Online Supervised Training of Spaceborne Vision during Proximity Operations using Adaptive Kalman Filtering

Tae Ha "Jeff" Park, Simone D'Amico
Space Rendezvous Laboratory (SLAB)
Department of Aeronautics & Astronautics, Stanford University

RENDEZVOUS SPACE
Stanford Space Rendezvous Laboratory
slab.stanford.edu

ICRA2024 YOKOHAMA | JAPAN

Introduction

This work presents an Online Supervised Training (OST) method to enable robust vision-based navigation about a noncooperative spacecraft. Spaceborne Neural Networks (NN) are susceptible to *domain gap* as they are primarily trained with synthetic images due to the inaccessibility of space. OST aims to fully close this gap by training a pose estimation NN online using incoming flight images during Rendezvous and Proximity Operations (RPO). The pose pseudo-labels for supervised fine-tuning are provided by the onboard navigation filter which tracks the target's relative orbital and attitude motion by using the NN as its measurement module. The experiments on real hardware-in-the-loop trajectory images show that OST can improve the NN performance on the target image domain given that OST is performed on images of the target viewed from a diverse set of directions during RPO.

The project website is available at https://taehajeffpark.com/ost/

ML-in-the-Loop Navigation in Space

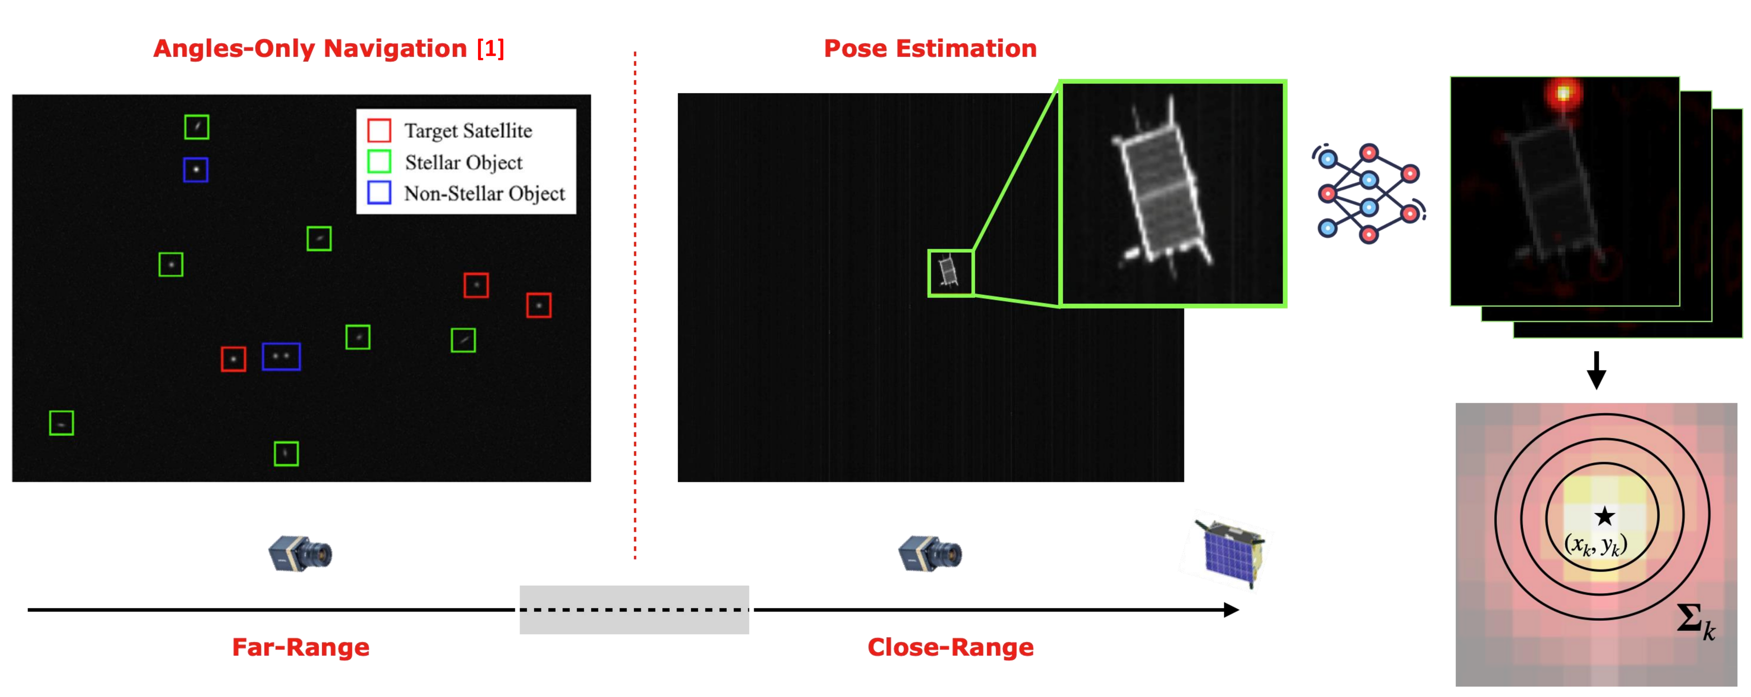

In order to accommodate as dynamic range as possible, all the way from the docking range to as far as the distance at which the target appears barely resolved, the image is first cropped around the target using the bounding box estimated from either the navigation filter or a separate object detection NN. The cropped image is then resized and fed into a NN which outputs heatmaps around the known surface keypoints of the target spacecraft.

Datasets

On-Ground, Offline Robust Training

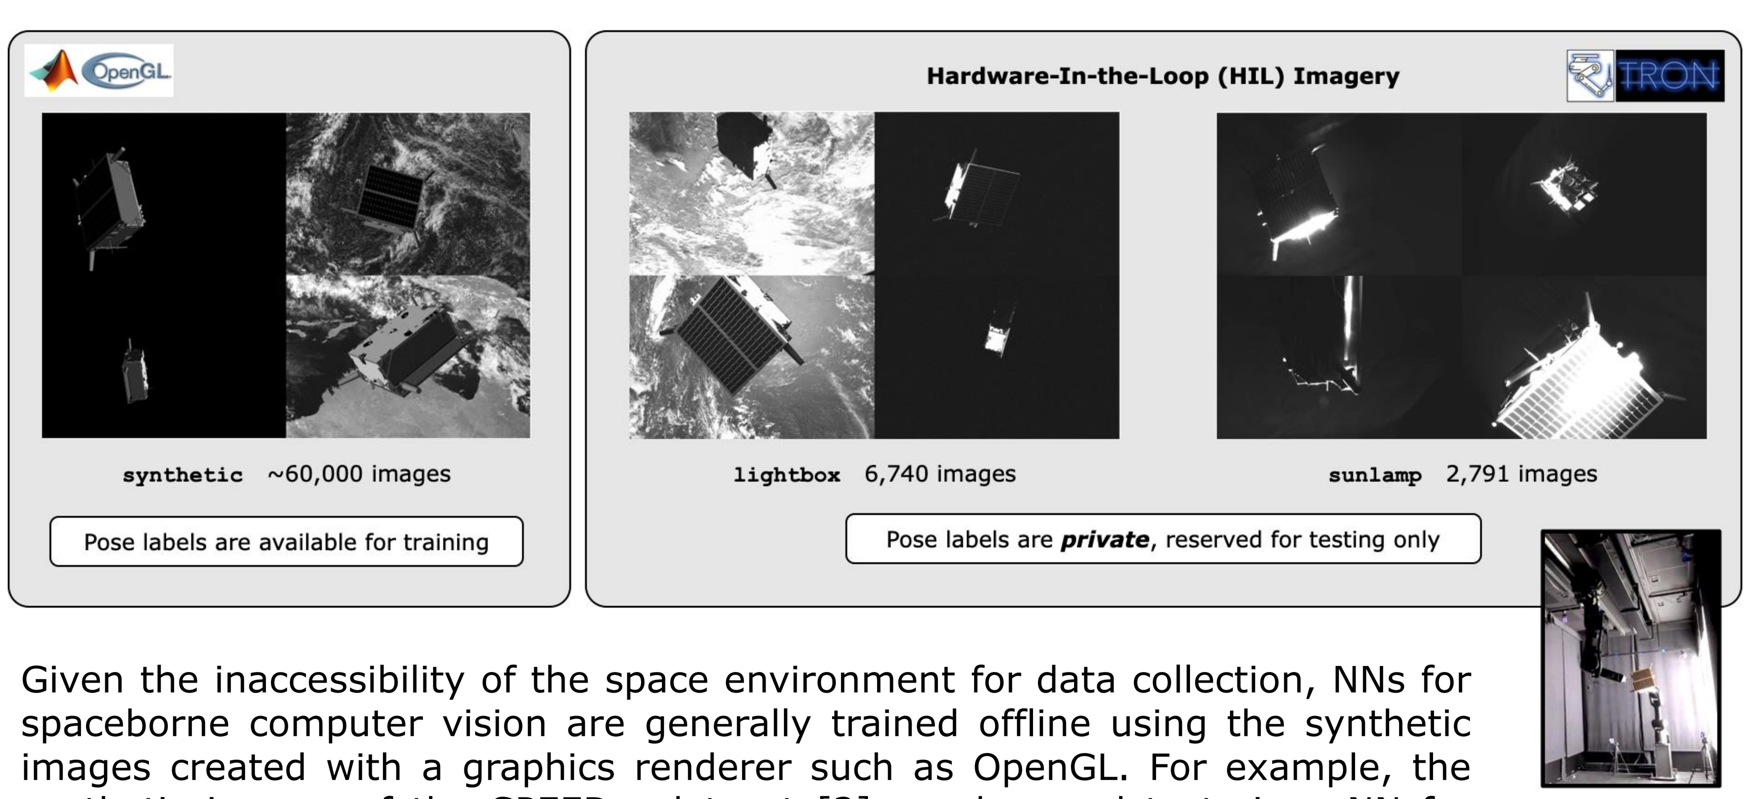

Given the inaccessibility of the space environment for data collection, NNs for spaceborne computer vision are generally trained offline using the synthetic images created with a graphics renderer such as OpenGL. For example, the synthetic images of the SPEED+ dataset [2] can be used to train a NN for monocular pose estimation of a non-cooperative spacecraft. Then, without spaceborne images to validate its performance, the Hardware-In-the-Loop (HIL) images from the robotic testbed such as TRON at SLAB [3] can be used as on-ground surrogates of otherwise unavailable flight images for robustness evaluation.

Online Supervised Training

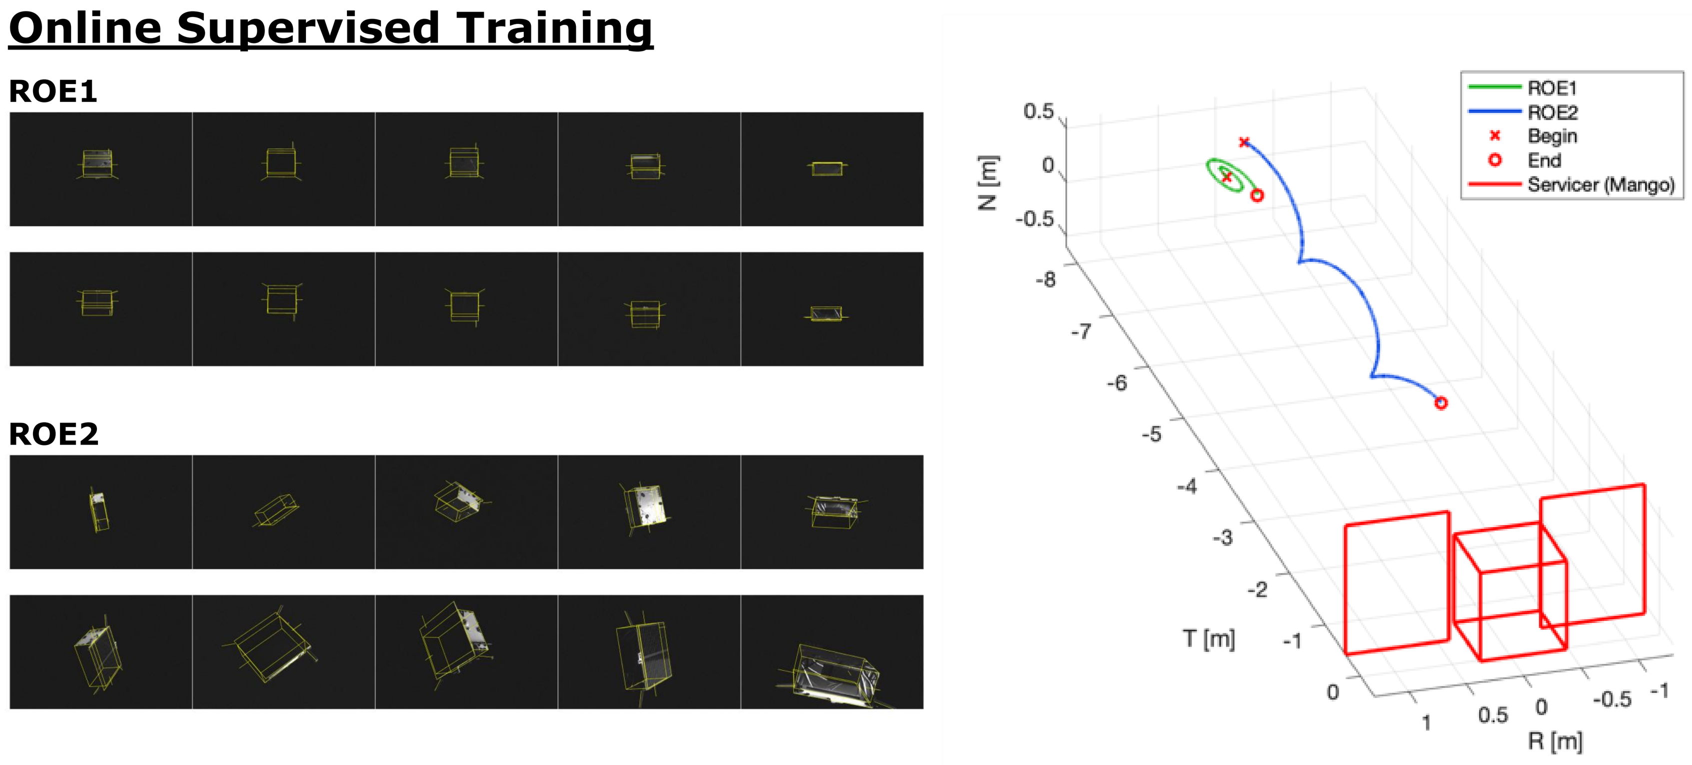

The SHIRT dataset [4] instead consists of two RPO trajectories in sun-synchronous LEO with two image domains: `synthetic` and `lightbox`. It is used for 1) offline robustness evaluation of the navigation filter with the NN as its measurement model and 2) performing OST. Note that ROE1 is more difficult as the target is located at a nearly constant distance from the camera while rotating along a single axis parallel to the image plane.

References
[1 J. Kruger, A. W. Koenig, S. D'Amico. The Starling Formation-Flying Optical Experiment (StarFOX): System Design, Experiment Design and Pre-Flight Verification. *Journal of Spacecraft and Rockets*, vol. 60, no. 6, pp.1755—1777 (2023)
[2] T. H. Park, M. Märtens, G. Lecuyer, D. Izzo, S. D'Amico. SPEED+: Next-Generation Dataset for Spacecraft Pose Estimation across Domain Gap. *IEEE Aerospace Conference*, Big Sky, Montana, March 5—12 (2022)
[3] T. H. Park, J. Bosse, S. D'Amico. Robotic Testbed for Rendezvous and Optical Navigation: Multi-Source Calibration and Machine Learning Use Cases. *2021 AAS/AIAA Astrodynamics Specialist Conference*, Big Sky, Virtual, August 9—11 (2021)
[4] T. H. Park and S. D'Amico. Adaptive Neural-Network-Based Unscented Kalman Filter for Robust Pose Tracking of Noncooperative Spacecraft. *Journal of Guidance, Control, and Dynamics*, vol. 46, no. 9, pp. 1671—1688 (2023)

Method

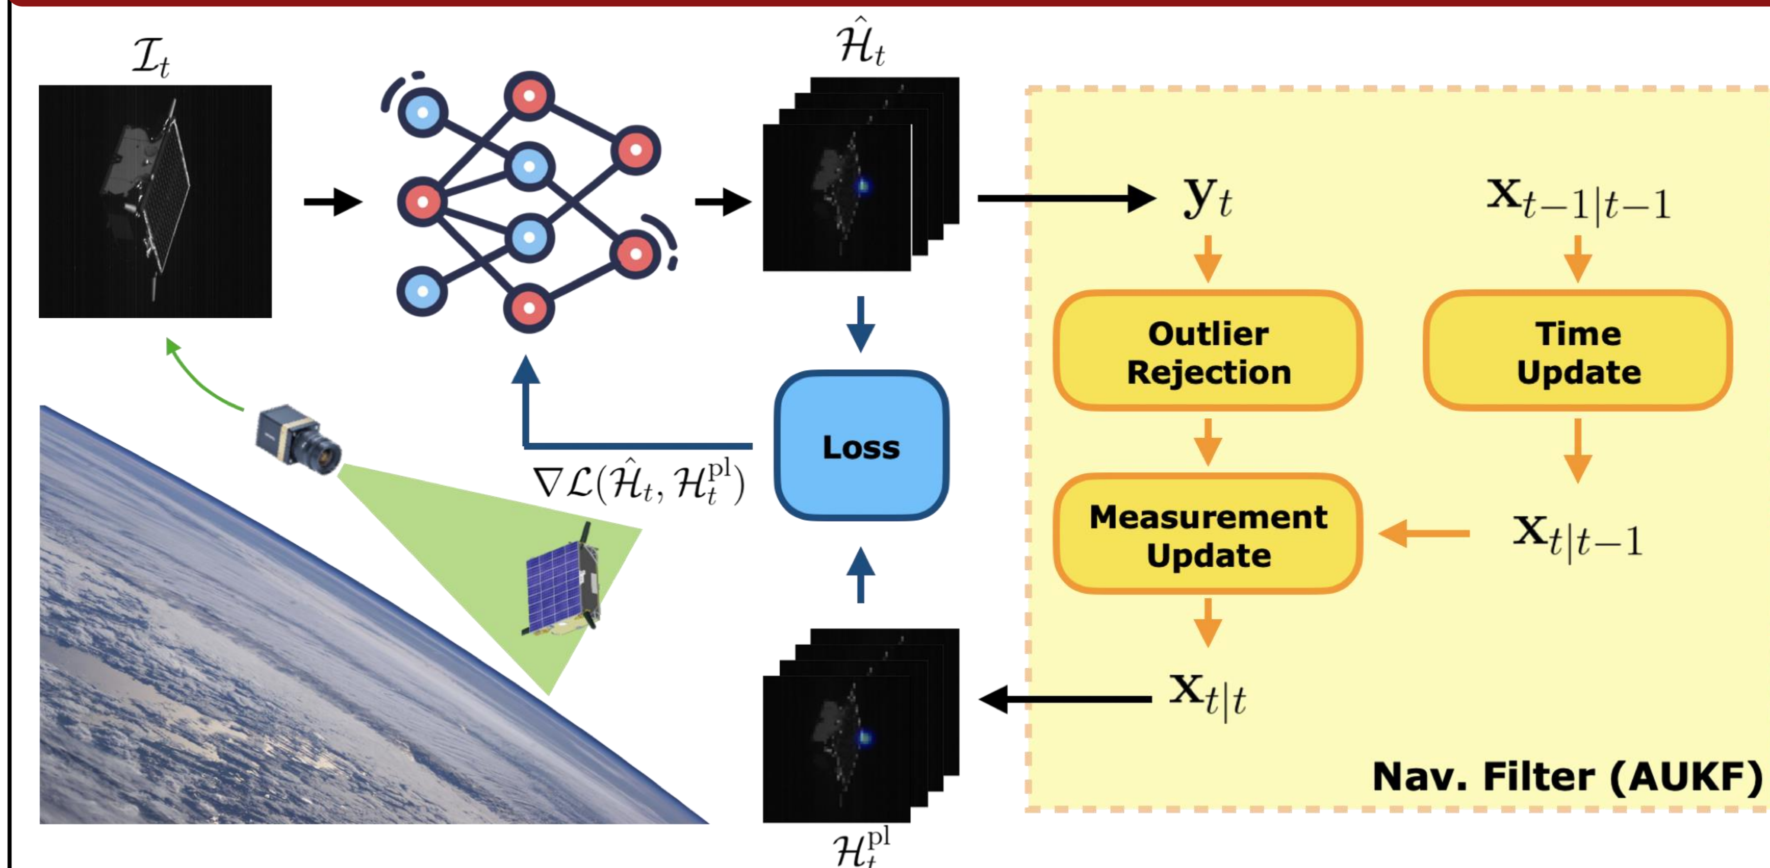

OST is performed by first predicting the heatmaps from an input flight image. The heatmaps are converted to keypoint measurements and uncertainties, and they are provided to the onboard navigation filter which performs time update, measurement update, outlier rejection and process noise update (see paper). The a posteriori state estimates are then used to generate heatmap pseudolabels by drawing a Gaussian blob around the keypoint locations. Finally, the NN is trained with a mean squared error loss function. The training only involves a single round of backpropagation on a single image to minimize incurred computational overhead. In this work, the training frequency is manually set to every 10 images.

Results

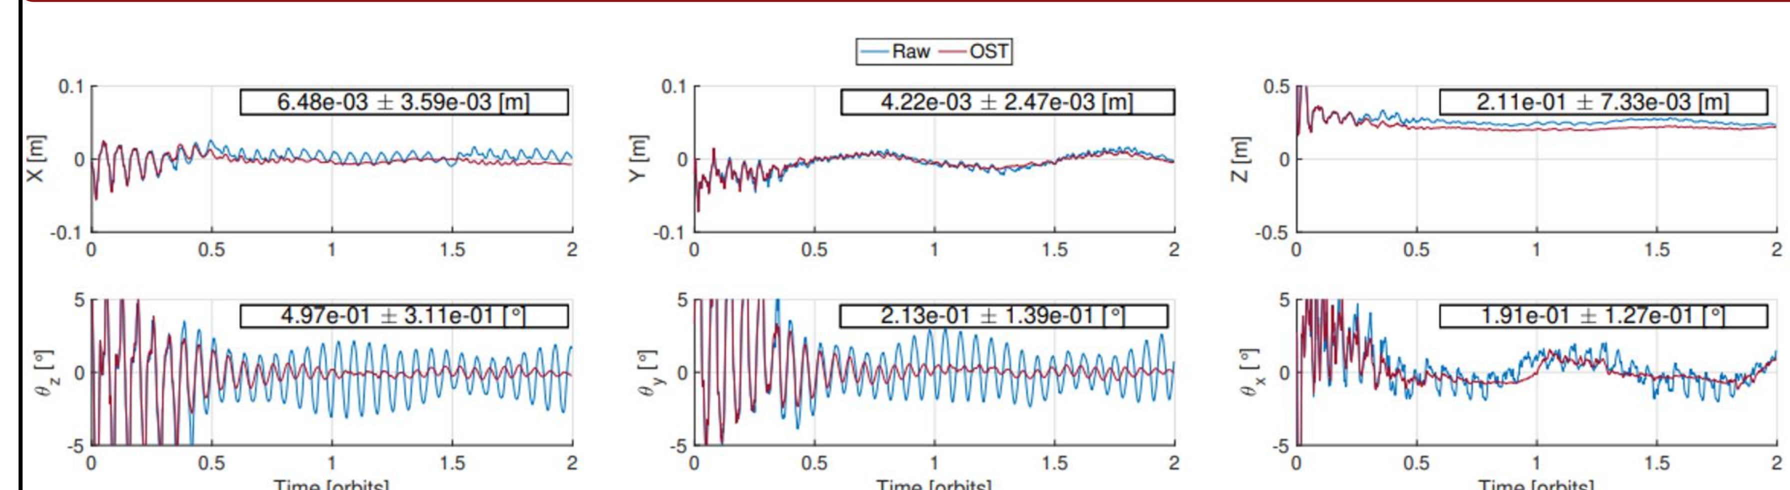

The performance improvement is most visible on a more difficult ROE1 scenario, where activating OST significantly improves the filter's orientation estimates.

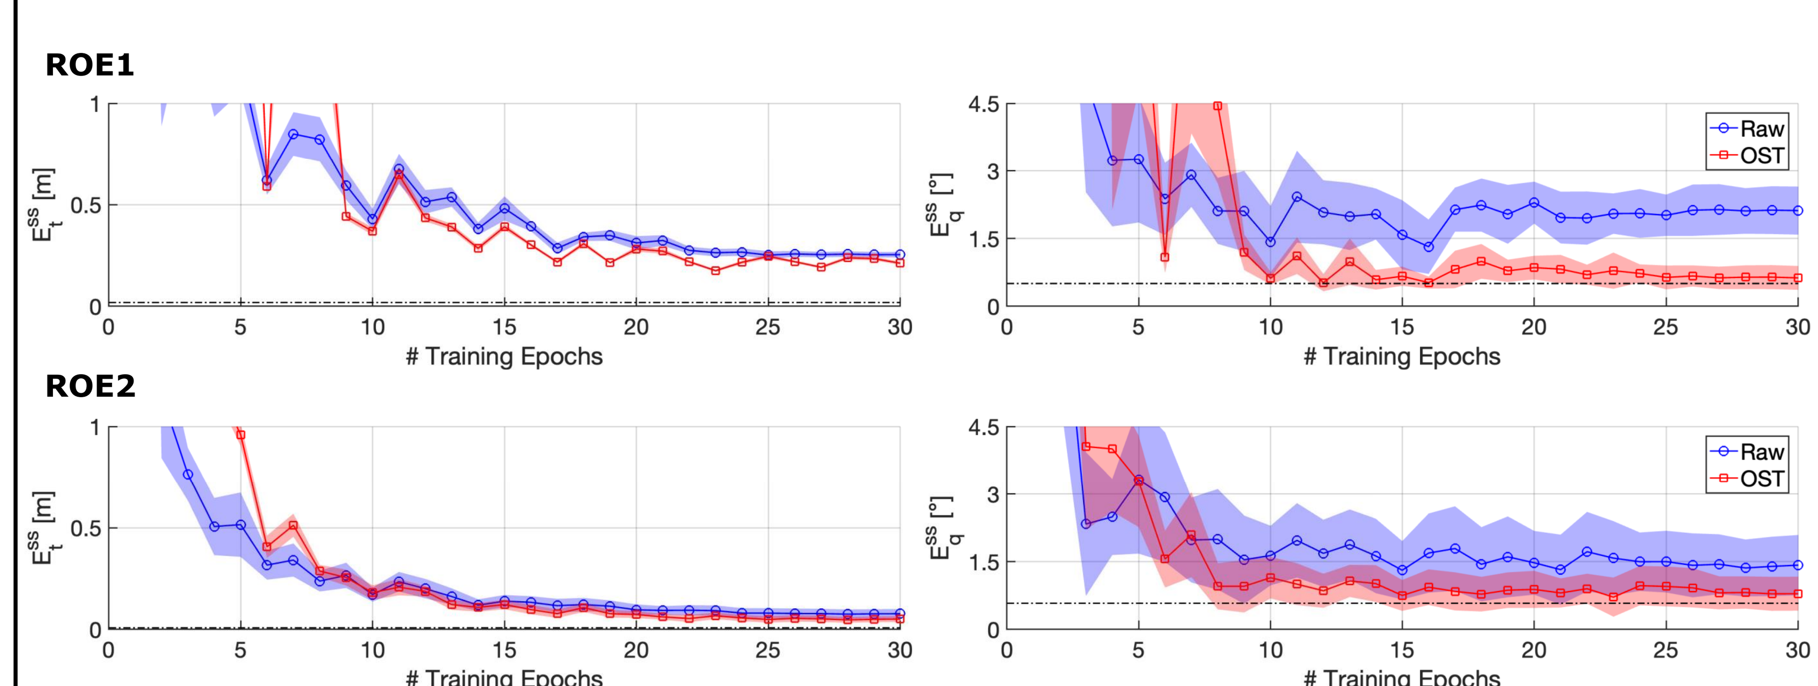

Different levels of domain gap are simulated by prematurely ending the training of a NN, and the filter's steady-state errors for translation and orientation components are shown. We can see that the NN trained for only 10 epochs out of a 30-epoch training schedule already improves the overall filter steady-state errors. Moreover, the orientation error is equivalent to that of the filter evaluated on synthetic images of the same trajectory (black horizontal lines), suggesting that OST fully closes the domain gap for orientation prediction.

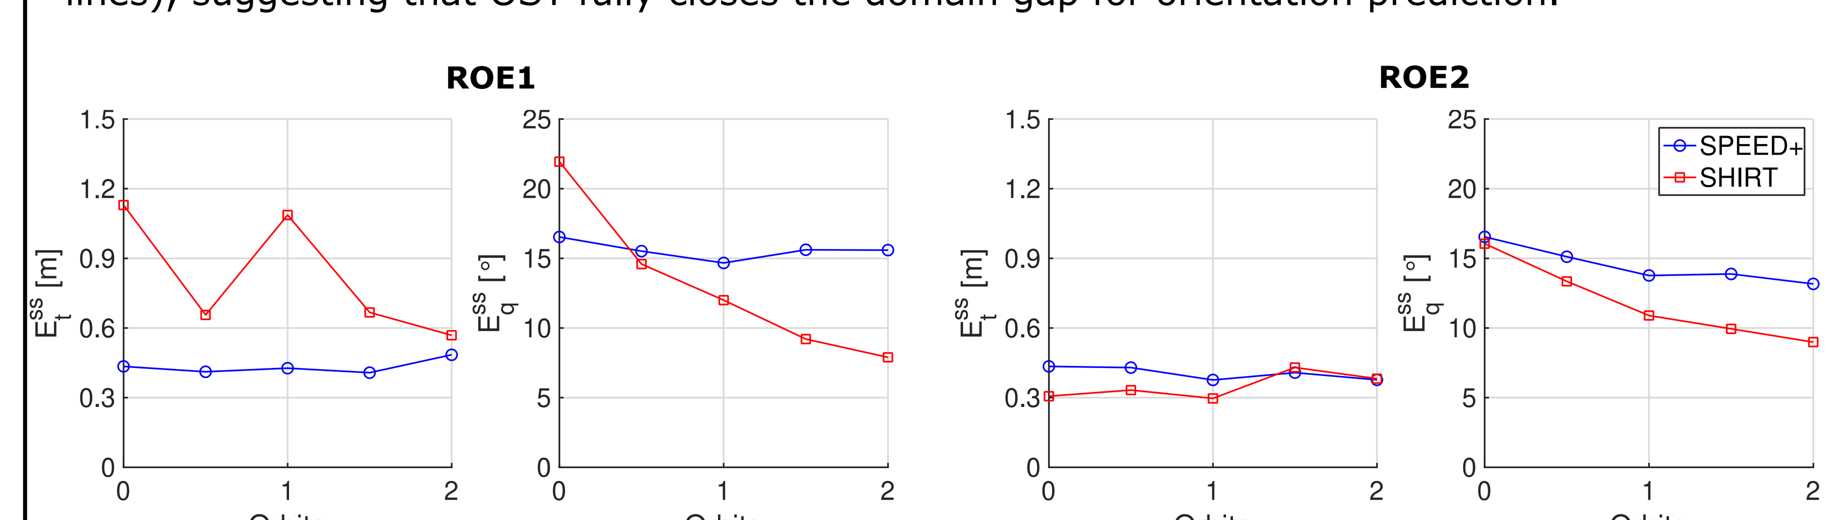

As NN is subject to longer duration of OST, its steady-state performance on SHIRT trajectories decreases monotonically, indicating overfitting to images encountered in RPO. However, its performance on SPEED+ `lightbox` images remains about the same, indicating the performance improvement is not at the expense of loss of generalizability. In fact, on ROE2, its performance on SPEED+ decreases, indicating the geometric diversity in OST training images can improve the general performance across the test domain.

Future Work

In the future, the proposed OST will be evaluated on much longer trajectories with HIL images in order to properly investigate the long-term coupling effect of OST and the navigation filter. Furthermore, OST will be improved by better reflecting the filter state uncertainties and incorporating more rigorous criteria to schedule OST during proximity operations.

This work is supported by the U.S. Space Force SpaceWERX Orbital Prime Small Business Technology Transfer (STTR) contract number FA8649-23-P0560 awarded to TenOne Aerospace in collaboration with SLAB.