Angles-Only Relative Orbit Determination Using Deep Recurrent Neural Networks

{jasulliv, damicos} @ stanford.edu
slab.stanford.edu

Josh Sullivan, Simone D'Amico

Space Rendezvous Laboratory

Department of Aeronautics and Astronautics, Stanford University, Stanford, CA 94305

Motivation

Future missions involving the interaction of multiple satellites present increasingly demanding relative navigation requirements which must be achieved autonomously using limited onboard resources. Vision-based navigation techniques deliver an effective response to these needs by providing an inherently passive, robust, and high-dynamic range capability which uses simple sensors that are already on board most spacecraft. Furthermore, because of their low cost, low power consumption, and small form factor as compared with other metrology systems, these sensors enable accurate relative navigation while complementing the current trend of spacecraft miniaturization.

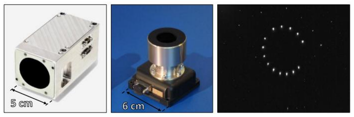

State-of-the-art star trackers (left and center). Overlaid images taken of a target spacecraft during the ARGON experiment (right) using a DTU star tracker.

Problem Statement and Challenges

This research project focuses on the far-range angles-only relative navigation problem, where an observing spacecraft is seeking to estimate the relative motion of a target space object. The features available to the observer include its own orbital and attitude state, as well as bearing angles which subtend the line-of-sight vector pointing to the target. In this problem, constrained dynamical observability generally makes it difficult to estimate the full 6D relative state from these simple feature sets. Common solutions approach the problem by iteratively linearizing the nonlinear dynamical system equations to employ least-squares batch estimation. However, this method inherently neglects system nonlinearities which can actually improve the dynamical observability. Furthermore, while it is rather simple to generate large and relatively high-fidelity data sets representative of real scenarios, no current approaches make use of deep learning.

Supervised Learning Problem and Data Generation

Project goal: Use deep recurrent neural networks (RNN) to learn the relationship between the observer-obtained measurements and the unique relative orbital state of the target. Evaluate the RNN model performance over a variety of hyperparameter choices.

Input data features: Each training example consists of temporally-ordered sequences of 12-dimensional feature data, consisting of:

- Observer absolute orbit $\in \boldsymbol{R}^6$, as parameterized by a set of classical orbital elements $(a, e, i, \Omega, \omega, M)^T$
- Observer absolute attitude $\in \boldsymbol{R}^4$, as parameterized by a set of attitude quaternions $(q_0, \boldsymbol{q})^T$
- Bearing angles $\in \boldsymbol{R}^2$, subtending the line-of-sight vector in the camera frame, given by the azimuth and elevation $(\alpha, \epsilon)^T$

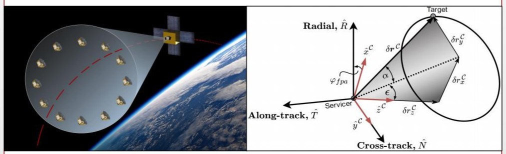

Outputs to be learned: This project is framed as a supervised learning problem, where the target relative state corresponding to the sequence of input data is provided as the truth.

- The target relative state $\in \boldsymbol{R}^6$, corresponding to the last measurement sequence is to be learned. It is parameterized by a set of Relative Orbital Elements: $(\delta a, \delta\lambda, \delta e_x, \delta e_y, \delta i_x, \delta i_y)$

Dataset generation: Large high-fidelity datasets of the observer and target orbital motion can be numerically propagated using the Space Rendezvous Satellite Simulator (S^3) which realistic sensor emulators which capture representative noise characteristics.

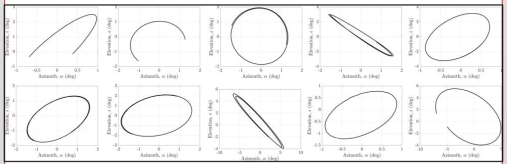

Hyperparameter study: This project focuses on building intuition on the applicability of RNNs for Relative Orbit Determination. Accordingly, four distinct hyperparameter studies were done:

1) <u>Learning rate</u> using the Adam optimization method.
2) <u>RNN layer depth</u> using fixed number of LSTM units
3) <u>Number of LSTM units</u> using fixed RNN layer depth
4) <u>Number of temporal sequences</u> of feature data.

Deep Learning Framework and Results

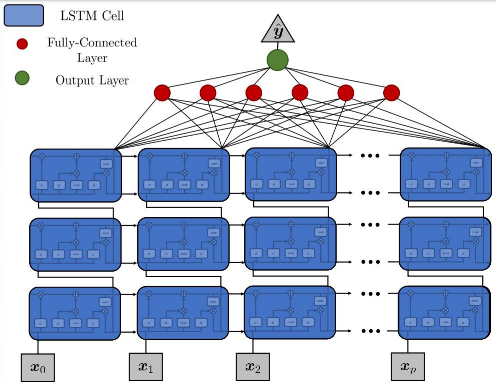

Test cases (underlined is nominal for other cases):

- Learning rate: 0.001, 0.01, <u>0.1</u>
- RNN layer depth with LSTM cells: 1 layer, <u>5 layers</u>, 10 layers
- Number of LSTM units: 64 units, <u>128 units</u>, 256 units
- Number of temporal sequences (p): <u>20 sequences</u>, 40 sequences
- 30,000 total examples, 80% train / 10% dev. / 10% test split
- Trained over 50 epochs

Hyperparameter	Case 1	Case 2	Case 3
Learning Rate	5.13 km	3.78 km	2.45 km
RNN Depth	7.42 km	2.45 km	1.89 km
LSTM Units	2.86 km	2.45 km	2.63 km
Temporal Sequences	2.45 km	1.64 km	-

Table 1: Average relative orbit determination *test* errors in kilometers.

Discussion: These results outperform several current analytic approaches which leverage only a reduced dynamics approach through measurement model linearization. In general, the learning rate was found to improve training and dev accuracy up to a certain point, and this is reflected in the test results. If more epochs were allowed for the smaller learning rates, errors would likely be smaller. RNN depth was shown to have a substantial effect on test accuracy, with deeper networks displaying the best results. Instead, number of LSTM units generally displayed minimal effect on test performances. Finally, adding more temporal sequences improved test performance. It is expected, however, that continuing to add more temporal sequences will eventually yield diminishing returns.