POLITECNICO MILANO 1863

SCUOLA DI INGEGNERIA INDUSTRIALE E DELL'INFORMAZIONE

Executive Summary of the Thesis

Adaptive Model Predictive Control for Autonomous Orbital Station Keeping via Sparse Identification of Nonlinear Dynamics

Laurea Magistrale in High Performance Computing Engineering - Ingegneria del calcolo ad alte prestazioni

Author: Daniele Bergamaschi

Advisor: Prof. Stefano Silvestrini

Co-advisor: Name Surname

Academic year: 2025-2026

1. Introduction

Future deep-space missions require autonomous control systems capable of operating in uncertain and partially known dynamical environments. Traditional space missions rely on centralized architectures [2, 3] in which navigation and control are supported by ground-based infrastructure and analytical models of the surrounding environment. However, communication constraints and limited prior knowledge of the local dynamics can significantly reduce the effectiveness of this approach, motivating the development of systems capable of learning the governing dynamics directly from data while simultaneously computing control actions. This thesis presents an adaptive station-keeping framework that combines Sparse Identification of Nonlinear Dynamics (SINDy) and Nonlinear Model Predictive Control (NMPC). A physics-informed Batch-SINDy procedure is first employed to identify the dominant orbital perturbations from commissioning data, while a Recursive-SINDy formulation continuously updates the model online as new measurements become available. The identified dynamics are subsequently integrated into an NMPC controller to generate optimal control actions under evolving environmental conditions.

The proposed methodology is validated in an Earth-orbiting benchmark scenario and demonstrates the feasibility of combining online model discovery and predictive control within a unified framework. The obtained results show that the proposed approach improves adaptability to model uncertainties while preserving physical interpretability.

2. SINDy and Recursive-SINDy Methods

Sparse Identification of Nonlinear Dynamics (SINDy) is a data-driven model discovery technique that identifies the governing equations of a dynamical system [1] directly from measurements. The method assumes that the observed dynamics can be represented by a sparse combination of candidate nonlinear functions and formulates the identification task as a sparse regression problem.

2.1. Batch-SINDy

In this work, Batch-SINDy is employed to identify the dominant perturbative contributions acting on the spacecraft from orbital data collected during the commissioning phase. The resulting identification problem can be expressed as

$$A \approx \Theta(X)\Xi,$$

where A contains the measured acceleration components, $\Theta(X)$ is the candidate function library, and Ξ is the sparse coefficient matrix. The corresponding optimization problem is formulated as

$$\min_{\Xi} ||A - \Theta(X)\Xi||_2^2,$$

whose solution identifies the dominant perturbative contributions governing the observed dynamics.

2.2. Recursive-SINDy

To enable online adaptation, the standard batch formulation is extended through a Recursive-SINDy approach, which updates the identified model as new measurements become available. The recursive formulation minimizes the exponentially weighted cost function

$$J_k(\Xi) = \sum_{i=1}^{k} \lambda_f^{k-i} ||\mathbf{a}_i - \Theta(\mathbf{x}_i)\Xi||_2^2,$$

where λ_f is the forgetting factor controlling the influence of past observations. The coefficient matrix is then recursively updated according to

$$\hat{\Xi}_k = \hat{\Xi}_{k-1} + \hat{K}_k \left(\mathbf{a}_k - \Theta(\mathbf{x}_k)\hat{\Xi}_{k-1}\right),$$

with $\hat{K}_k$ denoting the recursive gain matrix.
This recursive formulation allows the spacecraft to continuously refine its representation of the surrounding environment while preserving computational efficiency, making it suitable for real-time onboard applications.

3. Proposed Adaptive Framework

The proposed framework is organized into two main phases: an initial commissioning phase and an operational adaptive-control phase.

During the commissioning phase, the objective is to collect orbital data and identify an initial dynamical model through Batch-SINDy. Since the quality of the identified model strongly depends on both the candidate function library and the information content of the available data, a multi-criteria model selection strategy is introduced. Different candidate libraries are evaluated according to both prediction accuracy and numerical conditioning. The selected model must satisfy accuracy requirements while maintaining an acceptable condition number, ensuring robustness, reliability, and physical interpretability.
Once a satisfactory initial model has been identified, the spacecraft enters the operational phase. The identified dynamics are integrated within the NMPC controller and used to compute the initial station-keeping maneuvers. At the same time, a Physics-Informed Recursive-SINDy algorithm is initialized to continuously update the model parameters as new measurements become available.
To guarantee safe adaptation, model updates are not performed continuously. Instead, a hierarchical validation strategy is adopted. First, Recursive-SINDy is activated only after the spacecraft reaches the vicinity of the reference orbit. Subsequently, the updated coefficients must exhibit sufficient stability before being considered for model replacement. The candidate model is then compared against the currently deployed model using prediction-error metrics. A model update is accepted only if the observed improvement remains significant for a predefined period of time.
This approach allows the spacecraft to progressively refine its internal representation of the environment while preventing unnecessary or unstable model updates. As a result, the framework promotes both coefficient convergence and model convergence, enabling reliable long-term autonomous operations under uncertain dynamical conditions.

4. Validation Scenario

Although the proposed framework is motivated by the challenges of deep-space missions and proximity operations around small celestial bodies, the objective of this work is not to reproduce a specific mission scenario. Instead, the goal is

to develop and validate a general adaptive identification and control framework.
For this reason, Earth-centered orbital dynamics are adopted as a benchmark environment. This choice provides a realistic nonlinear dynamical system while simultaneously offering highly accurate reference models and well-known physical parameters, enabling a quantitative assessment of both model identification accuracy and physical interpretability.
The validation scenarios consider the dominant perturbative effects acting on low-altitude Earth-orbiting spacecraft, namely the central gravitational field, the J2 perturbation, and atmospheric drag. In particular, atmospheric drag represents an effective test case for adaptive model learning, since its magnitude varies significantly with orbital altitude and therefore introduces time-varying environmental conditions. Such characteristics make the selected benchmark suitable for evaluating the capability of the proposed framework to identify, adapt, and control uncertain dynamical systems.

5. Results

5.1. Multi-criteria Strategy Validation

The first objective of the experimental analysis is to evaluate the effectiveness of the proposed multi-criteria model selection strategy. The underlying hypothesis is that model quality depends not only on prediction accuracy but also on the information content of the available data and on the numerical conditioning of the identification problem.
To illustrate this aspect, a spacecraft orbiting at an altitude of 900 km with zero inclination was considered. The simulated environment includes the central gravitational field, the J2 perturbation, and atmospheric drag. Tables 1 and 2 compare the coefficients obtained using the model selected by the proposed strategy and those obtained using the complete candidate library. The selected model (Table 1) consists exclusively of the gravitational contribution and produces coefficients close to the normalized theoretical value of the gravitational parameter ($\mu \approx -1$). The small deviation from the theoretical value can be attributed to the correlation between the gravitational and J2 terms in an equatorial orbit, where the available trajectory data provide limited information for distinguishing the two perturbative effects.
In contrast, the model obtained using the complete candidate library (Table 2) distributes the identified dynamics across multiple perturbative terms, including drag and J2 contributions. Although this model achieves a comparable prediction accuracy, the resulting coefficients are difficult to interpret from a physical perspective. In particular, significant drag coefficients are identified despite the limited influence of atmospheric drag at the considered altitude, suggesting that the regression procedure is compensating for correlations among candidate functions rather than isolating the true perturbative contributions.
These results highlight the importance of considering both prediction accuracy and numerical conditioning during the model-selection process. By favoring the simplest model capable of accurately reproducing the observed dynamics, the proposed multi-criteria strategy improves coefficient interpretability while maintaining identification accuracy.

Table 1: Coefficients of the Gravity only Model selected.

Perturbation	Coef. Axis X	Coef. Axis Y	Coef. Axis Z
Gravity	-1.00124619	0	0
	0	-1.00124621	0
	0	0	0

Table 2: Estimated Coefficients of the full Model

Perturbation	Coef. Axis X	Coef. Axis Y	Coef. Axis Z
Gravity	-0.92060427	0	0
	0	-0.34847963	0
	0	0	0
J2	0.01129858	0	0
	0	0.28251460	0
	0	0	0
Drag	0	0.38143617	0
	-0.06299956	0	0
	0	0	0
J3	0	0	0
	0	0	0
	0	0	0
Dummy	0	0	0
	0	0	0
	0	0	0

5.2. Online Model Adaptation

While the previous results focused on the identification and accuracy of the initial model, the following experiments investigate the adaptation capabilities of the proposed framework once the spacecraft enters the operational phase. The analysis is performed for a Very Low Earth Orbit (VLEO) scenario at an altitude of 250 km, where the NMPC controller is activated.

At the beginning of the operational phase, the controller computes the initial control actions using the baseline model identified during the commissioning phase (Table 1). Once the spacecraft reaches the vicinity of the reference orbit, the Recursive-SINDy algorithm is activated and starts updating the model coefficients online. As new measurements become available, the trace of the covariance matrix progressively decreases and the estimated coefficients converge toward their theoretical values. The evolution of the identified coefficients is reported in Figure 1.

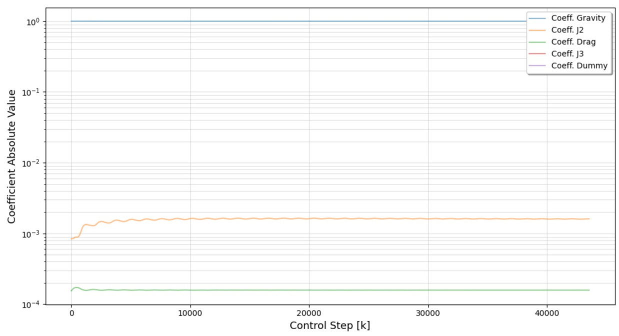

Figure 1: Online evolution of the model coefficients estimated by the Recursive-SINDy algorithm.

Although the coefficients exhibit a clear convergence trend, small oscillations remain present throughout the adaptation process. These fluctuations are primarily caused by variations in orbital altitude, which induce changes in the magnitude of perturbations such as atmospheric drag. As a result, some coefficients converge toward average values rather than fixed constants. Additional numerical noise is also introduced by the normalization and recursive estimation procedures. Nevertheless, the identified parameters remain stable and physically consistent, demonstrating the capability of the framework to adapt its internal model as environmental conditions evolve. The adaptation process starts from an initial model that only includes the gravitational contribution:

$$\boldsymbol{\Xi}^{(0)} = \begin{bmatrix} c_g \\ c_{J2} \\ c_{drag} \\ c_{J3} \\ c_{dummy} \end{bmatrix} = \begin{bmatrix} -1.00124623 \\ 0 \\ 0 \\ 0 \\ 0 \end{bmatrix}. \qquad (1)$$

and progressively converges toward a model in which both the J2 perturbation and atmospheric drag are correctly identified:

$$\boldsymbol{\Xi}^{*} = \begin{bmatrix} -1.00008941 \\ 1.53115145 \times 10^{-3} \\ -1.58438001 \times 10^{-4} \\ 0 \\ 0 \end{bmatrix}. \qquad (2)$$

These results indicate that the Recursive-SINDy algorithm successfully reconstructs the dominant perturbative effects acting on the spacecraft while preserving the sparsity of the identified model.

During the transition from the initial model to the final adapted model, updates are performed only when required by the proposed patience-based validation mechanism. Figure 2 reports the norm of the difference between the accelerations predicted by the active model and those predicted by the updated model. As the recursive identification process progresses, this quantity gradually decreases and eventually stabilizes below the predefined threshold of 5×10^{-7}.

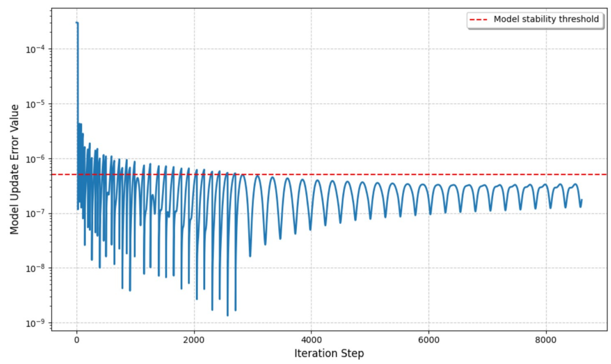

Figure 2: Norm of the difference between the accelerations predicted by the active model and the updated model.

This behavior can be explained by the combined effect of parameter convergence and recursive uncertainty reduction. After each accepted update, the deployed model becomes progressively more accurate, while the covariance trace continues to decrease, leading to increasingly smaller coefficient variations. Consequently, the prediction-error difference between consecutive models becomes negligible, preventing unnecessary model replacements. It is important to note that the number of model updates is strongly influenced by the selected stability threshold. Lower thresholds generally result in a larger number of updates and potentially higher model accuracy, at the expense of increased computational cost. Conversely, excessively small thresholds may cause the update mechanism to react to numerical fluctuations rather than genuine model improvements. The proposed validation strategy therefore provides a practical compromise between adaptation capability, computational efficiency, and robustness against numerical noise.

5.3. Adaptive vs Fixed Model

The final experiment evaluates the practical benefits of online model adaptation by comparing the proposed adaptive SINDy–NMPC framework against a baseline approach based on a fixed Batch-SINDy model identified during the commissioning phase. Tables 3 and 4 summarize the obtained results.

Table 3: Performance metrics of the fixed Batch-SINDy model.

Metric	Value
Delta-V (km/s)	0.485776
Mean Position Error (m)	203.34
Max Position Error (m)	204.65
Avg. NMPC Time (ms)	20.0
Model Discovery Time (s)	2.5
Dataset Size	24002
Memory Usage (MB)	1.65
Simulation Time (h)	51

Table 4: Performance metrics of the adaptive SINDy–NMPC framework.

Metric	Value
Delta-V (km/s)	0.271926
Mean Position Error (m)	10.44
Max Position Error (m)	13.87
Avg. NMPC Time (ms)	47.91
Model Discovery Time (s)	2.5
Dataset Size	24002
Memory Usage (MB)	1.65
Simulation Time (h)	51

Although the fixed model reduces the average NMPC execution time by approximately 58%, this computational advantage is achieved at the expense of significantly lower control performance. In particular, the station-keeping ΔV increases by approximately 79%, while the mean position error grows from 10.44m to 203.34m. These results demonstrate that the main benefit of the proposed framework lies in its ability to preserve model fidelity throughout the mission. By continuously updating the predictive model, the Recursive-SINDy algorithm enables the NMPC controller to maintain higher tracking accuracy while simultaneously reducing propellant consumption. The experiment therefore provides direct quantitative evidence that online adaptation substantially improves the overall performance of autonomous station-keeping operations.

6. Limitations

Despite the promising results, the framework remains dependent on the quality of the available data, the selection of several user-defined thresholds, and the computational cost associated with both the multi-criteria commissioning phase and the online NMPC model updates. Furthermore, the methodology has been validated only in simulation and should be further assessed in more realistic operational scenarios.

7. Conclusions

This work presented an adaptive framework that combines Sparse Identification of Nonlinear Dynamics (SINDy), Recursive-SINDy, and Nonlinear Model Predictive Control for autonomous orbital station-keeping in uncertain dynamical environments.

The results demonstrated that the proposed multi-criteria commissioning strategy improves the physical interpretability, robustness, and reliability of the identified models by jointly considering prediction accuracy and numerical conditioning during the model-selection process. Although the computational cost of the commissioning phase increases with the number of candidate library combinations, the resulting models provide a more accurate and physically meaningful representation of the underlying dynamics.

Furthermore, the online adaptation mechanism based on Recursive-SINDy successfully updates the predictive model during mission operations, enabling the controller to account for previously neglected perturbative effects. The comparison

with a fixed Batch-SINDy model showed that maintaining an updated dynamical model significantly improves tracking performance and reduces propellant consumption, at the cost of increased computational complexity. Nevertheless, the obtained execution times remain compatible with real-time applications, particularly for electric-propulsion spacecraft characterized by relatively slow dynamical evolution.
Overall, the presented results demonstrate the feasibility of integrating sparse model discovery, online adaptation, and predictive control within a unified autonomous framework. Future developments should focus on reducing the dependence on user-defined thresholds, improving the scalability of the commissioning procedure, and extending the validation to more realistic deep-space scenarios.

8. Acknowledgements

Un sentito ringraziamento al mio relatore, Prof. Stefano Silvestrini, per la disponibilità, il supporto e la guida forniti durante tutto lo sviluppo di questo lavoro.
Ringrazio inoltre la mia famiglia e i miei amici, che mi sono stati accanto durante questi anni, contribuendo non solo al mio percorso accademico ma anche alla mia crescita personale.

References

[1] Steven L. Brunton, Joshua L. Proctor, and J. Nathan Kutz. Discovering governing equations from data by sparse identification of nonlinear dynamical systems. *Proceedings of the National Academy of Sciences*, 113(15):3932–3937, 2016.

[2] Todd A. Ely, Jill Seubert, Nicholas Bradley, Ted Drain, and Shyam Bhaskaran. Radiometric autonomous navigation fused with optical for deep space exploration. *The Journal of the Astronautical Sciences*, 68(2):300–325, 2021. NASA Jet Propulsion Laboratory (JPL), California Institute of Technology.

[3] Vittorio Franzese and Francesco Topputo. Deep-space optical navigation exploiting multiple beacons. *The Journal of the Astronautical Sciences*, 69:368–384, 2022. Politecnico di Milano, Department of Aerospace Science and Technology.