Multi-objective Autonomous Spacecraft Motion Planning around Near-Earth Asteroids using Machine Learning

Tommaso Guffanti, tommaso@stanford.edu

people.stanford.edu/damicos

Problem Statement

A spacecraft (SC) capable to autonomously plan its motion while accounting for conflicting mission objectives in a Pareto optimal way would permit to accomplish complex mission tasks around highly uncharacterized celestial bodies such as near-Earth asteroids (NEAs). The two most common conflicting objectives of a space exploration mission are the maximization of the scientific output and the minimization of the control effort, i.e. the propellant required on board. If the targeted celestial body is sufficiently known and has a predictable orbital environment, both numerical and analytical tools can be leveraged on-ground to design spacecraft motion plans that account for the trade-off. On the contrary, if the celestial body and its orbit environment are largely uncertain, all plans elaborated on-ground may fail dramatically when implemented in space. A clear example are missions around NEAs having relevant dynamics parameters (i.e. mass, shape, rotation axis orientation and gravity coefficients) largely uncertain. In these missions, a spacecraft should be capable of autonomously plan its motion when an updated knowledge of the NEA environment is provided by the sensors and the navigation filter. In addition, the generated motion plan should account for the trade-off science-output vs. control-effort in an optimal sense.

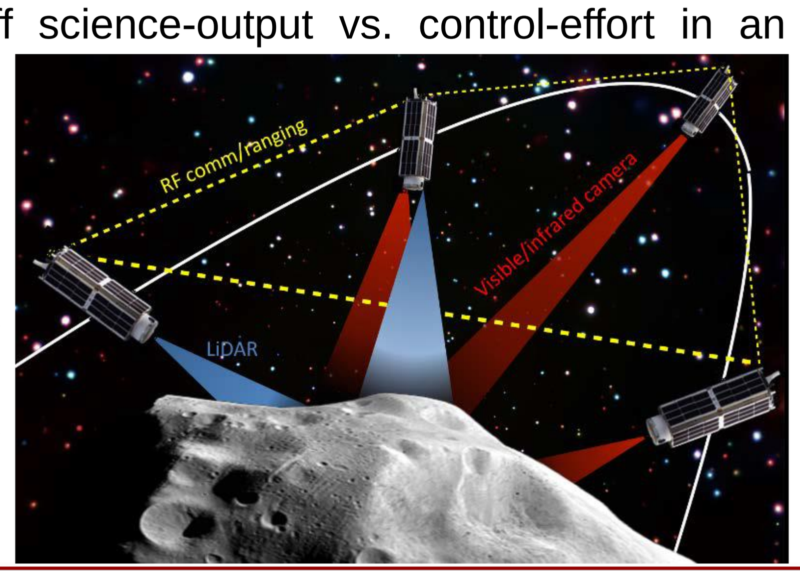

Autonomous Nanosatellite Swarming using Radio Frequency and Optical Navigation *credit: Stanford's Space Rendezvous LAB.*

Solution Approach

The solution approach takes inspiration from both Reinforcement and Supervised Learning. Let's define the state $s = [\mathbf{p_{NEA}}, \boldsymbol{\lambda}]$. Where $\mathbf{p_{NEA}}$ are the NEA uncertain dynamics parameters, in this case, its mass, its mean radius, its most relevant gravity coefficients and the orientation parameters of its rotation axis with respect to the inertial frame. In order to provide the spacecraft with Pareto Optimal autonomous planning capabilities, the following optimization problem should be solved on-line:

$$\boldsymbol{\pi}^*(s) = \boldsymbol{oe}_{SC}^*(s) = \underset{\boldsymbol{\pi}}{\operatorname{argmax}}\, R(s, \boldsymbol{\pi}) = \underset{\boldsymbol{\pi}}{\operatorname{argmin}}\, J(s, \boldsymbol{\pi}) = \\ = \underset{\boldsymbol{\pi}}{\operatorname{argmin}}\{\boldsymbol{\lambda}\, [J_{\Delta V}(\mathbf{p_{NEA}}, \boldsymbol{\pi}), J_{science}(\mathbf{p_{NEA}}, \boldsymbol{\pi})]^T\}$$

where $\boldsymbol{oe}_{SC}^*$ are the optimal spacecraft orbit elements around the NEA (a way to parameterize its position and velocity). This is a non-convex problem. It entails to run the multi-objective optimizer which embeds simulation of the spacecraft highly non-linear dynamics around the NEA. The considerable computational effort is out of the possibility of any spacecraft CPU. Here comes into play Supervised Learning. By performing *m*-runs of the multi-objective optimizer for various NEA dynamics parameters, a database of paired couples: dynamics parameters-Pareto fronts $(\mathbf{p}_{\mathbf{NEA}}^{(\mathrm{i})}, \mathrm{PF}^{(\mathrm{i})})$ for i = 1,…m, is generated. Each $\mathrm{PF}^{(\mathrm{i})}$ is discretized in n_{PF} Pareto points, each one corresponding to an optimal policy $\boldsymbol{\pi}_k^{*(\mathrm{i})}(\mathbf{p}_{\mathbf{NEA}}^{(\mathrm{i})})$ associated to the trade-off $\boldsymbol{\lambda}_k^{(\mathrm{i})}$, for k = 1,… n_{PF}. After m-runs of the multi-objective optimizer, $n_{PF} \times m$ state-policy pairs are generated as:

$$\left(\left[\mathbf{p}_{\mathbf{NEA}}^{(\mathrm{i})}, \boldsymbol{\lambda}_k^{(\mathrm{i})}\right], \boldsymbol{\pi}_k^{*(\mathrm{i})}\right) \quad \text{for } k = 1,\dots\, n_{PF} \quad \text{for } i = 1,\dots m$$

Results

The multi-objective optimizer has been run $m = 96$ times (~ 48 hours), therefore capturing 96 possible NEA dynamics parameters configurations. Taken uniformly through the whole spectrum of $\mathbf{p_{NEA}}$, 6 sample configurations have been extracted as the development set and 5 as the test set. This results in approximately 3400 state-policies pairs for the training set, 200 for the development set and 200 for the test set. The neural network (NN) weights have been trained on the training set using a mean square error (MSE) loss function. The NN hyperparameters have been tuned and optimized according to the performances provided on the development set. Input and output of the NN have been normalized between 0 and 1. An initial randomized search has made the NN model to converge to a configuration of 4 hidden layers, with tanh activation function for the first 3 layers and sigmoid activation function for the last layer. In this study, the focus is on predicting a subset of the optimal spacecraft orbit elements (optimal policy). In particular, $\boldsymbol{oe}_{SC}^* = [\boldsymbol{a_{SC}}, \boldsymbol{e_{SC}}, \boldsymbol{i_{SC}}]$ is considered, where $\boldsymbol{a_{SC}}$ is the orbit semi-major axis, $\boldsymbol{e_{SC}}$ is the orbit eccentricity and $\boldsymbol{i_{SC}}$ is the orbit inclination. Better results are obtained by training separately three NN each one optimized to predict one specific orbit element. Results of the hyperparameters optimization procedure for the three NN are reported in the following Figures. Finally, results of prediction accuracy on the test set in terms of root mean square error (RMSE) are reported in the following Tables.

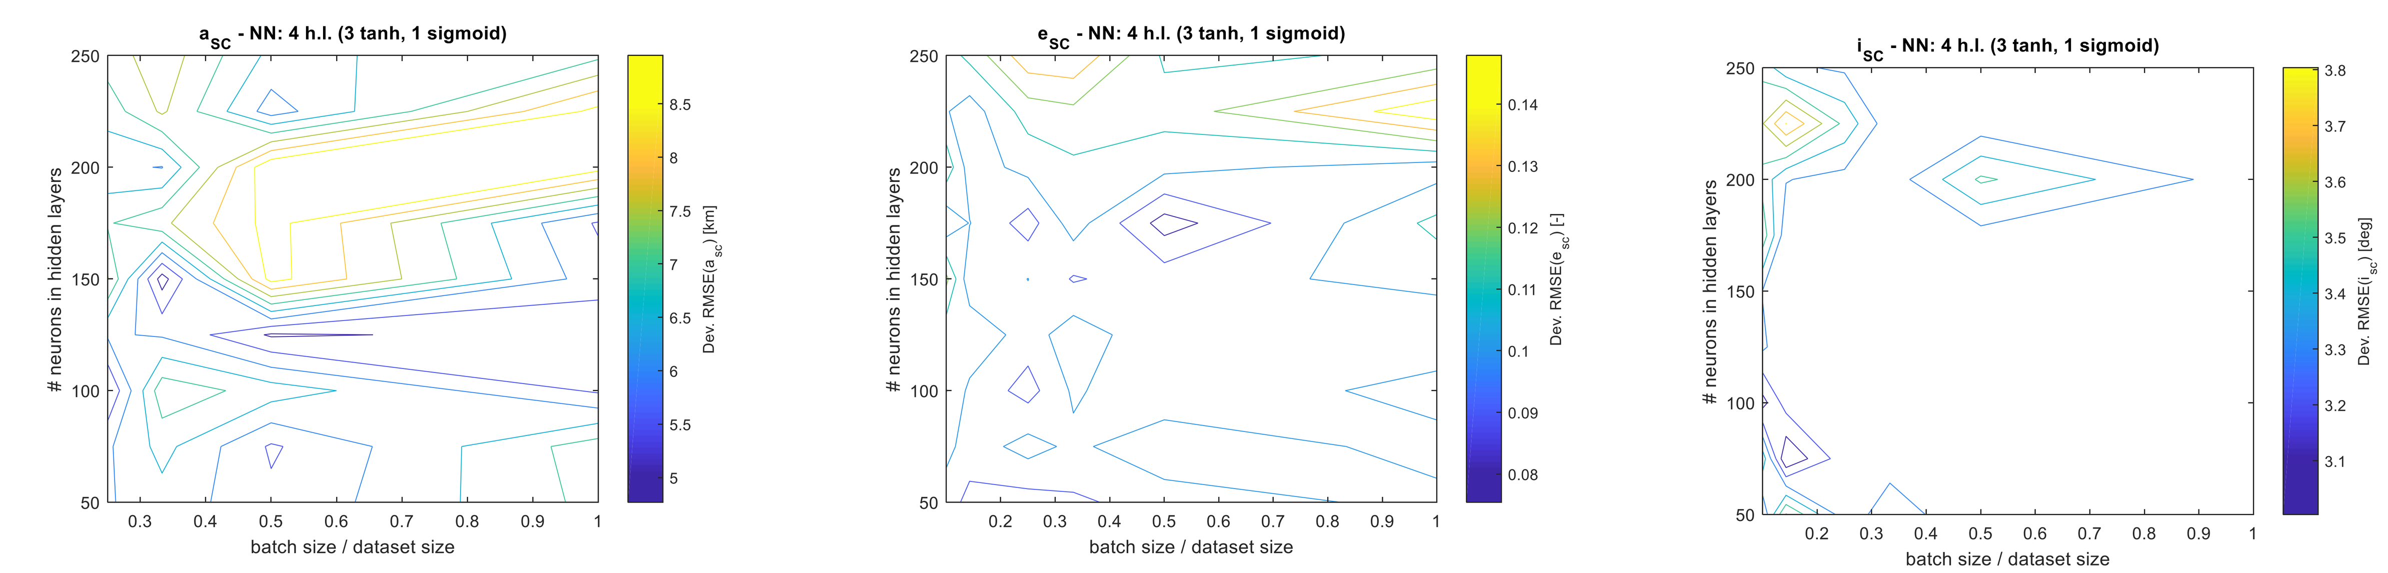

	RMSE_{train}	RMSE_{dev}	RMSE_{test}
a_{SC} [km]	2.36	4.77	3.13

	RMSE_{train}	RMSE_{dev}	RMSE_{test}
e_{SC} [-]	0.044	0.075	0.074

	RMSE_{train}	RMSE_{dev}	RMSE_{test}
i_{SC} [deg]	0.39	3.00	2.43

Multi-objective Motion Planning Dataset

The possible NEA dynamics parameters combinations have been discretized. For each defined combination, an heuristic multi-objective optimization algorithm (e.g. Multi-Objective Particle Swarm Optimization) is used to generate a Pareto front describing the trade-off offered by various motion plans according to two conflicting cost functions. The two cost functions to be minimized provide metric of: 1) the control effort required to realize the motion plan ($J_{\Delta V}$), 2) the inverse of the quality/quantity of scientific output perceivable through realization of the motion plan ($J_{science}$). To identify a specific point on a Pareto front, which corresponds to a specific trade-off between the two conflicting costs, the multi-objective problem can be scalarized as:

$$J = \lambda_1 J_{\Delta V} + \lambda_2 J_{science}$$

where $\lambda_1, \lambda_2 \in (0,1)$ represent weights associated to the two objectives.

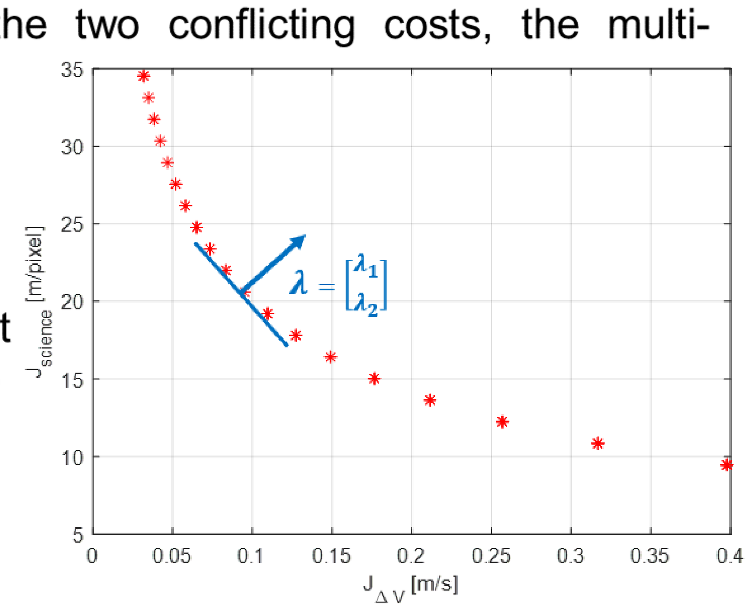

Autonomous motion planning unit

The $n_{PF} \times m$ state-policy pairs can be partitioned in a training set, a development set and a test set. A Neural Network (NN) is trained on the training set to then make predictions on the test set, the NN hyperparameters are optimized on the development set. The trained NN learns the functional relationship between states and policies. The NN can be then implemented on-board the spacecraft and used on-line between the navigation filter and the spacecraft controller as an autonomous motion planning unit (see Figure). This unit takes as input the most recently estimated $\mathbf{p_{NEA}}$ and the selected trade-off $\boldsymbol{\lambda}$, and outputs the optimal spacecraft orbital configuration: $\boldsymbol{\pi}^* = \boldsymbol{oe}_{SC}^*$. This output is passed to the controller which takes action to reach this target configuration.

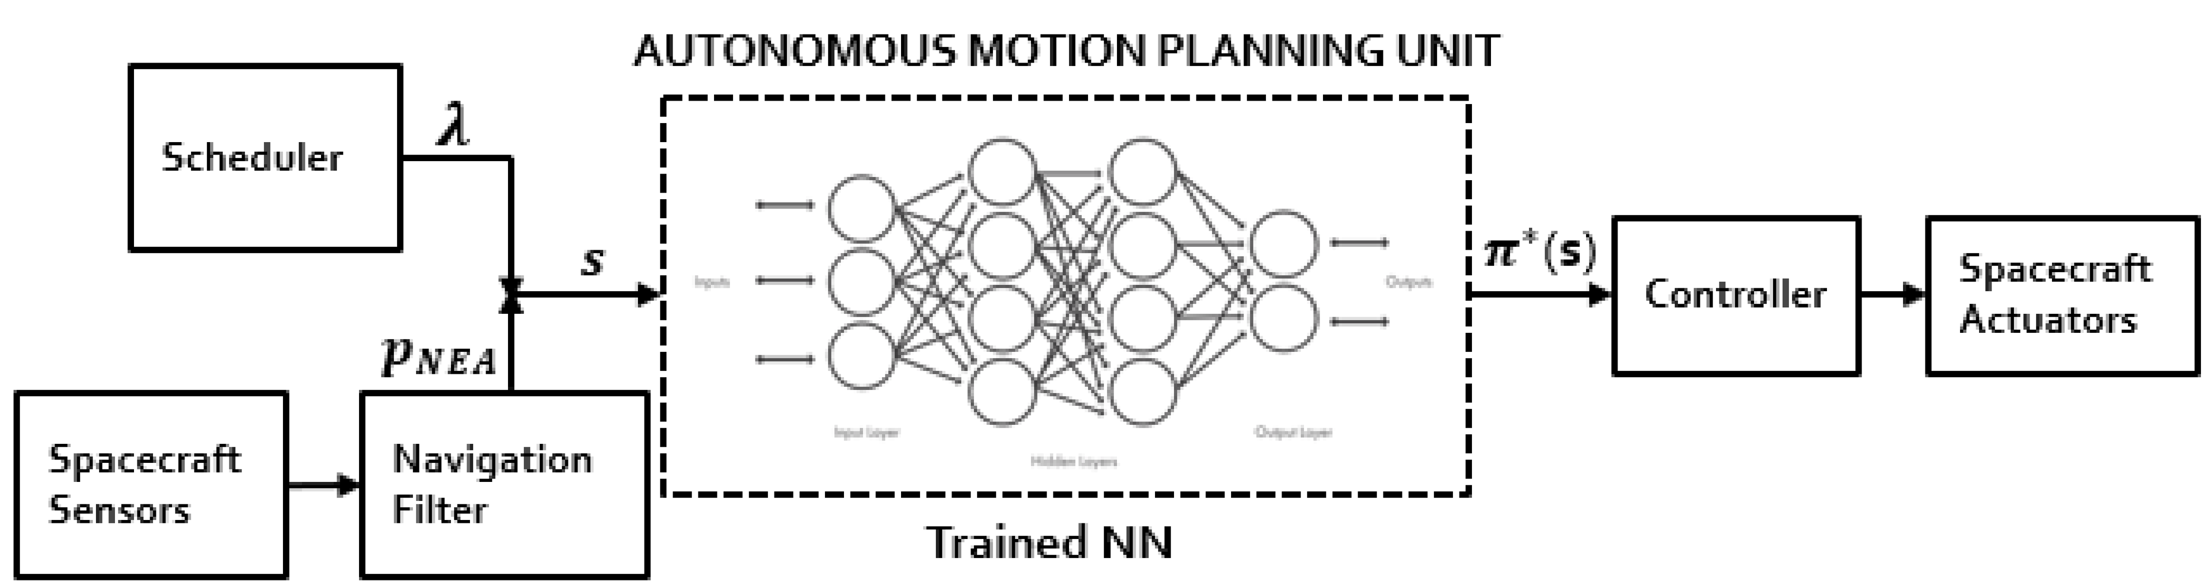

Conclusions

This project explores the use of neural networks (NN) to provide a spacecraft with autonomous multi-objective motion planning capabilities around a near-Earth asteroid (NEA). The trained NN has been shown to provide interesting but still moderate accuracy results. To improve the performances, the first way to follow is to enlarge the dataset, which up to now is limited to only 96 possible NEA dynamics parameters configurations. In addition, future work will explore ways to leverage information about the covariance of the estimated state, that the navigation filter outputs. In this sense, a possible way to go is to reformulate the problem as a stochastic Markov Decision Process.

Acknowledgments

I have developed this project advised by Prof. Simone D'Amico. The topic of this study has been motivated by the on-going research project: "Autonomous Nanosatellite Swarming using Radio Frequency and Optical Navigation", developed at Stanford's Space Rendezvous Laboratory in collaboration with NASA Ames Research Center and sponsored by NASA SSTP (Small Spacecraft Technology Program).