep learning approaches have achieved success in estimating spacecraft pose (position and attitude) from monocular imagery, they typically lack the velocity and angular rate estimates essential for trajectory prediction and control. This paper presents a pose-to-motion estimation framework that fuses pose measurements from a keypoint-based Convolutional Neural Network (CNN) pipeline with a Multiplicative Extended Kalman Filter (MEKF) to recover full six-degree-of-freedom (6-DoF) motion states. The measurement covariance is derived directly from Perspective-n-Point (PnP) reprojection geometry, providing adaptive uncertainty quantification. We evaluate the framework on trajectory sequence from the SPEED-UE-Cube dataset, demonstrating that the MEKF recovers velocity with RMSE of 0.002 m/s and angular rate with RMSE of 0.27°/s from noisy pose-only measurements. The filter simultaneously improves pose accuracy, reducing position RMSE from 6.96 m to 0.27 m and attitude RMSE from 34.4° to 7.8°. The results show that temporal filtering with geometry-derived measurement covariance enables reliable motion estimation from pose-only measurements, bridging the gap between single-frame pose estimation and the dynamic state information required for proximity operations.**

***Index Terms*—Uncooperative spacecraft pose estimation, motion estimation, Multiplicative Extended Kalman Filter, deep learning, proximity operations**

I. Introduction

Autonomous rendezvous and docking with an uncooperative spacecraft is a central enabling capability for forthcoming on-orbit servicing and logistics missions, including the supply and repair of the International Space Station and the automated inspection, maintenance, and assembly of space infrastructure. Unlike cooperative scenarios where the target is stabilized and equipped with markers, an uncooperative target may be non-responsive, poorly characterized, and tumbling, turning proximity operations into a tightly coupled estimation-and-control problem. The historical reliance on manual procedures for rendezvous and docking [1] highlights the operational gap that autonomous systems must address. Robust autonomy in this context requires a Guidance, Navigation, and Control (GNC) stack capable of inferring the target's relative state from intermittent and noisy sensing, since accurate relative position and attitude are essential for terminal docking alignment [2]–[5]. The urgency of this capability is amplified by the increasing density of objects in Low Earth Orbit and the growing population of defunct satellites, which raise the risk of cascading collision events known as the Kessler Syndrome [6]. Active Debris Removal and On-Orbit Servicing (ADR/OOS) missions can be viewed as uncooperative rendezvous-and-docking problems in which a servicer interacts with a dormant or dysfunctional target for removal, refueling, or repair [7]. These missions progress through phasing, approach, and capture stages [8]–[10], with the latter two dominated by close-range relative motion where autonomous navigation performance is most consequential [11]. Uncooperative targets impose stringent pose and motion requirements during terminal approach. Representative mission analyses indicate stringent requirements during terminal approach to uncooperative targets: relative attitude accuracy on the order of 0.3°, position accuracy below 0.2 m at 10 m separation, and translational velocity accuracy of approximately 0.01 m/s [11]–[13]. Safe contact is fundamentally limited by the estimator's ability to deliver accurate and stable relative state information in real time, particularly when the target is tumbling and measurement quality degrades due to self-occlusion and illumination changes.

II. Related Work

The use of Deep Learning (DL) methods, particularly Convolutional Neural Networks (CNNs), in monocular pose estimation has grown rapidly due to their strong performance in image classification tasks [14]. CNN-based pose estimation pipelines typically distinguish between an offline training phase and an in-flight inference stage, rather than separating image processing and pose estimation as traditional methods. Song [15] provides a comprehensive review of DL-based techniques for spacecraft relative navigation, classifying frameworks as either direct (end-to-end) methods that regress pose from images, or indirect (hybrid) methods that combine DL with classical algorithms such as Perspective-

n-Point (PnP) and Extended Kalman Filters (EKF). The survey by Pasqualetto Cassinis et al. [16] similarly categorizes monocular navigation methods into direct CNN-based pose regression [17], [18] and indirect approaches that detect image features via CNNs [19], [20] or classical image processing [21], [22], then solve a PnP problem against a known 3D target model. Recent works integrate DL pose estimators with recursive filters for improved robustness. Park and D'Amico [23] combined SPNv2, a multitask CNN, with an Unscented Kalman Filter (UKF) using Adaptive State Noise Compensation (ASNC) to tune process noise online, achieving sub-decimeter position and degree-level orientation accuracy on the SHIRT hardware-in-the-loop dataset. Candan and Servadio [24] proposed a marker-based pose estimation framework for ENVISAT using TinyCornerNET, a lightweight CNN for corner detection, paired with a UKF that inflates the process noise covariance during eclipse periods to maintain estimation robustness when markers become occluded.

III. Methodology

The proposed pipeline processes each incoming image with a deep-learning module that produces a pose estimate, which are fed into to a Multiplicative Extended Kalman Filter (MEKF) for state update. The complete prediction-update loop is illustrated in Fig. 1.

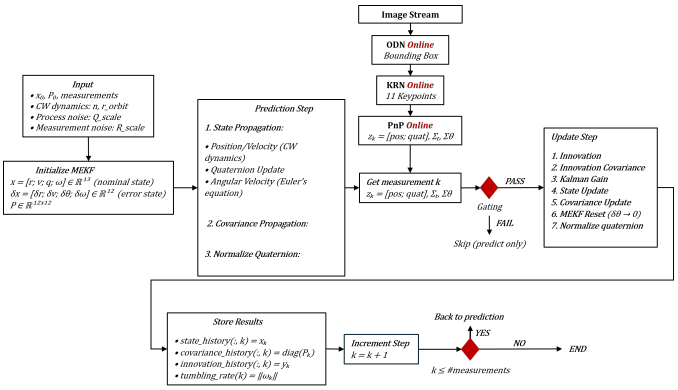

Fig. 1. System architecture. The ODN+KRN vision pipeline produces pose measurements $(\mathbf{p}^m, \mathbf{q}^m)$ with geometry-derived covariance $\mathbf{R}_k$, which the MEKF fuses with dynamics predictions to estimate the full state.

A. Dataset and Vision Pipeline

We use the SPEED-UE-Cube dataset [25], a photorealistic synthetic dataset for monocular pose estimation of a non-cooperative 3U CubeSat rendered in Unreal Engine 5. The dataset includes a full-orbit trajectory subset with 5 seconds capture intervals along simulated rendezvous trajectory, enabling tracking evaluation under realistic motion and illumination changes.

The original MATLAB Object Detection Network and Keypoint Regression Network (ODN+KRN) pipeline was reimplemented in PyTorch while preserving architectural equivalence with the YOLOv4-tiny baseline [26], including the CSPDarknet-tiny backbone and CIoU loss [27]. The ODN localizes the spacecraft and the KRN predicts 2D keypoints, which are solved via PnP to obtain pose measurements $(\mathbf{p}_k^m, \mathbf{q}_k^m)$.

B. Navigation Filter

We employ a 13-state MEKF that maintains the quaternion unit-norm constraint through multiplicative error-state updates [28]. The nominal state $\mathbf{x} = [\mathbf{p}^\top, \mathbf{v}^\top, \mathbf{q}^\top, \boldsymbol{\omega}^\top]^\top$ comprises position, velocity, attitude quaternion, and angular velocity, while the 12-dimensional error state replaces $\mathbf{q}$ with a minimal rotation vector $\delta\boldsymbol{\theta}$. State propagation uses Clohessy-Wiltshire dynamics [29] and the error-state covariance evolves as $\mathbf{P}_k^- = \mathbf{F}_k \mathbf{P}_{k-1} \mathbf{F}_k^\top + \mathbf{Q}_k$. The filter is initialized from the first pose measurement ($\hat{\mathbf{r}}_0 = \mathbf{r}_0^m$, $\hat{\mathbf{q}}_0 = \mathbf{q}_0^m$); since velocity and angular rate are not directly observable from a single image, they are initialized to zero with deliberately large diagonal entries in $\mathbf{P}_0 = \mathrm{blkdiag}(0.1\mathbf{I}_3, 1.0\mathbf{I}_3, 0.1\mathbf{I}_3, 0.5\mathbf{I}_3)$ to allow the filter to freely converge from measurements.

C. PnP-Derived Measurement Covariance

Rather than using fixed measurement noise, we derive $\mathbf{R}_k$ directly from the PnP solution geometry. Given N keypoint correspondences with reprojection residuals $\mathbf{e}_i = \hat{\mathbf{u}}_i - \pi(\mathbf{K}[\mathbf{R}|\mathbf{t}]\mathbf{X}_i)$, the pixel-space variance follows the maximum likelihood estimate [30]:

$$\hat{\sigma}^2_{\text{pix}} = \frac{\|\mathbf{e}\|^2}{2N-6}, \qquad \boldsymbol{\Sigma}_{\boldsymbol{\xi}} = \hat{\sigma}^2_{\text{pix}} \left(\mathbf{J}^\top \mathbf{J}\right)^{-1} \tag{1}$$

where $\mathbf{J} \in \mathbb{R}^{2N \times 6}$ is computed analytically using pinhole projection derivatives and Lie algebra rotation perturbations [31]. The translation and attitude blocks $(\boldsymbol{\Sigma}_{t,k}, \boldsymbol{\Sigma}_{\theta,k})$ form the block-diagonal $\mathbf{R}_k = \mathrm{blkdiag}(\boldsymbol{\Sigma}_{t,k}, \boldsymbol{\Sigma}_{\theta,k})$. The position residual is $\mathbf{y}_{p,k} = \mathbf{p}_k^m - \mathbf{p}_k^-$ and the attitude residual uses the quaternion log map $\mathbf{y}_{\theta,k} = \log(\mathbf{q}_k^m \otimes (\mathbf{q}_k^-)^{-1})$, yielding a small-angle vector compatible with $\boldsymbol{\Sigma}_{\theta,k}$.

D. Filter Parameter Optimization

The MEKF tuning parameters were optimized via Bayesian optimization with the Optuna framework [32] using the Tree-structured Parzen Estimator [33] (20 random trials followed by 1000 optimization trials). The objective is a weighted multi-metric cost:

$$J = w_p\,\mathrm{RMSE}_p + w_\phi\,\mathrm{RMSE}_\phi + w_\omega\,\mathrm{RMSE}_\omega + w_v\,\mathrm{RMSE}_v \tag{2}$$

with weights $(w_p, w_\phi, w_\omega, w_v) = (1.0, 0.21, 5.0, 333.33)$ chosen to normalize each term to comparable contribution given the unit and magnitude differences across position (m), attitude (deg), angular rate (deg/s), and linear velocity (m/s). The optimized parameters and ranges are summarized in Table I.

E. Baselines

We compare against two baselines: finite-difference motion estimation, which amplifies measurement noise, and the MEKF with fixed diagonal $\mathbf{R}$ optimized over identical parameter ranges as the PnP-derived variant. Comparing the two MEKF variants isolates the contribution of geometry-derived adaptive covariance from temporal filtering alone.

TABLE I
MEKF PARAMETERS AND BAYESIAN OPTIMIZATION RANGES.

Parameter	Range	Mode
σ_v (vel. proc. noise)	$[10^{-6}, 10^{-3}]$	Both
σ_θ (att. proc. noise)	$[10^{-4}, 0.05]$	Both
σ_ω (ang. vel. proc. noise)	$[10^{-7}, 0.005]$	Both
$\delta\theta_{\max}$	$[1, 30]$ deg	Both
$\delta r_{\max}$	$[0.005, 5]$ m	Both
$\delta v_{\max}$	$[0.0005, 0.5]$ m/s	Both
$\delta\omega_{\max}$	$[0.0005, 0.5]$ rad/s	Both
$e_{\text{reproj,max}}$	$[5, 100]$ px	Both
$\sigma_{\theta,\text{fix}}$	$[0.1, 10]$ deg	Fixed R
$\sigma_{r,\text{fix}}$	$[0.001, 1]$ m	Fixed R
k_Σ (cov. scale)	$[0.5, 20]$	PnP Cov.
$\sigma_{t,\min}/\sigma_{t,\max}$	$[0.005, 0.1]/[0.2, 2]$ m	PnP Cov.
$\sigma_{\theta,\min}/\sigma_{\theta,\max}$	$[0.1, 2]/[2, 30]$ deg	PnP Cov.

IV. RESULTS AND DISCUSSION

A. Pose Estimation Accuracy

Table II compares pose accuracy of raw deep learning measurements against MEKF-filtered estimates. The proposed MEKF with geometry-derived covariance and update clamping reduces position RMSE from 6.96 m to 0.271 m and attitude RMSE from 34.4° to 7.84° relative to raw DL predictions. The benefit of geometry-derived (PnP Cov.) over fixed covariance becomes apparent without update clamping: PnP covariance achieves 8× lower position RMSE (0.28 m vs 2.26 m) and 2.3× lower attitude RMSE (9.4° vs 22.1°). With clamping enabled, both methods reach comparable position accuracy, while PnP covariance retains an attitude advantage (7.84° vs 8.53° RMSE), indicating that geometry-derived covariance provides inherent robustness to outliers by adaptively reducing the Kalman gain for uncertain measurements.

TABLE II
POSE ESTIMATION PERFORMANCE.

		Trans. (m)		Att. (deg)	
Method	**Clamp**	Mean	RMSE	Mean	RMSE
Raw DL Pose	–	1.095	6.963	16.372	34.35
MEKF (Fixed $\mathbf{R}$)	Yes	**0.205**	0.272	7.01	8.53
MEKF (Fixed $\mathbf{R}$)	No	1.085	2.258	15.21	22.10
MEKF (PnP Cov.)	Yes	0.208	**0.271**	**6.64**	**7.84**
MEKF (PnP Cov.)	No	0.210	0.280	7.30	9.42

B. Motion Estimation Accuracy

As shown in Table III, finite differencing yields impractical velocity RMSE of 1.68 m/s and angular rate RMSE of 7.36°/s, which the MEKF reduces to 0.002 m/s and 0.274°/s. Without clamping, the performance gap between methods becomes substantial: PnP covariance maintains 0.318°/s angular rate RMSE compared to 0.517°/s for fixed covariance, confirming that geometry-derived covariance provides inherent outlier robustness that the fixed approach achieves only through explicit clamping.

C. Discussion

The clamped vs unclamped comparison highlights the role of measurement-quality-aware filtering. When clamping is disabled, the fixed method $\mathbf{R}$ is substantially degraded because uniform measurement trust prevents differentiation between reliable and unreliable observations, while the PnP covariance method maintains stable performance by assigning elevated uncertainty to geometrically ambiguous configurations. Throughout the 5930-second trajectory, the filter recovers all three angular velocity components without visible drift or divergence, indicating that the process noise model and geometry-derived covariance adequately characterize the system dynamics. Estimation errors remain within the bounds $\pm 2\sigma$ for the majority of the sequence, confirming reasonable uncertainty quantification, with a relative attitude RMSE of 2.7° demonstrating stable frame-to-frame tracking. The Z-axis position error shows a larger variance than the lateral axes, reflecting the inherent depth-scale ambiguity of monocular vision; the PnP-derived covariance captures this anisotropy and weights the depth estimates accordingly. The reprojection error threshold rejects approximately 4.8% of measurements exceeding 17.3 pixels, maintaining filter stability while preserving the majority of observations.

TABLE III
MOTION ESTIMATION PERFORMANCE.

		Vel. (m/s)		ω (deg/s)		Dir. (deg)	
Method	**Clamp**	$\bar{e}$	RMSE	$\bar{e}$	RMSE	$\bar{e}$	RMSE
Finite Diff.	–	0.288	1.675	3.275	7.358	49.68	67.17
Fixed $\mathbf{R}$	Yes	**0.002**	**0.002**	0.256	0.286	12.11	13.99
Fixed $\mathbf{R}$	No	0.005	0.008	0.440	0.517	22.77	27.78
PnP Cov.	Yes	**0.002**	**0.002**	**0.245**	**0.274**	**11.56**	**13.25**
PnP Cov.	No	0.002	0.004	0.270	0.318	13.18	15.65

V. CONCLUSION

This paper demonstrated that pose-only measurements from a deep learning pipeline can be elevated to full motion state estimation, including translational velocity and angular rate, through Multiplicative Extended Kalman Filtering. By deriving the measurement covariance directly from the PnP reprojection geometry, the filter adapts to the quality of each pose estimate. Evaluation of the SPEED-UE-Cube dataset showed that the MEKF substantially reduces the estimation errors of the velocity and angular rate compared to finite differencing while also improving pose accuracy through temporal smoothing, and geometry-derived covariance provides natural robustness to outliers, reducing the dependence on manually tuned safety mechanisms.

Future work will investigate learning-based covariance estimation, such as uncertainty derived from keypoint detection heatmaps or heteroskedastic networks that jointly predict pose and covariance, which could complement or replace the geometry-derived approach by capturing additional error sources not modeled by PnP reprojection analysis.

ACKNOWLEDGMENT

AI-assisted tools [34] were used for language editing and code debugging. All content was reviewed and verified by the authors.

REFERENCES

[1] J. H. Boynton and K. S. Kleinknecht, "Systems design experience from three manned space programs," *Journal of Spacecraft and Rockets*, vol. 7, no. 7, pp. 770–784, 1970.

[2] W. Fehse, *Automated Rendezvous and Docking of Spacecraft*. Cambridge university press, 2003, vol. 16.

[3] D. C. Woffinden and D. K. Geller, "Navigating the road to autonomous orbital rendezvous," *Journal of Spacecraft and Rockets*, vol. 44, no. 4, pp. 898–909, 2007.

[4] J. C. Morris, "Automated Spacecraft Docking Using a Vision-Based Relative Navigation Sensor," Ph.D. dissertation, Texas A&M University, 8 2010.

[5] M. E. A. A. E. Okasha, "Dynamics and Control of Satellite Relative Motion in Proximity," Ph.D. dissertation, Old Dominion University, 5 2012.

[6] D. J. Kessler and B. G. Cour-Palais, "Collision frequency of artificial satellites: The creation of a debris belt," *Journal of Geophysical Research: Space Physics*, vol. 83, no. A6, pp. 2637–2646, 6 1978.

[7] P. Colmenarejo, M. Graziano, G. Novelli, D. Mora, P. Serra, A. Tomassini, K. Seweryn, G. Prisco, and J. G. Fernandez, "On ground validation of debris removal technologies," *Acta Astronautica*, vol. 158, pp. 206–219, 5 2019.

[8] L. Kong and Y. Zhou, "Multiobjective Mission Planning for Multiple Geosynchronous Spacecraft Refueling," *International Journal of Aerospace Engineering*, vol. 2023, 2023.

[9] G. Arantes Jr. and L. S. Martins-Filho, "Guidance and control of position and attitude for rendezvous and dock/berthing with a noncooperative/target spacecraft," *Mathematical Problems in Engineering*, vol. 2014, 2014.

[10] E. Papadopoulos, F. Aghili, O. Ma, and R. Lampariello, "Robotic Manipulation and Capture in Space: A Survey," *Frontiers in Robotics and AI*, vol. 8, 7 2021.

[11] L. P. Cassinis, "Monocular Vision-Based Pose Estimation of Uncooperative Spacecraft," Ph.D. dissertation, Delft University of Technology, 2022. [Online]. Available: https://doi.org/10.4233/uuid:27dcbbc2-7d9e-4f67-925a-5e676ca4e43chttps://repository.tudelft.nl/islandora/object/uuid%3A27dcbbc2-7d9e-4f67-925a-5e676ca4e43c

[12] I. T. Mitchell, "Draper Laboratory Overview of Rendezvous and Capture Operations Motivation for Rendezvous," Draper Laboratory, Tech. Rep., 3 2010.

[13] J. Telaar, I. Ahrns, S. Estable, W. Rackl, M. De Stefano, R. Lampariello, N. Santos, P. Serra, M. Canetri, F. Ankersen, and J. Gil-Fernandez, "GNC architecture for the e.Deorbit mission," in *7th European Conference for Aeronautics and Space Sciences (EUCASS)*. Noordwijk; The Netherlands: ESA Publications Division, ESTEC, 2017, pp. 1–15.

[14] J. E. Ball, D. T. Anderson, and C. S. Chan, "A Comprehensive Survey of Deep Learning in Remote Sensing: Theories, Tools and Challenges for the Community," *Journal of applied remote sensing*, vol. 11, 9 2017. [Online]. Available: http://arxiv.org/abs/1709.00308http://dx.doi.org/10.1117/1.JRS.11.042609

[15] J. Song, D. Rondao, and N. Aouf, "Deep learning-based spacecraft relative navigation methods: A survey," *Acta Astronautica*, vol. 191, pp. 22–40, 8 2022. [Online]. Available: http://arxiv.org/abs/2108.08876http://dx.doi.org/10.1016/j.actaastro.2021.10.025

[16] L. Pasqualetto Cassinis, R. Fonod, and E. Gill, "Review of the robustness and applicability of monocular pose estimation systems for relative navigation with an uncooperative spacecraft," *Progress in Aerospace Sciences*, vol. 110, 10 2019.

[17] S. Sharma and S. D'Amico, "Neural Network-Based Pose Estimation for Noncooperative Spacecraft Rendezvous," *IEEE Transactions on Aerospace and Electronic Systems*, vol. 56, no. 6, pp. 4638–4658, 12 2020.

[18] K. Cosmas and A. Kenichi, "Utilization of fpga for onboard inference of landmark localization in cnn-based spacecraft pose estimation," *Aerospace*, vol. 7, no. 11, pp. 1–25, 11 2020.

[19] Y. Jia, E. Shelhamer, J. Donahue, S. Karayev, J. Long, R. Girshick, S. Guadarrama, and T. Darrell, "Caffe: Convolutional architecture for fast feature embedding," in *MM 2014 - Proceedings of the 2014 ACM Conference on Multimedia*. Association for Computing Machinery, 11 2014, pp. 675–678.

[20] S. Sonawani, R. Alimo, R. Detry, D. Jeong, A. Hess, and H. B. Amor, "Assistive relative pose estimation for on-orbit assembly using convolutional neural networks," *AIAA Scitech 2020*, vol. 1 PartF, pp. 1–11, 1 2020. [Online]. Available: http://arxiv.org/abs/2001.10673http://dx.doi.org/10.2514/6.2020-2096

[21] Z. Wang, Z. Zhang, X. Sun, Z. Li, and Q. Yu, "Revisiting Monocular Satellite Pose Estimation with Transformer," *IEEE Transactions on Aerospace and Electronic Systems*, vol. 58, no. 5, pp. 4279–4294, 10 2022.

[22] R. Volpe, G. B. Palmerini, and M. Sabatini, "A passive camera based determination of a non-cooperative and unknown satellite's pose and shape," *Acta Astronautica*, vol. 151, pp. 805–817, 10 2018.

[23] T. H. Park and S. D'Amico, "Adaptive Neural-Network-Based Unscented Kalman Filter for Robust Pose Tracking of Noncooperative Spacecraft," *Journal of Guidance, Control, and Dynamics*, vol. 46, no. 9, pp. 1671–1688, 6 2023.

[24] B. Candan and S. Servadio, "Relative Pose Estimation of an Uncooperative Target with Camera Marker Detection," *Aerospace*, vol. 12, no. 5, 5 2025.

[25] Z. Ahmed, H. Park, A. Bhattacharjee, R. Fazel-Rezai, R. Graves, O. Saarela, R. T. ——, K. Vemulapalli, and S. D'amico, "SPEED-UE-CUBE: A MACHINE LEARNING DATASET FOR AUTONOMOUS, VISION-BASED SPACECRAFT NAVIGATION," in *46th Rocky Mountain AAS Guidance, Navigation and Control Conference*, 2024, pp. 24–027. [Online]. Available: https://github.com/tpark94/speed-ue-cube-baseline

[26] A. Bochkovskiy, C.-Y. Wang, and H.-Y. M. Liao, "YOLOv4: Optimal Speed and Accuracy of Object Detection," 4 2020. [Online]. Available: http://arxiv.org/abs/2004.10934

[27] Z. Zheng, P. Wang, W. Liu, J. Li, R. Ye, and D. Ren, "Distance-IoU Loss: Faster and Better Learning for Bounding Box Regression," 11 2019. [Online]. Available: http://arxiv.org/abs/1911.08287

[28] F. Landis Markley and J. L. Crassidis, *Fundamentals of Spacecraft Attitude Determination and Control*. Springer, 2014, vol. 33. [Online]. Available: http://www.springer.com/series/6575

[29] W. H. CLOHESSY and R. S. WILTSHIRE, "Terminal Guidance System for Satellite Rendezvous," *Journal of the Aerospace Sciences*, vol. 27, no. 9, pp. 653–658, 9 1960.

[30] R. Hartley and A. Zisserman, *Multiple view geometry in computer vision* . Cambridge University Press, 2004, vol. 2.

[31] T. D. Barfoot, *State Estimation for Robotics*. Cambridge University Press, 2017.

[32] T. Akiba, S. Sano, T. Yanase, T. Ohta, and M. Koyama, "Optuna: A Next-generation Hyperparameter Optimization Framework," in *Proceedings of the ACM SIGKDD International Conference on Knowledge Discovery and Data Mining*. Association for Computing Machinery, 7 2019, pp. 2623–2631.

[33] J. Bergstra, R. Bardenet, Y. Bengio, and B. Kégl, "Algorithms for Hyper-Parameter Optimization," in *NIPS'11: Proceedings of the 25th International Conference on Neural Information Processing Systems*, 12 2011, pp. 2546 – 2554.

[34] Anthropic, "Claude," 2025. [Online]. Available: https://www.anthropic.com/claude