

Abstract

This PhD thesis presents the development and verification of Guidance, Navigation, and Control (GNC) algorithms for the Space Rider Observer CubeSat (SROC) mission, with specific focus on terminal rendezvous and docking with the Space Rider vehicle. The research is motivated by the growing importance of autonomous proximity operations in applications such as spacecraft inspection, in-orbit servicing, debris removal, and reusable orbital platforms. Within this framework, SROC is conceived as a 12U technology demonstrator aimed at validating enabling technologies for operations in the vicinity of a larger spacecraft. The main objective of the thesis is to assess the feasibility of an integrated GNC architecture for SROC based on monocular visual navigation, state estimation, and constrained control, and to verify its suitability for safe and reliable terminal rendezvous and docking under realistic orbital conditions.

A central contribution is the development of a monocular vision-based navigation solution based on a commercial off-the-shelf camera. Relative pose estimation is carried out with a state-of-the-art convolutional neural network (CNN) algorithm. The adopted architecture predicts both the relative pose of the target and a set of body keypoints, thus enabling a multi-task formulation that improves robustness under difficult visual conditions.

The network training dataset consists of synthetic imagery produced by a custom Blender-based pipeline. The rendering environment was designed to reproduce representative orbital conditions, including target attitude, relative geometry, illumination variations, and camera-related effects. Two synthetic datasets were generated, one for the observation phase and one for the final-approach phase. In order to ensure that the navigation algorithm remains robust in practical settings, the methodology prioritizes narrowing the synthetic-to-real domain gap through extensive data augmentation and hyperparameter optimization. Furthermore, the training logic was refined by evaluating several workflows, ranging from sequential fine-tuning across mission phases to the use of mixed-domain datasets.

The trained models were tested on real images acquired at ESA's GNC Rendezvous, Approach and Landing Simulator using a scaled mock-up of Space Rider. The results showed promising pose-estimation performance and satisfactory robustness, while also highlighting the influence of dataset characteristics and residual domain mismatch. In parallel, the feasibility of onboard deployment was investigated by exporting the trained

model and integrating it into a C++ inference pipeline. Inference-time and memory-usage tests on flight-representative embedded hardware confirmed the practical compatibility of the proposed solution with onboard computational constraints.

The visual estimates produced by the CNN, together with the uncertainty associated with keypoint detection, were then integrated into an Extended Kalman Filter for relative state estimation. The navigation chain was integrated in a high-fidelity MATLAB/Simulink simulator reproducing the final approach under realistic orbital conditions, including translational and rotational dynamics, actuator models, environmental perturbations, and measurement noise. For the control layer, a chance-constrained stochastic Model Predictive Control strategy was implemented, propagating the estimation covariance along the prediction horizon to enforce the docking-corridor and safety constraints under explicit navigation uncertainty.

The robustness of the GNC architecture was assessed through an extensive Monte Carlo campaign accounting for initialization errors, measurement uncertainties, and model dispersions. The resulting simulations quantified the response of the integrated closed-loop architecture to the considered uncertainties. Within the modeled operating conditions, the proposed GNC architecture was shown to satisfy the adopted terminal performance and safety requirements.