International Astronautical Congress, Sydney, Australia, 29 Sep-3 Oct 2025.

IAC-25-C.1.3.5.x101635

Low-thrust trajectory optimization via Covector Mapping Principle: bridging direct and indirect methods for spacecraft autonomous guidance

Niccolò Michelotti[a]*, **Carmine Giordano**[a]†, **Francesco Topputo**[a]‡

[a] *Department of Aerospace Science and Technology, Politecnico di Milano, Via La Masa 34, 20156, Milano, Italy.*

* PhD Student, niccolo.michelotti@polimi.it, corresponding author
† Assistant Professor, carmine.giordano@polimi.it
‡ Full Professor, francesco.topputo@polimi.it

Abstract

Low-thrust spacecraft trajectory optimization is an established problem in applied mathematics and space mission design. The current state-of-the-art methods for solving it primarily fall into two categories: indirect and direct methods. This work investigates the Covector Mapping Principle in order to bridge these two approaches, allowing the transformation of NLP Lagrange multipliers into costates. Pseudospectral direct collocation and indirect shooting methods are implemented, and costate mapping is verified with excellent accuracy for different problem formulations. By combining these two methodologies, we mitigate their respective limitations and pave the way for a robust, flexible, and computationally efficient hybrid guidance algorithm for onboard spacecraft applications. Benchmarks against traditional indirect methods are performed to demonstrate that the proposed hybrid approach improves convergence and robustness. The algorithm is tested and simulated in different deep-space mission scenarios, highlighting its potential for onboard implementation in autonomous spacecraft. The hybrid algorithm is deployed on relevant hardware for processor-in-the-loop experiments to evaluate its performance under realistic constraints and limitations imposed by the onboard computational resources.

1. Introduction

Over the last few decades, the space industry has witnessed disruptive innovations in space exploration. Among them, smaller and more affordable satellites, known as CubeSats, have given universities and small companies the opportunity to reach orbit due to their low design, manufacturing, and launch costs. So far, their use has been mostly limited to near-Earth missions, however, there's growing interest in taking low-thrust CubeSats into deep-space. To overcome the traditional high costs of infrastructure and personnel associated with the ground-based control of interplanetary missions, CubeSats will need to carry out more operations on their own, relying on onboard autonomous guidance, navigation, and operations. Developing these capabilities could pave the way to a new era of exploration, where swarms of small, self-sufficient spacecraft travel the Solar System, collecting data and performing tasks that were once only possible for large, expensive probes. In this context, the ERC-funded *EXTREMA* project [1] aims to enable self-driving spacecraft, proposing a paradigm shift of guidance, navigation, and control operations from ground to onboard Cubesats. In particular, the development of autonomous guidance algorithms will have to focus on guaranteeing robustness, efficiency, and optimality, to overcome the challenges of space exploration with limited onboard resources [2, 3] and within environments often characterized by complex dynamics and high uncertainties [4].

Low-thrust spacecraft trajectory optimization requires solving nonlinear optimal control problems, which are a well-established and challenging class of problems in applied mathematics [5, 6]. They have been extensively studied over the past century due to their broad applicability in numerous scientific and engineering disciplines, including aerospace engineering.

In this context, the current state of the art can be mainly classified into two categories, direct and indirect methods [7]. Direct ones consist in transcribing the original optimal control problem into a discrete-time Nonlinear Programming Problem and then solving it by finding the optimality conditions, also known as Karush-Kunh-Tucker conditions [8]. On the other hand, indirect methods directly solve for the optimality conditions derived from the continuous-time OCP by applying the Pontryagin Maximum Principle, and then using numerical methods to find the results [5]. The two approaches have been extensively implemented and tested in several applications related to low-thrust trajectory optimization, and both showed advantages and disadvantages depending on the specific application. Direct methods have wider convergence regions but usually require a large computational effort due to the solution of high-dimensional nonlinear programming problems [9]. In addition, they do not provide any hints on the verification of the optimality of the original continuous-time problem. Indirect methods, instead, ensure both primal and dual feasibility and optimality, however, they are characterized by limited convergence due to the high sensitivity of the costate and their limited physical insights [10].

This work explores the relationship between direct and indirect approaches to the solution of OCP, exploiting the theory of the Covector Mapping Principle, which has been formulated in the early 2000s [11, 12]. We analyze this relationship applied to low-thrust spacecraft trajectory optimization within the context of on-board autonomous guidance, where optimality, robustness, and computational efficiency become of paramount importance.

In the literature, previous works have implemented hybrid strategies that leverage the results of the Covector Mapping Principle in order to solve trajectory optimization problems. Guo et al. [13] implemented the direct PS collocation to solve the low-thrust spacecraft energy-optimal problem and used the solution as an initial guess for the fuel-optimal indirect shooting. Spada et al. [14] implemented a direct-indirect hybrid strategy for fuel-optimal powered descent guidance using pseudospectral convex methods as a first step and then refined the solution with indirect shooting. Busi et al. [15] explored the relationship between direct and indirect solutions using non-complete methods for energy-optimal problems, however, the mapping of the Lagrange multipliers costates was affected by an unexplained scaling factor. This work further explores the application of the Covector Mapping Theorem and costates mapping in the context of low-thrust spacecraft trajectory optimization, providing a numerical in-depth analysis of its advantages compared to traditional guidance methods, and highlighting the potential for onboard applications.

Section 2 summarizes the theory of the Covector Mapping Principle and its application to nonlinear optimal control problems through the implementation of pseudospectral schemes. In Section 3, the mapping between the costates obtained with the indirect method and the Lagrange multipliers of the direct method is verified. Section 4 presents a hybrid direct-indirect algorithm, which is tested over a dataset of hundreds of interplanetary scenarios. Section 5 discusses the performance of the hybrid algorithm tested and compares it with that of an indirect approach. Section 6 presents some considerations about the deployment of the algorithm on relevant hardware to test and assess the performance for autonomous on-board guidance purposes. Conclusions of the work are drawn in Section 7.

2. State-of-the-Art

2.1 The Covector Mapping Principle

An optimal control problem for low-thrust trajectory optimization can be formulated as follows [6]. The objective is to determine the state $\boldsymbol{y}(t) \in \mathbb{R}^n$ and the control $\boldsymbol{u}(t) \in \mathbb{R}^m$ that optimize a performance index while satisfying the physical and dynamics constraints of the system. This translates as minimizing the cost functional.

$$J = \Phi(\boldsymbol{y}(t_0), t_0, \boldsymbol{y}(t_f), t_f) + \int_{t_0}^{t_f} g(\boldsymbol{y}(t), \boldsymbol{u}(t), t) dt, \quad (1)$$

subject to the dynamics constraints

$$\dot{\boldsymbol{y}}(t) = \boldsymbol{f}(\boldsymbol{y}(t), \boldsymbol{u}(t), t), \quad (2)$$

the boundary conditions

$$\phi(\boldsymbol{y}(t_0), t_0, \boldsymbol{y}(t_f), t_f) = 0, \quad (3)$$

and the inequality path constraints

$$\boldsymbol{C}(\boldsymbol{y}(t), \boldsymbol{u}(t), t) \leq 0. \quad (4)$$

This corresponds to the original continuous-time OCP, which is denoted as Problem P [11]. By deriving the optimality conditions and applying the Pontryagin Maximum Principle (PMP), one obtains the two-point boundary value problem, which is identified as Problem P^{λ}. It is twice the dimensions of the original problem due to the introduction of the dual variables (i.e. costates) whose dimension is twice the dimension of the state of the original problem. For this reason, the process from Problem P to P^{λ} is denoted as dualization. Typically, P^{λ} is not solvable in closed-form and, therefore, an approximate solution has to be found through numerical methods. In this way, Problem $P^{\lambda N}$ is defined, where N represents the discretization of the problem. The equivalence between $P^{\lambda N}$ and P^{λ} is related to the convergence of the approximation implemented to solve the two-point boundary value problem. This way of approaching the solution of an OCP is generally referred to as the *indirect approach*.

The solution of an OCP can also be tackled the other way around. The continuous-time Problem P is first discretized in a discrete-time grid to obtain the Problem P^N.

This corresponds to the transcription of the optimal control problem into a Nonlinear Programming Problem (NLP). The NLP described by Problem P^N can now be solved with numerical methods which exploit the definition of the Karush-Kunh-Tucker (KKT) optimality conditions. In this way, the formulation of Problem $P^{N\lambda}$ is obtained. In literature, this way of solving the OCP is referred to as the *direct approach*.

In general, the KKT optimality conditions of an NLP do not generate discrete adjoint equations that are consistent with the one generated through the optimality conditions and PMP of the original continuous-time OCP, and, consequently, the KKT multipliers are not equal to the costates [8]. In other words, Problem $P^{\lambda N}$ and Problem $P^{N\lambda}$ are not equivalent and, therefore, the operations of discretization and dualization (i.e. derivation of the optimality conditions) are not commutative.

However, it has been proved that the equivalence between these two problems can be made true by a correct choice of the discretization scheme and a proper coordinate transformation. As a consequence, the *non-commutation gap* can be closed through the introduction of a mapping function Γ^N between the solutions of Problem $P^{\lambda N}$ and Problem $P^{N\lambda}$ [11]. This is the essence of the Covector Mapping Principle (CMP), whose mathematical formulation is often referred to as Covector Mapping Theorem (CMT) and depends on the discretization adopted. A schematic of the complete CMP and of the two approaches to solve OCP is shown in Figure 1. Methods that satisfy the CMT are denoted as *complete methods*.

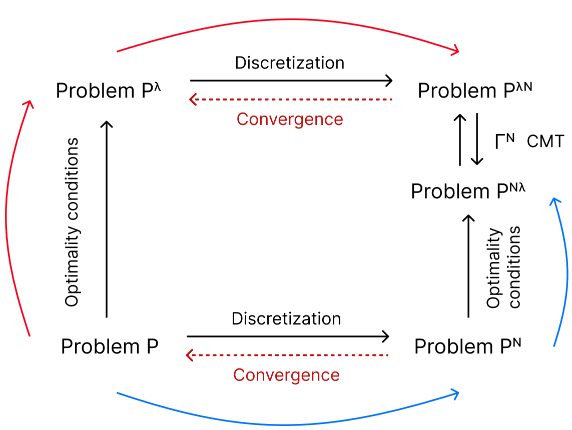

Fig. 1: Schematic of the Covector Mapping Principle [11]. Red arrows are related to the indirect approach, while the blue ones are related to the direct approach.

2.2 LGR pseudospectral method and CMT

Pseudospectral (PS) methods are among the discretization schemes that have been proved to satisfy the Covector Mapping Theorem [12], and therefore are complete methods. They have been widely implemented in the solution of OCP due to their convergence with quasi-exponential rate, in the case of smooth enough solutions [16, 17]. Therefore, fewer discretization points are required with respect to other transcription methods. In addition, the implementation of local collocation PS schemes increases the sparsity of Jacobian matrices and reduces the overall computational effort [18]. Figure 2 shows a comparison of the discretization and collocation points for the most common families of PS methods.

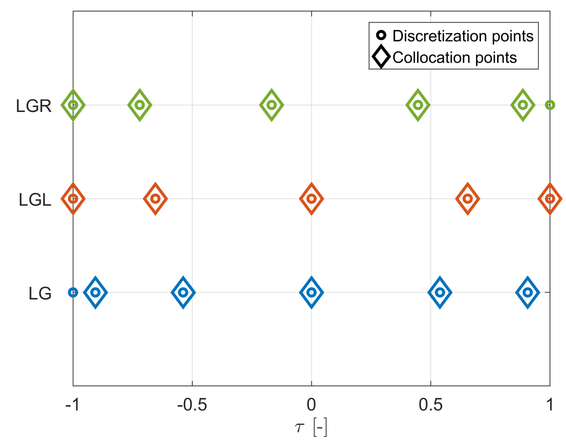

Fig. 2: Comparison of discretization and collocation points for Legendre-Gauss (LG), Legendre-Gauss-Lobatto (LGL) and Legendre-Gauss-Radau (LGR) PS methods, for $N = 5$ collocation points.

Among the different families of PS methods, this work exploits the Legendre-Gauss-Radau (LGR) discretization, which has been proven to accurately map KKT multipliers of the NLP with the costates of the continuous-time optimal control problem [17]. LGR dicretization points $(\tau_1, \ldots, \tau_{N+1})$ are defined in the normalized time interval $[-1; +1]$ such that $\tau_1 = -1$ and $\tau_{N+1} = +1$ (which is not a collocation point). Therefore, the physical time domain $[t_0; t_f]$ has to be mapped to $[-1; +1]$ with the following transformation

$$t = \frac{t_f - t_0}{2}\tau + \frac{t_f + t_0}{2} \tag{5}$$

The LGR collocation points $(\tau_1, \ldots, \tau_N)$ correspond to the roots of the equation

$$P_{N-1}(\tau) + P_N(\tau) = 0 \tag{6}$$

States and control variables are approximated by a summation of Lagrange polynomials $L_i(\tau)$ of degree N over the LGR discretization points. Note that the control at the final point $\tau = 1$ is not provided. States and control variables are expressed by Lagrange polynomial approximations

$$\boldsymbol{y}(\tau) \approx \boldsymbol{Y}(\tau) = \sum_{i=1}^{N+1} \boldsymbol{Y}_i L_i(\tau) \tag{7}$$

$$\boldsymbol{u}(\tau) \approx \boldsymbol{U}(\tau) = \sum_{i=1}^{N} \boldsymbol{U}_i L_i(\tau) \tag{8}$$

where $\boldsymbol{Y}_i = \boldsymbol{y}(\tau_i)$, $\boldsymbol{U}_i = \boldsymbol{u}(\tau_i)$, and $L_i(\tau)$ are the Lagrange polynomials defined as

$$L_i(\tau) = \prod_{\substack{j=1 \\ j\neq i}}^{N} \frac{\tau - \tau_j}{\tau_i - \tau_j} \tag{9}$$

By differentiating Eq. 7, the state dynamics can be imposed at the collocation points through the dynamics defects, expressed as

$$\sum_{i=1}^{N+1} D_{ki}\boldsymbol{Y}_i - \frac{t_f - t_0}{2}\boldsymbol{f}(\boldsymbol{Y}_k, \boldsymbol{U}_k, \tau; t_0, t_f) = \boldsymbol{0}, \quad k = 1, ..., N \tag{10}$$

where $D_{ki} = \dot{L}_i(\tau_k)$ is the element of the Radau pseudospectral differentiation matrix $\boldsymbol{D}$. Note that, since the state approximation uses $N+1$ points and the collocation is done at only the N LGR points (i.e. $\tau = 1$ is excluded), the matrix $\boldsymbol{D}$ has dimensions N x $(N+1)$. Discretization must also be performed on the integral term of the cost function. Using the LGR quadrature rule, the approximated cost function is

$$\begin{aligned} J =& \Phi(\boldsymbol{Y}(\tau_1), \tau_1, \boldsymbol{Y}(\tau_{N+1}), \tau_{N+1}) + \\ &+ \frac{t_f - t_0}{2}\sum_{k=1}^{N} w_k g(\boldsymbol{Y}_k, \boldsymbol{U}_k, \tau; t_0, t_f) \end{aligned} \tag{11}$$

where w_k are the quadrature weights associated with the k^{th} LGR collocation point and are computed as

$$w_1 = \frac{2}{N^2}, \quad w_k = \frac{1}{N^2}\frac{1-\tau_k}{(L_{N-1}(\tau_k))^2}, \quad k = 2, ..., N \tag{12}$$

In particular, it is possible to analytically demonstrate the existence of a Covector Mapping Theorem (CMT) for the LGR pseudospectral discretization by deriving the optimality conditions of the continuous-time and discrete-time formulations of the OCP and comparing the two of them. The analytical manipulation of optimality conditions, which is extensively derived in [19], leads to the definition of the CMT for a Radau PS method with N+1 discretization points, which is

$$\begin{aligned} \widetilde{\boldsymbol{\lambda}}_{N+1} &= \boldsymbol{D}_{N+1}^{\top}\boldsymbol{\Lambda} \\ \widetilde{\boldsymbol{\lambda}}_k &= \frac{\boldsymbol{\Lambda}_k}{w_k}, \quad 1 \le k \le N \\ \widetilde{\boldsymbol{\gamma}}_k &= \frac{\boldsymbol{\Gamma}_k}{w_k}, \quad 1 \le k \le N \\ \widetilde{\boldsymbol{\psi}} &= \boldsymbol{\Psi} \end{aligned} \tag{13}$$

The left-hand sides of Eq. 13 are the estimates of the Lagrange multipliers associated with the continuous-time OCP. In particular, $\boldsymbol{\lambda_k}$, known also as costates, are related to the dynamics constraints, $\boldsymbol{\gamma}_k$ are related to the path constraints, and $\boldsymbol{\psi}$ is related to the boundary constraints. The right-hand sides of Eq. 13 contain instead the respective Lagrange multipliers associated with the KKT conditions of the NLP problem, namely $\boldsymbol{\Lambda_k}, \boldsymbol{\Gamma}_k, \boldsymbol{\Psi}$. Lastly, w_k are the quadrature weights associated with the k^{th} LGR collocation point and $\boldsymbol{D}_{N+1}$ is the last column of the differentiation matrix D of the LGR discretization. In summary, the CMT in Eq. 13 proves the existence of an analytical map between the solution of the original continuous-time OCP and its discretized version.

3. Costates estimation in low-thrust trajectory optimization

The validity of the Covector Mapping Theorem, presented in Section 2, is numerically verified in the solution of trajectory optimization problems for interplanetary transfer with low-thrust spacecrafts. In particular, direct and indirect independent approaches are used to solve the OCP and compared to each other. The direct method is implemented through a Radau pseudospectral collocation, based on the theory reported in Section 2. In addition, a local collocation approach with hp-adaptive mesh refinement has been implemented in order to handle the solution of non-smooth problems, by varying iteratively the segment widths h and the polynomial degree p [18, 20]. An initial guess of the direct optimization is provided with a simple shape-based cubic interpolation of the boundary conditions. The NLP solver used in this analysis is the large-scale interior-point solver IPOPT [21]; however, successful tests have also been carried out with the solver SNOPT [22]. The indirect approach is based on a state-of-the-art shooting algorithm, which allows for correct handling of the bang-bang control profile through the implementation of an accurate switching-detection technique that places an integration node exactly at the switching point. In addition, two continuation methods are used to improve the convergence, namely the energy-to-fuel homotopy and the hyperbolic tangent smoothing (HTS). The indirect shooting algorithm uses initial guesses computed with the adjoint control transformation (ACT) technique. The authors refer to [10, 23] for an exhaustive description of the implementation of the indirect shooting solver and formulation of the problem. In this work, the two-body dynamical model is considered; therefore, the dynamics of state and costates is described as

$$\begin{bmatrix} \dot{\boldsymbol{r}} \\ \dot{\boldsymbol{v}} \\ \dot{m} \\ \dot{\boldsymbol{\lambda}}_r \\ \dot{\boldsymbol{\lambda}}_v \\ \dot{\lambda}_m \end{bmatrix} = \begin{bmatrix} \boldsymbol{v} \\ -\frac{\mu}{r^3}\boldsymbol{r} - u\frac{T_{max}}{m}\frac{\boldsymbol{\lambda}_v}{\lambda_v} \\ -u\frac{T_{max}}{I_{sp}g_0} \\ -\frac{3\mu}{r^5}\left(\boldsymbol{r}\cdot\boldsymbol{\lambda}_v\right)\boldsymbol{r} + \frac{\mu}{r^3}\boldsymbol{\lambda}_v \\ -\boldsymbol{\lambda}_r \\ -\frac{u\lambda_v T_{max}}{m^2} \end{bmatrix} \tag{14}$$

where $\mu = 1.3271 \cdot 10^{11}$ km^3/s^2 is the Sun gravitational constant and $g_0 = 9.81$ m/s^2 is the Earth gravity acceleration. The maximum thrust, specific impulse, and initial mass are representative of a deep-space CubeSat with electric propulsion [24], namely $T_{max} = 3$ mN, $I_{sp} = 3000$ s, and $m_0 = 25$ kg. The scenario considers a low-thrust heliocentric transfer from Earth to Venus with different problem formulations: time-optimal, energy-optimal, and fuel-optimal. The objective functions for the three cases are

$$J = t_f, \qquad \text{for time-optimal} \tag{15}$$

$$J = \frac{1}{T_{max}\ Isp\ g_0}\int_{t_0}^{t_f} \boldsymbol{u}(t)^2 dt, \qquad \text{for energy-optimal} \tag{16}$$

$$J = \frac{1}{Isp\ g_0}\int_{t_0}^{t_f} \boldsymbol{u}(t) dt, \qquad \text{for fuel-optimal} \tag{17}$$

The initial and final boundary constraints are

$$\boldsymbol{r}(t_0) - \boldsymbol{r}_E(t_0) = 0,\ \ \boldsymbol{v}(t_0) - \boldsymbol{v}_E(t_0) = 0,\ \ m(t_0) - m_0 = 0 \tag{18}$$

$$\boldsymbol{r}(t_f) - \boldsymbol{r}_V(t_f) = 0,\ \ \boldsymbol{v}(t_f) - \boldsymbol{v}_V(t_f) = 0, \tag{19}$$

where $(\boldsymbol{r}_E, \boldsymbol{v}_E)$ and $(\boldsymbol{r}_V, \boldsymbol{v}_V)$ are the states of Earth and Venus, respectively. The OCP is solved with the direct and indirect approaches independently, leading to the same results and comparable performance. The CMT in Eq. 13 is then applied to the NLP Lagrange multipliers of the dynamics defects in order to verify the mapping with the costates. For brevity, the results reported below refer only to the energy-optimal and fuel-optimal cases with a time of flight of $TOF = 1000$ days; however, analogous considerations are valid for the time-optimal case. Figures 3 and 4 show the energy-optimal trajectory, mass history, and control profile of the interplanetary transfer obtained with the direct and indirect approaches. The energy-optimal problem is a smooth problem that is easily solved by pseudospectral methods, even with a single global collocation interval. In this case, $N = 40$ collocation nodes are enough for achieving convergence of the results. Figure 5 shows the accurate mapping of the time evolution of the Lagrange multipliers of the dynamics defects into the costates through the CMT.

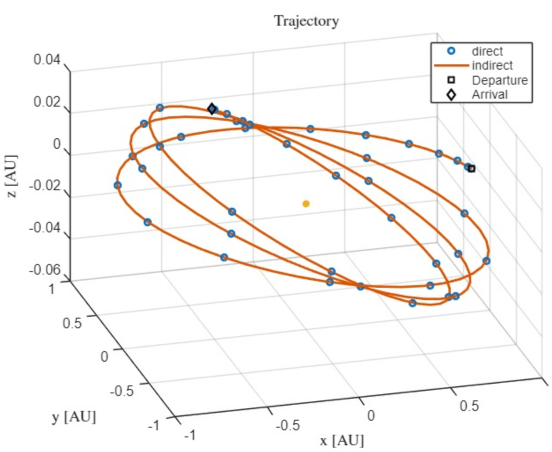

Fig. 3: Trajectory of the energy-optimal Earth-Venus transfer computed with direct and indirect methods (Ecliptic J2000 reference frame centered on the Sun, z-axis not to scale).

The fuel-optimal problem is instead a non-smooth problem, which requires a local collocation approach with a higher number of nodes and, eventually, an adaptive mesh refinement process in order to correctly grasp the bang-off-bang control structure. On the indirect side, energy-to-fuel homotopy is exploited. Figure 6 shows the mass and control history of the fuel-optimal transfer obtained with the two methods. In this case, the PS collocation has been solved with three intervals and a total of 120 nodes. Figure 7 shows that, also in the fuel-optimal case, the estimation

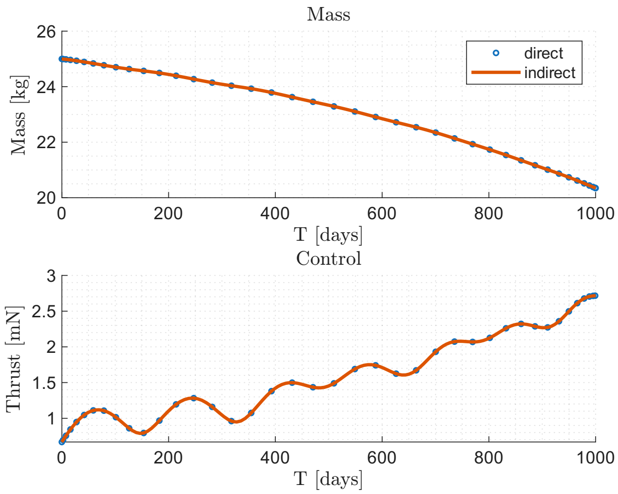

Fig. 4: Thrust profile of the energy-optimal Earth-Venus transfer computed with direct and indirect methods.

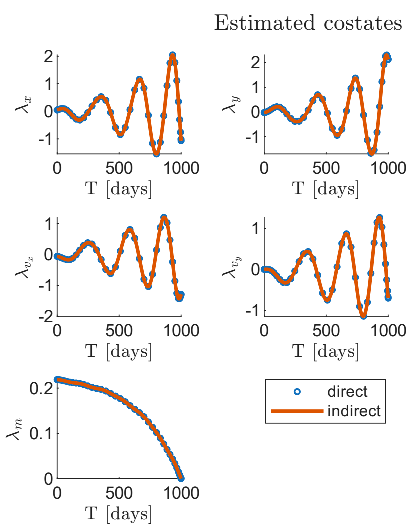

Fig. 5: Comparison of costates and mapped NLP Lagrange multipliers for the energy-optimal Earth-Venus transfer.

of the costates is achieved with excellent accuracy.

While performing this analysis, the main disadvantages of both approaches have been encountered. In particular, the solution of trajectory optimization problems via indirect shooting is highly sensitive to the initial guess, often resulting in convergence to different local optima between different runs. In contrast, the direct method exhibits lower sensitivity to initial conditions. However, it requires solving large-scale NLP problems involving multiple segments and a high number of collocation nodes to adequately capture the bang-off-bang control structure.

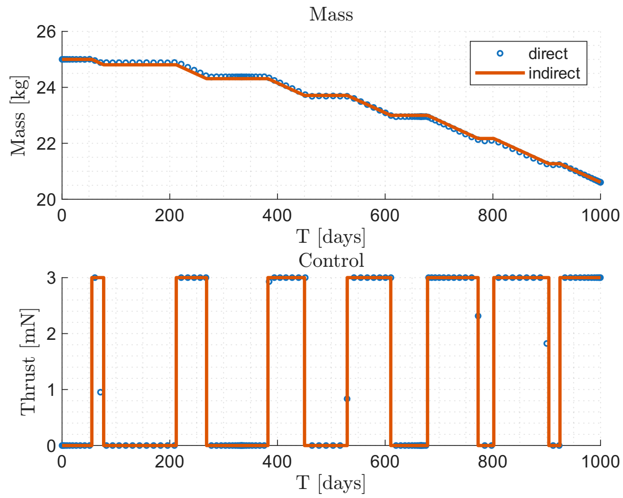

Fig. 6: Thrust profile of the fuel-optimal Earth-Venus transfer computed with direct and indirect methods.

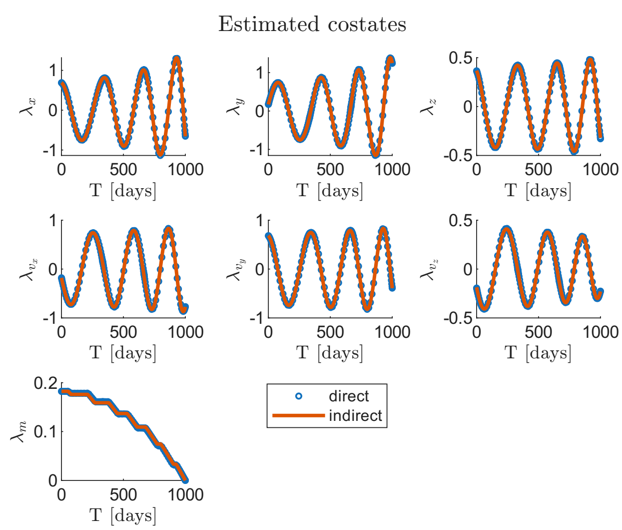

Fig. 7: Comparison of costates and mapped NLP Lagrange multipliers for the fuel-optimal Earth-Venus transfer.

4. Direct-indirect hybrid guidance algorithm

The verification of costates mapping in the context of low-thrust trajectory optimization paves the way for exploiting the relationship between the solutions of the direct and indirect approaches. In particular, it enables the development of hybrid algorithms that leverage the strengths of both formulations while mitigating their respective limitations. The relationship between direct and indirect methods is explored to traverse through the schematic of the CMP in Figure 1 in any direction, allowing transitions between the solutions of both formulations. The results obtained with one approach can be exploited to warm-start the other one and vice versa. For instance, the initial guess for the direct collocation can be provided in terms of trajectory, control history, and, eventually, mapped Lagrange multipliers of the dynamics defects. For the indirect shooting, instead, the initial guess is simply the set of mapped initial costates. In this work, a direct-indirect hybrid algorithm is developed, exploiting the warm-starting of an indirect solver, using the knowledge provided from the solution of the original problem with a direct method. The aim is to solve a minimum-fuel low-thrust trajectory optimization problem with greater robustness and efficiency than traditional guidance methods, while avoiding the disadvantages of approaches that use direct or indirect methods separately.

The proposed hybrid algorithm proceeds as follows:

1. Starting from the original optimal control problem, formulate the transcription of its energy-optimal version [25].

2. Solve the energy-optimal problem using the LGR pseudospectral direct collocation method [26]. Thanks to the quasi-exponential convergence of PS methods for smooth problems, the solution can be efficiently computed with a reduced number of nodes and segments. A shape-based initial guess is provided through a cubic interpolation of the boundary conditions.

3. Exploit the theory of the Covector Mapping Principle to map the Lagrange multipliers of the dynamics defects of the NLP solution to the costates of the continuous-time OCP [12].

4. Use the initial set of mapped costates to warm-start the indirect shooting method. When accurately mapped, this initial guess enables immediate convergence to the energy-optimal solution, often requiring only few iterations for numerical refinement [13].

5. Exploit the energy-to-fuel homotopy within the indirect shooting framework to compute the fuel-optimal solution of the original problem [23].

The outlined hybrid strategy offers several advantages. Firstly, the sensitivity of indirect methods to initial guesses is bypassed through warm starting. In addition, the direct collocation avoids solving the non-smooth fuel optimal problem, which would have required a large number of nodes and eventually a mesh refinement strategy. This would have significantly increased the computational effort compared to the one required to solve a smooth energy-optimal problem. Instead, with this approach, the minimum fuel solution is obtained through a robust and efficient energy-to-fuel homotopy, which identifies the switching instants with excellent accuracy without requiring a thrust regularization process, as in direct collocation approaches [27] Lastly, as the final step of the optimization process is the indirect method, the primal and dual feasibility and optimality of the solution are already verified. Note that this hybrid algorithm is just one possibility for exploiting the connection between the direct and indirect methods via the CMP to solve the OCP. For example, a potential area for future research is to warm-start the direct method using the results of the indirect method.

5. Results discussion

The hybrid algorithm outlined in Section 4 is tested on a large number of scenarios. The dataset of scenarios is composed of randomly generated interplanetary mission scenarios developed for the ERC-funded *EXTREMA* project [1]. The mission scenarios are related to hypothetical deep-space CubeSats with masses ranging from 10 to 25 kg and performing multi-revolution interplanetary transfers from two random points of the Solar System with a time of flight from 1000 to 3000 days [28]. Even though the scenarios are not particularly relevant from a practical point of view, they are useful to test the guidance algorithms under different conditions and to assess their robustness. The hybrid algorithm performance is compared against that of a state-of-the-art indirect shooting method initialized using the ACT-based guess, and exploiting hyperbolic tangent smoothing and energy-to-fuel homotopy. To average the stochasticity of the initial guess of the indirect method, 10 independent runs are conducted for each scenario. The hybrid approach, instead, is fully deterministic, and a discretization with three segments and 20 nodes each is used for each scenario. The results presented below focus on three key performance metrics: robustness of the algorithm, optimality of the solution, and computational efficiency. The complete dataset contains 1000 scenarios. Note that since the scenarios are randomly generated, some are hardly feasible, and therefore, fuel-optimal convergence could become cumbersome for some of them. The hybrid algorithm achieved successful convergence in 80.70% of cases, whereas the mean convergence rate of the indirect method across 10 independent runs was 66.93%.

Figure 8 shows the difference in normalized fuel mass and CPU time between the hybrid and indirect approaches.

Note that the normalized initial wet mass of the spacecraft is equal to one. In the plots, positive values (which are colored in red) indicate scenarios in which the hybrid algorithm achieved lower fuel consumption and shorter computational times with respect to the indirect shooting. Note that mass and CPU time values for the indirect method are related to the mean of the results of 10 runs for each scenario. On average across all scenarios, the hybrid algorithm achieves a 12% reduction in fuel consumption and a 55.6% reduction in CPU time compared to indirect shooting.

It is worth mentioning that the results of indirect shooting are highly sensitive to the initial guess, which contains stochastic terms to allow for different attempts when one fails. In particular, Figure 9 shows the histograms of the standard deviation of normalized fuel mass and CPU time across the 10 independent runs for each scenario. The widespread distribution of standard deviations means that the indirect method is highly likely to converge to one of several local optima, depending on the initial guess. Differently, the hybrid algorithm, which is fully deterministic and has null standard deviation across different runs, is more prone to end up finding the global minimum, as shown in Figure 8a, where the hybrid approach almost always returns a solution with lower mass consumption.

To further show the greater robustness and reliability of the hybrid algorithm compared to the indirect method, Figure 10 compares the optimization time and normalized fuel consumption of 12 randomly sampled scenarios. The blue boxes of the indirect solution summarize the statistics of the results across the 10 runs. Instead, the result of the hybrid algorithm is simply plotted as a red line for each scenario. It is clear that the hybrid algorithm primarily provides a globally optimal solution in a shorter amount of time, without the issues of local optimization and high variability in computational time.

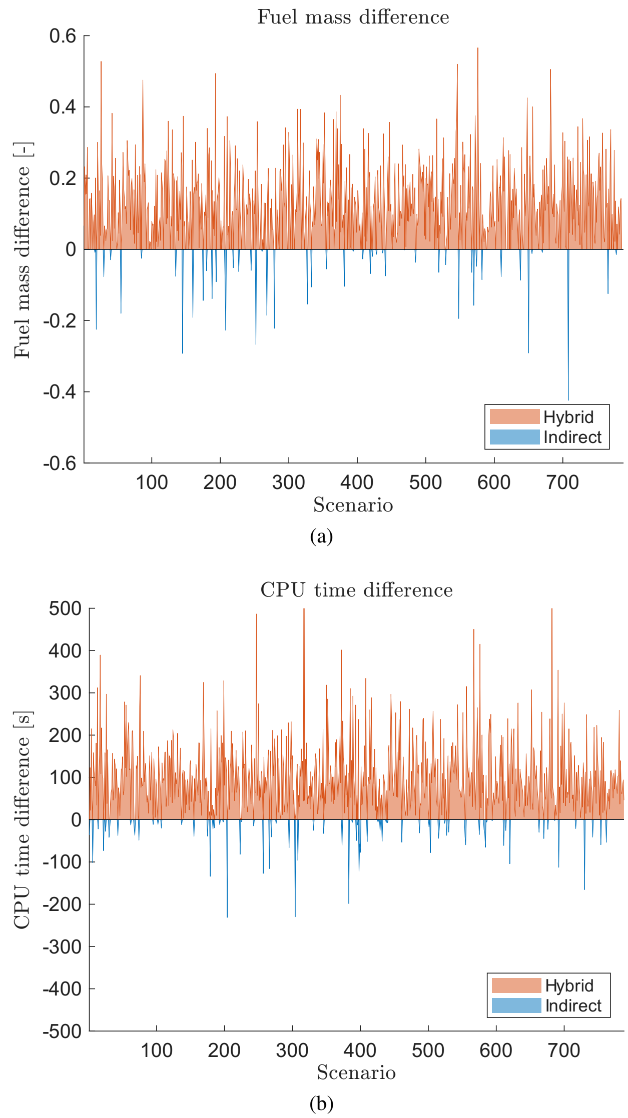

Fig. 8: Difference of normalized fuel mass (a) and CPU time (b) for scenarios converged with both methods. For the indirect method, the values of fuel mass and CPU time are the means of the values obtained in 10 independent runs.

6. Onboard guidance considerations

The development and testing of onboard trajectory optimization algorithms for autonomous systems is among the new challenges in the computational guidance field. The requirements of performance and reliability are often strict and difficult to achieve, especially for indirect methods [29]. The proposed direct-indirect hybrid algorithm demonstrates characteristics and performance capabilities that make it a valuable option for onboard guidance applications. The algorithm allows for finding optimal solutions with high robustness and reduced computational effort. In particular, solving smooth problems with pseudospectral collocation is highly efficient and requires a small number of nodes, which translates to better memory usage. The optimality of the fuel-optimal solution is already verified thanks to the use of indirect shooting as the last step of the process, which directly solves for the necessary optimality conditions of the original problem. In addition, the control profile is ready to be commanded to the thruster without any regularization process, since the switching instants are accurately identified. All these features make

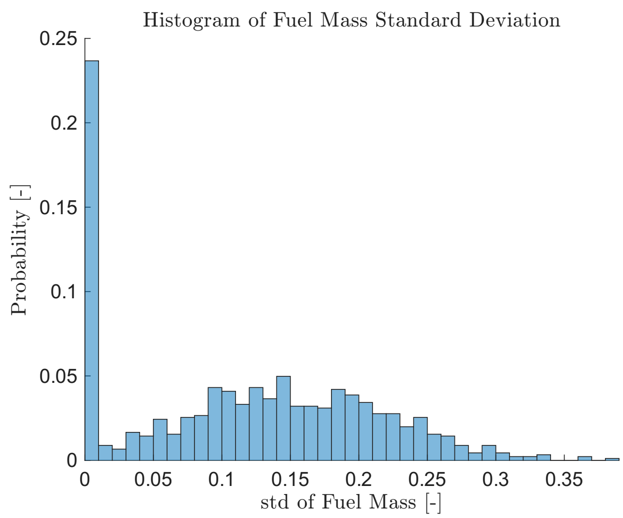

(a)

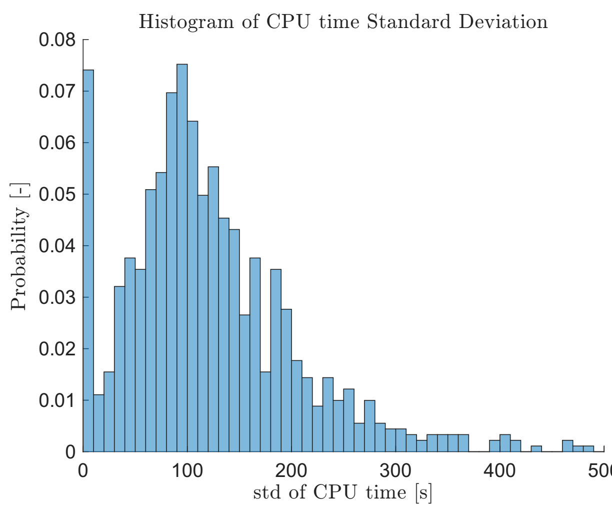

(b)

Fig. 9: Histograms of the standard deviations of the normalized fuel mass and CPU time obtained with the indirect method, across 10 independent runs for each scenario.

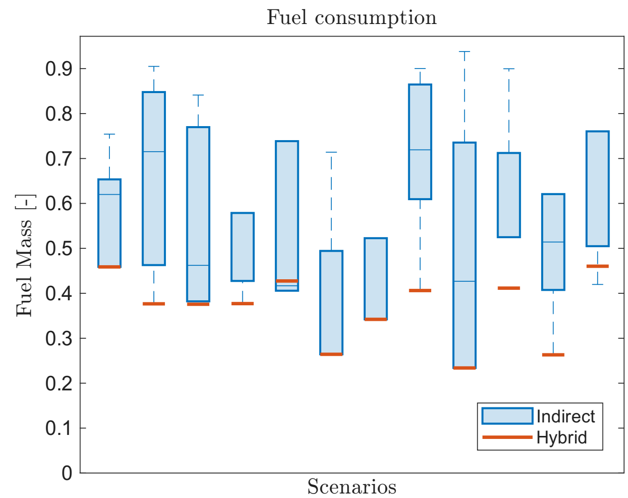

(a)

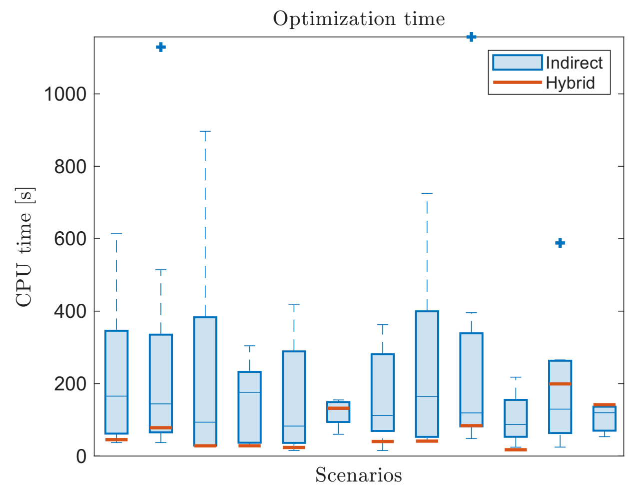

(b)

Fig. 10: Comparison of normalized fuel mass (a) and CPU time (b) for 12 scenarios. The boxes of the indirect solutions show the statistics of the 10 independent runs. The central line represents the median, while the lower and upper edges correspond to the 25th and 75th percentiles. The whiskers extend to the most extreme values that are not classified as outliers, and any outliers are shown separately using the "+" marker. Red lines represent the hybrid solutions.

the hybrid approach a strong candidate for further investigation in the context of on-board autonomous guidance systems. To allow processor-in-the-loop tests on representative hardware of onboard resources, the structure of the hybrid algorithm needs to be modified in order to account for correct memory allocation, configuration files, and correct implementation of the solvers. In particular, the coding of the indirect shooting part follows the one performed in [29], whereas the pseudospectral collocation and the hybrid interface have been deployed from scratch using the C++ generated code from `MATLAB` and integrating the NLP solver SNOPT for C++. The evaluation board selected for the deployment is the AMD Kria KD240 with

Zynq UltraScale+™ MPSoC [1], which is representative of state-of-the-art processors for the next generation of deep-space CubeSats. The accuracy of the results obtained with PIL tests is consistent with the ones obtained on personal computers. However, as expected, the computational time on the evaluation board is significantly higher. In particular, the tests on the Kria board revealed CPU times 7 to 10 times higher than those of a desktop PC with Intel Core i7-10700 CPU running at 2.90GHz and equipped with 16 GB of RAM.

The next steps will be to integrate the deployed algorithm in hardware-in-the-loop test campaigns, within the context of the *EXTREMA* project [1]. The guidance algorithm will interact with the flight software, navigation algorithms and cameras [30], thruster-in-the-loop experiments, and attitude simulation system [31].

7. Conclusions

In this work, the application of the Covector Mapping Principle to low-thrust trajectory optimization for onboard guidance has been explored. The LGR pseudospectral collocation method is implemented to verify the mapping of NLP Lagrange multipliers with the costates of an indirect shooting method. Indeed, the costate estimation has been verified to an excellent degree of accuracy in different scenarios involving time-, energy- and fuel-optimal problem formulations. The advantages and drawbacks of direct and indirect methods are highlighted, and a direct-indirect hybrid approach exploiting the Covector Mapping Theorem is proposed. The hybrid algorithm is tested on a dataset of randomly generated interplanetary scenarios, and its performance is compared with that of a state-of-the-art indirect shooting method. The hybrid approach shows greater robustness, optimality, and computational efficiency, overcoming the main disadvantages of both direct and indirect methods. Its features make it ideal for onboard guidance applications in the next generation of autonomous deep-space CubeSats. The algorithm has therefore been deployed on relevant hardware and tests have been carried out, paving the way for future hardware-in-the-loop test campaigns.

Acknowledgements

This research is part of EXTREMA, a project that has received funding from the European Research Council (ERC) under the European Union's Horizon 2020 research and innovation program (Grant Agreement No. 864697).

[1] `https://www.amd.com/en/products/system-on-modules/kria/k24/kd240-drives-starter-kit.html`. Last access on 11^{th} September 2025.

References

1. Di Domenico, G. *et al.* in *Modeling and Optimization in Space Engineering: New Concepts and Approaches* 167–199 (Springer International Publishing, 2023).
2. Andreis, E., Franzese, V. & Topputo, F. Onboard orbit determination for deep-space cubesats. *Journal of guidance, control, and dynamics* **45,** 1466–1480 (2022).
3. Morelli, A. C., Morselli, A., Perico, D., Topputo, F., *et al.* The EXTREMA Autonomous Guidance Algorithm for Low-Thrust Interplanetary Spacecraft. in *74th International Astronautical Congress (IAC 2023)* (2023), 1–10.
4. Michelotti, N., Rizza, A., Giordano, C. & Topputo, F. Uncertainty Propagation Performance in Proximity Operations Around Small Bodies. *The Journal of the Astronautical Sciences* **71,** 55 (2024).
5. Bryson, A. E. Applied Optimal Control: Optimization, Estimation and Control (Routledge, 1975).
6. Hull, D. G. Optimal control theory for applications (Springer Science & Business Media, 2003).
7. Rao, A. V. A survey of numerical methods for optimal control. *Advances in the astronautical Sciences* **135,** 497–528 (2009).
8. Enright, P. J. & Conway, B. A. Discrete approximations to optimal trajectories using direct transcription and nonlinear programming. *Journal of Guidance, Control, and Dynamics* **15,** 994–1002 (1992).
9. Topputo, F. & Zhang, C. Survey of Direct Transcription for Low-Thrust Space Trajectory Optimization with Applications. *Abstract and Applied Analysis* **2014,** 851720 (2014).
10. Zhang, C., Topputo, F., Bernelli-Zazzera, F. & Zhao, Y.-S. Low-thrust minimum-fuel optimization in the circular restricted three-body problem. *Journal of Guidance, Control, and Dynamics* **38,** 1501–1510 (2015).
11. Ross, I. M. A historical introduction to the Covector Mapping Principle. Calhoun: The NPS Institutional Archive. 2005.

12. Gong, Q., Ross, I. M., Kang, W. & Fahroo, F. Connections between the covector mapping theorem and convergence of pseudospectral methods for optimal control. *Computational optimization and applications* **41,** 307–335 (2008).
13. Guo, T., Jiang, F. & Li, J. Homotopic approach and pseudospectral method applied jointly to low thrust trajectory optimization. *Acta Astronautica* **71,** 38–50 (2012).
14. Spada, F., Sagliano, M. & Topputo, F. Direct–indirect hybrid strategy for optimal powered descent and landing. *Journal of Spacecraft and Rockets* **60,** 1787–1804 (2023).
15. Busi, S., Cunis, T., Topputo, F., Fichter, W., *et al.* Analysis of the Covector Mapping Principle using Non-Complete Methods. in *Euro GNC Conference 2024* (2024), 1–15.
16. Gong, Q., Fahroo, F. & Ross, I. M. Spectral algorithm for pseudospectral methods in optimal control. *Journal of Guidance, Control, and Dynamics* **31,** 460–471 (2008).
17. Garg, D. *et al.* A unified framework for the numerical solution of optimal control problems using pseudospectral methods. *Automatica* **46,** 1843–1851 (2010).
18. Patterson, M. A., Hager, W. W. & Rao, A. V. A ph mesh refinement method for optimal control. *Optimal Control Applications and Methods* **36,** 398–421 (2015).
19. Garg, D. Advances in global pseudospectral methods for optimal control. PhD thesis (University of Florida, FL, 2011).
20. Darby, C. L., Hager, W. W. & Rao, A. V. An hp-adaptive pseudospectral method for solving optimal control problems. *Optimal Control Applications and Methods* **32,** 476–502 (2011).
21. Wächter, A. & Biegler, L. T. On the implementation of an interior-point filter line-search algorithm for large-scale nonlinear programming. *Mathematical programming* **106,** 25–57 (2006).
22. Gill, P. E., Murray, W. & Saunders, M. A. SNOPT: An SQP algorithm for large-scale constrained optimization. *SIAM review* **47,** 99–131 (2005).
23. Wang, Y. & Topputo, F. Indirect optimization of power-limited asteroid rendezvous trajectories. *Journal of Guidance, Control, and Dynamics* **45,** 962–971 (2022).
24. Topputo, F. *et al.* Envelop of reachable asteroids by M-ARGO CubeSat. *Advances in Space Research* **67,** 4193–4221 (2021).
25. Hargraves, C. R. & Paris, S. W. Direct trajectory optimization using nonlinear programming and collocation. *Journal of guidance, control, and dynamics* **10,** 338–342 (1987).
26. Garg, D. *et al.* Direct trajectory optimization and costate estimation of finite-horizon and infinite-horizon optimal control problems using a Radau pseudospectral method. *Computational Optimization and Applications* **49,** 335–358 (2011).
27. Morelli, A. C., Morselli, A., Giordano, C. & Topputo, F. Convex trajectory optimization using thrust regularization. *Journal of Guidance, Control, and Dynamics* **47,** 339–346 (2024).
28. Merisio, G. EXTREMA: Random autonomous interplanetary mission scenarios for EXTREMA Simulation Hub (ESH) hardware-in-the-loop simulations. Version 1.3.0 (Zenodo, Aug. 2023). `https://doi.org/10.5281/zenodo.8300807`.
29. Mannocchi, A., Giordano, C., Topputo, F., *et al.* Closed-loop guidance for interplanetary CubeSats with indirect methods. in *74th International Astronautical Congress (IAC 2023)* (2023), 1–11.
30. Ornati, F., Panicucci, P., Andreis, E., Topputo, F., *et al.* RETINA: a highly-versatile optical facility for camera-in-the-loop testing of spaceborne vision-based sensors. in *46th AAS Guidance, Navigation and Control Conference* (2024), 1–19.
31. Mannocchi, A. *et al.* An Overview of EXTREMA Project Pillar II: Validating Autonomous Guidance Algorithms via HIL Testing. in *75th International Astronautical Congress (IAC 2024)* (2024), 1–11.