Doctoral Dissertation
Doctoral Program in Aerospace Enginering (37^{th} cycle)

Deep Learning for Space Trajectory Optimization

Andrea Forestieri

* * * * *

Supervisor
Prof. Lorenzo Casalino

Doctoral examination committee
Prof. Nicola Baresi, Referee, University of Surrey
Prof. Hexi Baoyin, Referee, Tsinghua University
Prof. Manuela Battipede, Politecnico di Torino
Prof. Massimo Canale, Politecnico di Torino
Prof. Alessandro Zavoli, Sapienza Università di Roma

Politecnico di Torino
2025

I hereby declare that, the contents and organisation of this dissertation constitute my own original work and does not compromise in any way the rights of third parties, including those relating to the security of personal data.

. . . .

Andrea Forestieri
Turin, 2025

Summary

Space trajectory design and optimization encompass a broad spectrum of challenges, from determining optimal maneuvers to orchestrating complex, large-scale mission planning. This dissertation addresses both ends of this spectrum by integrating deep learning techniques, demonstrating how machine learning can offer efficient, robust, and scalable solutions to problems that have traditionally been approached with handcrafted optimization methods.

In local trajectory optimization, the focus is on low-thrust transfers, a regime where conventional numerical methods often face convergence issues. Neural networks emerge as a powerful alternative when infused with principles from optimal control theory, allowing the learning process to incorporate fundamental dynamical constraints. This integration not only accelerates convergence but also enhances accuracy, making neural networks particularly suitable for solving intricate multi-target low-thrust missions. The results confirm that, even for sequences involving up to fifty consecutive transfers, the total mission time error when using a trained network remains negligible. This level of precision is essential in real-world applications, where even small deviations in transfer times can accumulate, potentially compromising mission feasibility.

These findings highlight how deep learning can bridge the gap between computational speed and solution quality, challenging the notion that fast surrogate models must necessarily sacrifice accuracy. By merging data-driven methodologies with the rigorous mathematical structure of optimal control, this dissertation illustrates how machine learning can effectively complement or even replace traditional approaches in local trajectory design.

Shifting to global trajectory optimization, the focus turns to reinforcement learning, where the principal challenge lies in strategically selecting and sequencing flybys to maximize a given performance index. The proposed algorithms are rigorously benchmarked on the eleventh edition of the Global Trajectory Optimization Competition, an open, highly competitive event within the scientific community. The results are striking: the machine learning-based algorithms outperform multiple top-ranked human-designed solutions, successfully balancing the number of flybys with the fuel consumption required to reach them. Remarkably, these achievements are obtained using a single workstation, underscoring the feasibility of reinforcement learning for

cost-effective, scalable mission planning. The self-play approach used in training further suggests that, with greater computational resources, even stronger solutions could be obtained.

The research presented in this dissertation lays a strong foundation for further exploration and refinement of machine learning in space trajectory optimization.

For local trajectory optimization, future research might focus on expanding the framework to broader mission scenarios. Extending the method to accommodate a wider range of initial conditions, more complex dynamical models, and higher-dimensional state spaces without compromising accuracy or efficiency remains a promising direction. Addressing these aspects would make neural networks even more versatile for practical mission applications.

For global trajectory optimization, the success of reinforcement learning in this study suggests an increasingly prominent role for machine learning in future mission design. A key direction for further research is understanding how the proposed algorithms balance exploration and exploitation. Standard formulations tend to rapidly converge to favored solutions, which, while beneficial in structured environments, can hinder the discovery of equally promising alternatives. Adjusting exploration dynamics to maintain a more balanced search could lead to further improvements in solution quality, particularly in high-dimensional, combinatorial problems such as interplanetary mission planning.

By tackling both local and global trajectory optimization, this dissertation provides a cohesive perspective on how deep learning can be systematically and effectively integrated into mission design. The demonstrated performance gains in accuracy, scalability, and computational efficiency mark a significant step forward, indicating that future space missions, from low-thrust transfers to complex multi-planetary tours, can increasingly leverage data-driven methodologies to achieve more ambitious objectives within tighter operational constraints.