

Reachability for Low-Thrust Trajectories via Maximum Initial Mass

Giacomo Acciarini⁽¹⁾, Dario Izzo⁽¹⁾, and Zhong Zhang⁽²⁾

⁽¹⁾Advanced Concepts Team, European Space Agency, European Space Research and Technology Centre (ESTEC), Keplerlaan 1, 2201 AZ Noordwijk, The Netherlands

⁽²⁾Politecnico di Milano, Piazza Leonardo da Vinci 32, Milan, Italy

¹giacomo.acciarini@gmail.com

Abstract - Reachability analysis plays a central role in low-thrust spacecraft trajectory optimization by identifying which target states can be achieved under constraints on time, thrust, and propellant. Traditional approaches construct reachable sets by solving numerous optimal control problems over grids of terminal states, requiring extensive forward simulations with fixed initial conditions. While effective, this method is computationally expensive and becomes impractical for high-dimensional systems or nonlinear dynamics, such as those encountered in cislunar environments or solar sail missions. This work introduces a dual formulation of the reachability problem. Instead of directly computing the full reachable set, we determine, for fixed transfer time and boundary conditions, the maximum allowable initial mass (or, for solar sails, a scalar sail-strength parameter) that enables a successful transfer. A target is deemed reachable if the spacecraft's initial mass does not exceed this threshold. This transforms the problem into a scalar optimization for each target, yielding a well-behaved scalar field that encodes the same feasibility information as the classical reachable set while exhibiting smoother numerical behavior. We leverage indirect maximum-initial-mass (MIM) formulations for both electric low-thrust and solar-sail dynamics and show how they can be used as efficient reachability oracles. Building on this, we then construct data-driven surrogate models that approximate the MIM-based reachability indicator. In particular, we investigate fully connected architectures and show that residual networks provide the best trade-off between accuracy, training stability, and complexity in this

setting. The resulting surrogates enable rapid reachability evaluation while preserving the advantageous numerical properties of the dual formulation, offering a practical tool for preliminary mission design and feasibility assessment.

I. INTRODUCTION

Low-thrust propulsion technologies, such as electric thrusters and solar sails, enable highly efficient spacecraft trajectories at the price of complex, long-duration transfers governed by nonlinear dynamics and tight control constraints. In this context, *reachability* questions, whether a given target state can be achieved within a prescribed time-of-flight, thrust budget, and propellant mass, are central to preliminary mission design, target selection, and trade-space exploration.

Classical reachability analyses for low-thrust systems typically rely on solving many two-point boundary-value problems over grids of terminal states, thereby constructing approximations of the reachable set in the state space [1, 2, 3, 4]. These approaches require a large number of indirect or direct optimal-control solves and quickly become impractical for high-dimensional systems, long time horizons, or complex dynamical environments. Moreover, minimum-time formulations introduce strong nonlinearities and sensitivity near the reachable-set boundary, which can hinder convergence and complicate surrogate modeling. For this reason, there have been works that started proposing approximators of the reachable set [5, 6, 3].

In parallel, recent work has shown the usefulness of dual formulations based on maximum initial mass for low-thrust transfer design [7, 8, 9, 10], as well as the promise of machine-learning surrogates that replace repeated optimal-control solves by fast evaluations. However, an explicit treatment of maximum-initial-mass

(MIM) reachability for general low-thrust systems, its properties, and its interplay with learned surrogates is still lacking.

The present work addresses this gap with two main contributions. First, we formulate reachability for low-thrust trajectories in terms of a maximum initial mass (or, for solar sails, a scalar sail-strength parameter) defined over fixed boundary conditions and time-of-flight. This yields an implicit scalar-field representation of the reachable set, which we analyze and compare to its "dual" minimum-time formulations, highlighting its improved numerical behavior and regularity. These maximum initial mass problems are then solved via indirect methods, leveraging Pontryagin Minimum Principle (PMP) for both low-thrust propulsion and solar-sail dynamics. Second, we construct and evaluate machine-learning surrogates for these dual problems, demonstrating the performances that residual networks with smooth activations (in particular residual networks with **Softplus** units) achieve for MIM-based reachability. Together, these elements provide an efficient and accurate framework for reachability analysis in low-thrust mission design.

II. REACHABILITY IN LOW-THRUST TRAJECTORIES

The setting we study is cases in which the dynamics is driven by a state-dependent dynamics vector, $\mathbf{f}$, and a control-dependent component, $\mathbf{g}$:

$$\dot{\mathbf{x}} = \mathbf{f}(\mathbf{x}(t)) + \mathbf{g}(\mathbf{x}(t), \mathbf{u}) \quad (1)$$

where $\mathbf{x} \in \mathbb{R}^{n_x}$ and $\mathbf{u}$ are the state and control vectors, respectively. In cases where the mass of the spacecraft varies, we can find its time evolution augmenting the system with the following ordinary differential equation:

$$\dot{m} = -\frac{T_{\max}}{I_{sp}g_0} \|\mathbf{u}(t)\|_2, \quad (2)$$

where $\|\mathbf{u}\|_2$ is the l_2 -norm magnitude of the throttle vector, with $T_{\max}$ being the maximum thrust allowable, m_0 is the initial mass, $g_0 = 9.80665 \text{ m/s}^2$ the gravitational acceleration at the sea level of the Earth, and I_{sp} the specific impulse. In our experiments, we will focus on two different scenarios: a low-thrust trajectory optimization case with electric propulsion, where the underlying dynamics is a two-body one, with added acceleration coming from the low-thrust engine, and a solar sailing case, in which there is no mass loss due to propellant, as the thrust is generated leveraging the solar radiation pressure of the Sun. Both these cases will be detailed in Sec. III.A and III.B.

Classically, the reachable set, $\mathcal{R}$, is then defined as the collection of states that can be reached from an initial state and time, with a fixed duration $[t_0, t_f]$, with a control that is admissible (i.e., that satisfies the constraints) $\mathbf{u} \in \mathcal{U}$:

$$\begin{aligned} \mathcal{R}(t_f; \mathbf{x}_0, m_0) := \left\{ \mathbf{x}_f \in \mathbb{R}^{n_x} \mid \exists \mathbf{u}(\cdot) \in \mathcal{U} \text{ s.t.} \right. \\ \dot{\mathbf{x}}(t) = \mathbf{f}(\mathbf{x}(t)) + \mathbf{g}(\mathbf{x}(t), \mathbf{u}(t)), \\ \dot{m}(t) = -\frac{T_{\max}}{I_{sp}g_0} \|\mathbf{u}(t)\|_2, \\ \mathbf{x}(t_0) = \mathbf{x}_0, \quad m(t_0) = m_0, \\ \left. \mathbf{x}(t_f) = \mathbf{x}_f \right\}. \end{aligned} \quad (3)$$

Usually, solving this problem is computationally intensive, as it requires exploring a high-dimensional space of admissible controls and terminal states. In practice, this often involves discretizing the target state space and solving a large number of boundary-value or optimal control problems, each subject to nonlinear dynamics, control bounds, and, potentially, mass depletion effects. Moreover, the structure of $\mathcal{R}$ can be highly nonconvex and sensitive to the choice of final time, making its accurate characterization, especially at the boundaries, very challenging, as discussed profusely in the literature [1, 2, 3, 4]. In fact, the classical reachability formulation reduces the problem to a binary feasibility one, with a hard discontinuous boundary: causing numerical difficulties particularly for the solutions of the two point boundary value problem that lay at the boundary of the set. These limitations encourage the introduction of alternative formulations that shift the focus from explicitly constructing the reachable set to evaluating reachability through scalar quantities, which are possibly defined everywhere and continuous (and differentiable).

Dual formulation via maximum initial mass. Motivated by this, we formulate the problem as a maximum initial mass one, where the objective is, given an initial and final state ($\mathbf{x}_0$ and $\mathbf{x}_f$, respectively) and a transfer window time $[t_0, t_f]$, to find the maximum initial mass that makes the transfer possible:

$$\begin{aligned} m_{\text{MIM}}(\mathbf{x}_0, \mathbf{x}_f, t_f) := \sup \left\{ m_0 \in \mathbb{R} \mid \exists \mathbf{u}(\cdot) \in \mathcal{U} \text{ s.t.} \right. \\ \dot{\mathbf{x}}(t) = \mathbf{f}(\mathbf{x}(t)) + \mathbf{g}(\mathbf{x}(t), \mathbf{u}(t)), \\ \dot{m}(t) = -\frac{T_{\max}}{I_{sp}g_0} \|\mathbf{u}(t)\|_2, \\ \mathbf{x}(t_0) = \mathbf{x}_0, \quad m(t_0) = m_0, \\ \left. \mathbf{x}(t_f) = \mathbf{x}_f \right\}, \end{aligned} \quad (4)$$

where m_f is dictated by the control history and mass-flow law. So instead of fixing m_0 and asking “is $\mathbf{x}_f$ reachable?”, we fix $\mathbf{x}_0, \mathbf{x}_f, t_f$ and ask “what is the largest m_0 that still allows reaching $\mathbf{x}_f$?”. Under *mild* assumptions on $\mathbf{g}$, the latter problem admits a solution, albeit with mass values that may not be feasible with the given spacecraft mass. Note that while this reduces the reachability problem to a threshold condition on the mass, the two problems are formally equivalent, and the function m_{MIM} encodes the same feasibility information as the boundary of the reachable set. This provides an implicit scalar-field representation of the reachable set, more analogous to value-function representations in Hamilton–Jacobi analysis [11, 12, 13]. Level sets of this scalar field correspond to constant mass margins, and the reachable set boundary for a spacecraft of mass m_{sc} is obtained by thresholding at $m_{\text{MIM}} = m_{\text{sc}}$.

Relation to minimum-time reachability. An important observation is that minimum-time solutions also characterize the boundary of the reachable set [14]. In the classical minimum-time setting one fixes $(\mathbf{x}_0, \mathbf{x}_f, m_0)$ and minimizes the final time t_f , obtaining a time-to-reach value function on the target space. This formulation is often numerically more challenging than the maximum-initial-mass approach, due to the high sensitivity of the target state with respect to the optimal transfer time: small changes in the target state can induce large changes in the optimal time-of-flight, especially near the reachable-set boundary. As a result, the time-to-reach function can exhibit sharper gradients and stronger nonsmoothness, and the associated two-point boundary-value problem (TPBVP) is often more prone to convergence issues and ill-conditioning. In contrast, in the present MIM formulation the time horizon is fixed and the unknowns are the initial mass and the costates, rather than the final time and the costates: the shooting function remains of the same dimension but is typically better conditioned because the time-scaling freedom is removed. The control structure remains similar to the time-optimal case, but the scalar quantity of interest is m_{MIM} rather than t_f . Empirically, this dual scalar field varies more smoothly with the target state in practical low-thrust problems [7, 8, 9, 10, 15], which makes it better suited as a regression target and as a basis for surrogate reachability models.

III. THE MAXIMUM INITIAL MASS PROBLEM

A. Electric Low-Thrust Propulsion

We consider the controlled dynamics introduced in Sec. II, where the drift and control-dependent terms are denoted by $\mathbf{f}$ and $\mathbf{g}$, and the spacecraft mass evolves ac-

cording to the low-thrust mass-flow law. The state is augmented with the mass,

$$\mathbf{z}(t) = \begin{bmatrix} \mathbf{x}(t) \\ m(t) \end{bmatrix} \in \mathbb{R}^{n_x+1}, \quad (5)$$

and the dynamics reads

$$\begin{aligned} \dot{\mathbf{r}}(t) &= \mathbf{v}(t) \\ \dot{\mathbf{v}}(t) &= -\mu \frac{\mathbf{r}(t)}{r^3(t)} + \frac{T_{\max}}{m} \mathbf{u}(t), \\ \dot{m}(t) &= -\frac{T_{\max}}{I_{sp} g_0} \|\mathbf{u}(t)\|_2, \end{aligned} \quad (6)$$

where $\mathbf{a}_g(\mathbf{x})$ denotes the gravitational acceleration (two-body in this work) and we used the standard low-thrust relation $\mathbf{g}(\mathbf{x}, \mathbf{u}) = \frac{T_{\max}}{m} \mathbf{u}$, with admissible controls $\mathbf{u}(\cdot) \in \mathcal{U}$ satisfying the thrust magnitude bound $\|\mathbf{u}(t)\|_2 \leq 1$.

For a fixed transfer time $[t_0, t_f]$ and fixed boundary states $\mathbf{x}(t_0) = \mathbf{x}_0, \mathbf{x}(t_f) = \mathbf{x}_f$, the maximum initial mass problem is

$$\begin{aligned} \max_{m_0, \mathbf{u}(\cdot)} \quad & m_0 \\ \text{s.t.} \quad & \dot{\mathbf{r}}(t) = \mathbf{v}(t), \\ & \dot{\mathbf{v}}(t) = -\mu \frac{\mathbf{r}(t)}{r^3(t)} + \frac{T_{\max}}{m} \mathbf{u}(t), \\ & \dot{m}(t) = -\frac{T_{\max}}{I_{sp} g_0} \|\mathbf{u}(t)\|_2, \\ & \mathbf{x}(t_0) = \mathbf{x}_0, \quad \mathbf{x}(t_f) = \mathbf{x}_f, \\ & m(t_0) = m_0, \\ & \|\mathbf{u}(t)\|_2 \leq 1 \quad \forall t \in [t_0, t_f]. \end{aligned} \quad (7)$$

Reachability is then encoded by the scalar field $m_{\text{MIM}}(\mathbf{x}_0, \mathbf{x}_f, t_f)$: a target state $\mathbf{x}_f$ is reachable for a given spacecraft iff its actual initial mass does not exceed this maximum admissible initial mass.

1. Pontryagin minimum principle conditions

We recast (7) as a minimization problem by considering the Mayer cost Φ

$$J = \Phi = -m(t_0) = -m_0, \quad (8)$$

and apply Pontryagin’s Minimum Principle (PMP) in its standard form for fixed-time, endpoint-constrained problems [16]. We split the state as $\mathbf{x} = [\mathbf{r}^\top, \mathbf{v}^\top]^\top$, with position $\mathbf{r}$ and velocity $\mathbf{v}$, and introduce costates $\boldsymbol{\lambda}_r(t), \boldsymbol{\lambda}_v(t)$, and $\lambda_m(t)$ associated with $\mathbf{r}(t), \mathbf{v}(t)$, and $m(t)$, respectively. The Hamiltonian reads

$$\begin{aligned} H(\mathbf{x}, m, \boldsymbol{\lambda}_r, \boldsymbol{\lambda}_v, \lambda_m, \mathbf{u}) &= \boldsymbol{\lambda}_r^\top \mathbf{v} + \boldsymbol{\lambda}_v^\top (\mathbf{a}_g(\mathbf{r}) + \frac{T_{\max}}{m} \mathbf{u}) \\ &\quad - \lambda_m \frac{T_{\max}}{I_{sp} g_0} \|\mathbf{u}\|_2, \end{aligned} \quad (9)$$

where $\mathbf{a}_g(\mathbf{r})$ denotes the gravitational acceleration (two-body in this work).

The transversality conditions follow from the fact that J depends only on the initial mass, while the initial and final states $\mathbf{x}(t_0)$ and $\mathbf{x}(t_f)$ are fixed. We obtain

$$\lambda_m(t_0) = -\frac{\partial \Phi}{\partial m_0} = 1, \quad (10)$$

while $\lambda_m(t_f) = 0$. All components of λ_r and λ_v at both endpoints are free and are determined by enforcing the two-point boundary conditions.

The remaining costate dynamics are given, as usual, by the negative partial derivatives of H with respect to the state variables. Because $\mathbf{a}_g$ does not depend on m , the costate associated with the mass has the explicit evolution law

$$\dot{\lambda}_m(t) = \lambda_v^\top(t) \frac{T_{\max}}{m^2(t)} \mathbf{u}(t). \quad (11)$$

As we discuss next, the optimal control direction aligns with $-\lambda_v$, so $\lambda_v^\top \mathbf{u} \leq 0$ along optimal trajectories. Since $T_{\max} > 0$ and $m(t) > 0$, this implies that $\dot{\lambda}_m(t) \leq 0$ almost everywhere, i.e. the mass costate is monotonically decreasing and starts from $\lambda_m(t_0) = 1$.

2. Optimal control structure

The Hamiltonian (9) is affine in the control $\mathbf{u}$, subject to the thrust magnitude bound $\|\mathbf{u}(t)\|_2 \leq 1$. Minimization of H yields an optimal thrust direction aligned with the velocity costate,

$$\hat{\mathbf{u}}^*(t) = -\frac{\lambda_v(t)}{\|\lambda_v(t)\|_2}, \quad \|\mathbf{u}^*(t)\|_2 = 1, \quad (12)$$

whenever $\lambda_v(t) \neq \mathbf{0}$; see previous work for detailed derivations of the optimal thrust structure in thrust-bounded low-thrust transfers [11, 17, 18].

Introducing the usual switching function $\sigma(t)$ that multiplies the thrust magnitude, one can show that, for the MIM formulation considered here and for the boundary conditions of interest, $\sigma(t)$ remains strictly negative along optimal trajectories. Consequently, the optimal solution uses maximum available thrust everywhere.

It should be noted that, under the full-thrust structure together with a fixed transfer time, the propellant consumption can be directly determined as $\Delta m = m_0 - m_f = (t_f - t_0)T_{\max}/I_{sp}g_0$. Therefore, an alternative shooting setting can be adopted: since m_f is directly determined by m_0 , the terminal mass is fixed, such that the unknown variable is effectively changed from m_f to λ_{m_f} . This alternative formulation yields the same optimal trajectory, although λ_m may differ by a scaling factor. The equivalence between these two settings arises

from the full-thrust structure, which is commonly derived in time-optimal control problems. A detailed explanation is beyond the scope of the present paper; some relevant discussions can be found in [19], while a more comprehensive analysis will be presented in future work.

3. Forward-backward shooting

In the numerical implementation, the PMP boundary-value problem is solved in a forward-backward fashion using a midpoint-matching shooting formulation [18, 20, 17] and leveraging the smoothing techniques for the control and shooting functions discussed in [21]. The state $\mathbf{x}$ is represented in modified equinoctial elements, and the dynamics and adjoint equations are integrated up to the mismatch point, both in forward time from t_0 and in backward time from t_f . The forward-backward shooting method has proven very effective both for direct methods [22] and for indirect ones [18, 23].

Let $\lambda_{r,0}$, $\lambda_{v,0}$ denote the initial position and velocity costates at t_0 , and $\lambda_{r,f}$, $\lambda_{v,f}$ the corresponding costates at t_f . The decision vector collects the unknown initial mass and costates at both boundaries:

$$\mathbf{X} = [\lambda_{r,0}^\top, \lambda_{v,0}^\top, m_0, \lambda_{r,f}^\top, \lambda_{v,f}^\top, m_f]^\top. \quad (13)$$

For a given $\mathbf{X}$, the augmented state and costates are integrated from t_0 to the midpoint $t_m = \frac{1}{2}(t_0 + t_f)$ and from t_f backward to t_m . Continuity of both state and costate at t_m then defines a system of nonlinear equations

$$\mathbf{F}(\mathbf{X}) = \mathbf{z}^{\text{fwd}}(t_m; \mathbf{X}) - \mathbf{z}^{\text{bck}}(t_m; \mathbf{X}) = \mathbf{0}, \quad (14)$$

where $\mathbf{z} = [\mathbf{x}^\top, m, \lambda_r^\top, \lambda_v^\top, \lambda_m]^\top$ collects state and costates. Solving $\mathbf{F}(\mathbf{X}) = \mathbf{0}$ provides the costate boundary conditions and the associated maximum initial mass m_{MIM} for the prescribed transfer. In the numerical solver, this root-finding problem is tackled using a variant of the Minpack hybrid algorithm [24].

B. Solar Sailing

For the solar-sail case we consider a spacecraft subject to the gravitational attraction of the central body and to a radiation-pressure acceleration generated by an ideal flat sail. The state is again

$$\mathbf{x}(t) = \begin{bmatrix} \mathbf{r}(t) \\ \mathbf{v}(t) \end{bmatrix} \in \mathbb{R}^6, \quad (15)$$

with $\mathbf{r}$ and $\mathbf{v}$ the position and velocity in an inertial frame. The control is the sail normal direction $\mathbf{i}(t) \in \mathbb{R}^3$, constrained by

$$\|\mathbf{i}(t)\|_2 = 1, \quad (16)$$

and the sail strength is parametrised by a scalar coefficient $k_0 > 0$ proportional to the sail lightness number.

Let $r(t) = \|\mathbf{r}(t)\|_2$ and $\mathbf{u}_r(t) = -\mathbf{r}(t)/r(t)$ denote the radial unit vector pointing towards the Sun. The sail acceleration model [25] adopted in this work is

$$\begin{aligned}\dot{\mathbf{r}}(t) &= \mathbf{v}(t), \\ \dot{\mathbf{v}}(t) &= -\mu \frac{\mathbf{r}(t)}{r^3(t)} - k_0 \left(\frac{r_0}{r(t)} \right)^2 \cos^2 \alpha(t) \mathbf{i}(t),\end{aligned}\quad (17)$$

where μ is the standard gravitational parameter, r_0 is a reference distance, and $\alpha(t)$ is the cone angle between $\mathbf{i}(t)$ and $\mathbf{u}_r(t)$, so that $\cos \alpha(t) = \mathbf{i}^T(t) \mathbf{u}_r(t)$. The control appears only through the sail orientation $\mathbf{i}(t)$, while k_0 is treated as a scalar design parameter scaling the solar-radiation-pressure acceleration (and inversely proportional to the sail mass).

We consider a fixed time of flight $[t_0, t_f]$, with the initial state $(\mathbf{r}_0, \mathbf{v}_0)$ prescribed and the final position constrained to a given target $\mathbf{r}_f$, while the terminal velocity is free. To obtain a scalar reachability descriptor analogous to the maximum initial mass in the low-thrust case, we use the sail-strength parameter k_0 itself as the optimisation variable and define a Lagrange cost as:

$$J = \int_{t_0}^{t_f} k_0 dt, \quad (18)$$

so that we minimise k_0 for a fixed transfer time (equivalently, we maximise the sail mass for a given sail area). In the following, we augment the state by treating k_0 as a constant state satisfying $\dot{k}_0 = 0$.

1. Pontryagin minimum principle

We introduce costates $\boldsymbol{\lambda}_r(t)$ and $\boldsymbol{\lambda}_v(t)$ associated with $\mathbf{r}(t)$ and $\mathbf{v}(t)$, respectively, and augment the constant state k_0 , and introduce the associated auxiliary costate $I_{k_0}(t)$ to encode the optimality condition with respect to k_0 . The augmented state–costate is

$$\mathbf{z}(t) = \begin{bmatrix} \mathbf{r}(t) \\ \mathbf{v}(t) \\ k_0 \\ \boldsymbol{\lambda}_r(t) \\ \boldsymbol{\lambda}_v(t) \\ I_{k_0}(t) \end{bmatrix}. \quad (19)$$

The Hamiltonian for the minimisation problem can be expressed as

$$H = \boldsymbol{\lambda}_r^T \mathbf{v} + \boldsymbol{\lambda}_v^T \left(-\mu \frac{\mathbf{r}}{r^3} - k_0 \left(\frac{r_0}{r} \right)^2 \cos^2 \alpha \mathbf{i} \right) + k_0, \quad (20)$$

with $\cos \alpha = \mathbf{i}^T \mathbf{u}_r$. Since $\dot{k}_0 = 0$, the term $I_{k_0} \dot{k}_0$ vanishes identically and does not appear explicitly in the Hamiltonian.

The additional state and costate equations associated with k_0 are $\dot{k}_0 = 0$, and $\dot{I}_{k_0} = -\partial H / \partial k_0$. From (20), we obtain

$$\dot{I}_{k_0} = -1 + \left(\frac{r_0}{r} \right)^2 \cos^2 \alpha \boldsymbol{\lambda}_v^T \mathbf{i}. \quad (21)$$

Since k_0 is unknown and free, the associated transversality conditions are

$$I_{k_0}(t_0) = 0, \quad I_{k_0}(t_f) = 0. \quad (22)$$

These two boundary conditions enforce the first-order optimality of the constant parameter k_0 .

The other costate equations are obtained as usual from the partial derivatives of the Hamiltonian w.r.t the state and are integrated together with the state and $I_{k_0}(t)$.

2. Optimal sail orientation

The control enters the Hamiltonian only through the sail term,

$$H_i = -k_0 \left(\frac{r_0}{r} \right)^2 \cos^2 \alpha \boldsymbol{\lambda}_v^T \mathbf{i}, \quad (23)$$

subject to the constraint $\|\mathbf{i}\|_2 = 1$. Minimising H with respect to $\mathbf{i}$ yields a closed-form expression for the optimal sail normal $\mathbf{i}^*$ as a function of the costates and the radial direction [25]. In particular, the optimal *cone angle* α^* and *clock angle* define a unit vector $\mathbf{i}^*$ in the plane spanned by the radial unit vector $\mathbf{u}_r$ and the transverse component of $\boldsymbol{\lambda}_v$. The resulting law coincides with the classical optimal orientation for ideal solar sails, whereby the sail normal lies in the plane generated by $\mathbf{u}_r$ and $\boldsymbol{\lambda}_v$ and balances the radial and transverse contributions to minimise the Hamiltonian [26].

3. Single shooting

As in the electric low-thrust case, the necessary conditions for optimality are enforced via a shooting formulation. In this case, however, the control structure is smooth, and the associated costate dynamics exhibit reduced sensitivity to the initial costate values. As a result, we do not find forward–backward schemes strictly necessary to regularize the boundary-value problem, for the time horizons and transfer types that we investigate. In our numerical experiments, the single-shooting formulation converged reliably, which motivated its adoption in the solar-sail case.

The decision vector collects the initial costates and the parameter k_0 ,

$$\mathbf{X} = [\boldsymbol{\lambda}_r^T(t_0), \boldsymbol{\lambda}_v^T(t_0), k_0]^T, \quad (24)$$

and the augmented PMP dynamics is integrated over the fixed time interval $[t_0, t_f]$. The residuals enforce termi-

nal position and velocity matching and the parameter-stationarity condition,

$$\begin{aligned} \mathbf{r}(t_f; \mathbf{X}) - \mathbf{r}_f &= \mathbf{0}, \\ \mathbf{v}(t_f; \mathbf{X}) - \mathbf{v}_f &= \mathbf{0}, \\ I_{k_0}(t_f; \mathbf{X}) &= 0, \end{aligned} \quad (25)$$

resulting in a system of seven nonlinear equations in seven unknowns. As before, this root-finding problem is solved via the Minpack hybrid algorithm, yielding the optimal parameter k_0 and the associated sail trajectory and costate history. The corresponding maximum mass (or minimum sail area) for the chosen dynamical model can then be obtained by inverting the relationship between k_0 and the characteristic sail acceleration.

IV. MACHINE LEARNING SURROGATES

We now construct data-driven surrogate models for the maximum-initial-mass formulation introduced above. The overarching goal is to replace repeated indirect optimal-control solves by fast evaluations, both for the two-body low-thrust reachability problem and for the solar-sail transfers, in the spirit of previous work on neural approximators for asteroid-belt transfers [7, 8, 9, 15]. A key contribution of this paper is to show that residual multilayer perceptrons, equipped with smooth activations, provide particularly effective surrogates for these dual reachability indicators, combining high accuracy with relatively modest parameter counts.

A. Two-body low-thrust reachability surrogate

In the two-body low-thrust setting we do not regress a continuous performance index, but directly learn a binary reachability indicator.

A transfer is labelled reachable if the maximum initial mass exceeds a prescribed spacecraft mass m_{sc} (here $m_{sc} = 1500$ kg), and unreachable otherwise. Formally, the classification label is

$$y = \mathbb{I}\left(\frac{m_{MIM}}{m_{sc}} \geq 1\right) \in \{0, 1\}, \quad (26)$$

so that $y = 1$ denotes a reachable target, whereas $y = 0$ corresponds to an infeasible transfer.

For each transfer, we solve a forward-backward shooting indirect formulations, described in Sec. III.A, and a set of normalised features, and the true maximum initial masses are then stored.

The input feature vector $\phi \in \mathbb{R}^{11}$ concatenates

- the normalised time of flight and specific impulse, denoted `tof_n` and `isp_n`;
- the three components of the Lambert initial $\Delta \mathbf{v}$ at the departure asteroid, expressed in the local RTN

basis and normalised (`dv_lambert_start_r_n`, `t_n`, `n_n`);

- the three components of the Lambert arrival $\Delta \mathbf{v}$ at the target asteroid in the same RTN frame (`dv_lambert_target_r_n`, `t_n`, `n_n`);
- the normalised eccentricity of the departing orbit (`ecc_n`);
- the cosine and sine of the true anomaly at departure (`cos_f`, `sin_f`).

All features are standardised based on the dataset statistics and subsequently clipped to a finite range, which improves conditioning and limits the influence of outliers. As shown in previous works, rotational invariant and Lambert solutions inputs are very useful to improve the accuracy of the surrogate [9]. On the other hand, the feature space is also maintained as compact as possible, as this has been also shown to have benefits [15].

We assemble the dataset into disjoint training, validation, and test sets with proportions 80%, 10%, and 10%, respectively. The class label distribution is not imbalanced, as roughly half of the random asteroid-pair transfers are reachable in the considered time of flights.

B. Solar-sail surrogate

In the solar-sail case (Sec. III.B), the scalar quantity of interest is the minimum sail-strength parameter k_0^* that renders a transfer feasible within a fixed time of flight. For each sampled transfer we solve the PMP boundary-value problem with parameter stationarity, obtaining an optimal value k_0^* and the associated state-costate history. The value of k_0^* directly encodes reachability in terms of sail mass or sail area: smaller values correspond to larger admissible spacecraft masses for a given sail area [25].

We construct a feature vector ϕ encoding the initial and final heliocentric states, the time of flight, and coarse information about the heliocentric distance range covered by the transfer. The surrogate model is then trained to approximate k_0^* itself.

We consider four classes of fully connected neural networks mapping ϕ to a logit target:

Residual MLP (ResNet). The main architecture used in this work is a residual multilayer perceptron, whose residual blocks are shown in Fig. 1.

An input projection maps ϕ to a hidden representation of width d_h , followed by a stack of residual blocks of the form

$$\mathbf{h}_{\ell+1} = W_{2,\ell} \sigma(W_{1,\ell} \mathbf{h}_\ell + \mathbf{b}_{1,\ell}) + \mathbf{b}_{2,\ell} + \mathbf{h}_\ell, \quad (27)$$




Fig. 1: Residual block

where σ is a nonlinearity (we use **SELU**, **SiLU**, **Softplus**, or **GELU**) and ℓ indexes the residual blocks. A final linear layer maps the last hidden state to the scalar output.

Plain MLP. As a strong non-residual baseline we employ a conventional MLP with L hidden layers of width d_h , using the **SELU** and **ReLU** activations in all hidden layers and Kaiming-normal initialisation.

SIREN network. Finally, we include a sinusoidal representation network (SIREN) [27], consisting of a first SineLayer that maps ϕ to a hidden width d_h , followed by several hidden SineLayers with the same width and a final linear output layer. Each SineLayer applies a linear transformation followed by a scaled sinusoidal activation ($\sin(\omega_0 \cdot)$), with (ω_0) controlling the effective frequency and the weight initialisation range. We use the initialisation rules proposed in [27] and treat ω_0 as a hyperparameter.

C. Training objective and optimisation

For the two-body reachability surrogate we adopt a binary-classification objective. Let $f_\theta(\phi) \in \mathbb{R}$ denote the logit output of a network with parameters θ . The loss function is the binary cross-entropy with logits,

$$\mathcal{L}(\theta) = \frac{1}{N} \sum_{i=1}^N \text{BCEWithLogits}(f_\theta(\phi_i), y_i), \quad (28)$$

with a small amount of label smoothing (here 5%) applied to both classes to mitigate overconfidence and reduce the impact of mislabeled samples. During validation and testing, we threshold the sigmoid of the logit at 0.5 to obtain hard predictions and report accuracy, precision, recall, and F_1 -score, together with the confusion-matrix entries.

In all experiments we use the AdamW optimiser with weight decay, an initial learning rate of order 10^{-3} , and a cosine-annealing schedule down to a small floor value. Mini-batch training is employed, with batch sizes in the few-thousand range. The best checkpoint is selected based on the lowest validation loss, and all reported test metrics correspond to that checkpoint.

For the solar-sail surrogate, the same optimization setup is used as the low-thrust electric propulsion scenario. In this case the network output approximates the

scalar field $k_0^*(\phi)$, which plays the same role for solar sails as m_{MIM} does for electric low-thrust transfers.

D. Why residual networks outperform SIREN and plain MLPs

Across both the two-body reachability and the solar-sail regression experiments, residual MLPs consistently achieve lower validation loss, higher classification F_1 -scores, and tighter error distributions than SIREN and plain MLP architectures of comparable width and depth.

This empirical observation is consistent with several known advantages of residual architectures.

First, residual networks are well-suited to such functions because each block learns a perturbation of the identity, so the layerwise Jacobian takes the form

$$J_{\ell+1} = \frac{\partial \mathbf{h}_{\ell+1}}{\partial \mathbf{h}_\ell} \approx I + J_{F_\ell}, \quad (29)$$

with F_ℓ the residual mapping. The resulting product of Jacobians across layers tends to remain better conditioned, keeping singular values closer to one and alleviating the vanishing- and exploding-gradient problems that affect deep plain networks [28]. This leads to more stable optimisation and smoother decision boundaries, which also translate into improved generalisation in regions with sparse training data.

Second, the residual parameterisation naturally biases the network towards learning corrections around a simple baseline mapping. A residual MLP can be interpreted as the integration of small vector-field updates, which yields stable forward dynamics and well-conditioned backward sensitivity.

Third, while SIREN networks are highly expressive and excel at representing high-frequency structure, their sinusoidal activations tend to produce highly oscillatory Jacobians, which can sometimes lead to unstable training results and/or poorer generalization capabilities.

Plain MLPs instead would either require substantially deeper architectures, which exacerbates optimisation difficulties, or significantly wider layers to reach comparable accuracy, increasing computational cost.

Finally, from a practical standpoint, the residual architectures show markedly more robust behaviour under changes in hyperparameters and random seeds. A larger fraction of training runs converge to high-performing solutions without manual tuning, and the spread in validation metrics across independent runs is reduced.

V. EXPERIMENTS

A. Two-body electric low-thrust reachability

These results were generated leveraging a two-body problem dynamics where the central body is the Sun, with a

family	activation	N	accuracy	precision	recall	F1	loss
mlp	softplus	64	0.8493	0.8397	0.8537	0.8466	0.3461
residual	gelu	64	0.9773	0.9750	0.9785	0.9768	0.0808
residual	relu	64	0.9730	0.9703	0.9744	0.9724	0.0914
residual	selu	64	0.9730	0.9703	0.9744	0.9724	0.0914
residual	softplus	64	0.9777	0.9749	0.9795	0.9772	0.0815
residual	softplus	128	0.9791	0.9769	0.9802	0.9786	0.0774
residual	softplus	256	0.9791	0.9769	0.9802	0.9786	0.0774
residual	softplus	512	0.9785	0.9763	0.9796	0.9780	0.0789
siren	gelu	64	0.9290	0.9248	0.9300	0.9274	0.1881
siren	silu	64	0.9301	0.9254	0.9316	0.9285	0.1856
siren	softplus	64	0.9290	0.9248	0.9300	0.9274	0.1881

Tab. 1: Validation metrics for different architectures and activations for the electric low-thrust propulsion case.

maximum thrust of 0.1 N, and specific impulse $I_{sp} \in [1000, 5000]$ s. In terms of starting and target orbits, these were generated within the following Keplerian elements: semi-major axis $a \in [0.5, 3.5]$ AU, eccentricity $e \in [0, 1)$, inclination $i \in [0, 2\pi]$, true anomaly $\nu \in [0, 2\pi]$, right ascension of the ascending node $\Omega \in [0, 2\pi]$, and argument of perigee $\omega \in [0, 2\pi]$. 0, 1, 2, and 3 revolutions transfers are generated, with equal proportions in the training, validation, and test sets. A transfer is deemed feasible if the maximum initial mass is above the spacecraft mass (here assumed to be 1500 kg). Once a solution is generated, the same homotopy strategy as detailed in [15] is used to continue the family of solutions. Tab. 1 reports validation metrics for different architectures and activation functions on the electric low-thrust reachability classification task. Among models with the smallest hidden width ($N = 64$), the residual network with **Softplus** activation achieves the best validation performance, with accuracy 0.9777, precision 0.9749, recall 0.9795, F_1 -score 0.9772, and cross-entropy loss 0.0815. This configuration is highlighted in bold in the table and already provides a very competitive trade-off between accuracy and number of learnable parameters.

Increasing the width of the residual **Softplus** network further improves the metrics, but with diminishing returns. The best overall model in terms of validation loss is the residual **Softplus** network with $N = 512$ neurons per hidden layer, which attains accuracy 0.9785, precision 0.9763, recall 0.9796, F_1 -score 0.9780, and loss 0.0789. While this represents a slight improvement over the $N = 64$ counterpart, the gains are not dramatic when compared to the substantial increase in parameter count and computational cost, suggesting that relatively compact residual networks already capture most of the structure of the dual reachability indicator.

On the held-out test set (comprising 10% of the overall data), the selected residual **Softplus** model yields the




Fig. 2: Comparison of ground truth and surrogate reachability on the test set. The surrogate captures the main feasible/infeasible structure and sharp decision boundary, with most errors and low-confidence predictions occurring near the threshold, particularly close to the mass limit $m_{sc} = 1500$ kg.

following confusion matrix:

$$CM_{\text{test}} = \begin{bmatrix} 468418 & 11138 \\ 9143 & 445759 \end{bmatrix}, \quad (30)$$

where the rows correspond to the true labels (unreachable, reachable) and the columns to the predicted labels. This corresponds to an overall test accuracy of approximately 0.978, precision of about 0.976, recall of about 0.980, and F_1 -score of about 0.978, in line with the validation metrics. Most errors occur near the reachability boundary: the number of false positives (predicted reachable but actually infeasible) and false negatives (predicted infeasible but actually reachable) is small compared to the number of correctly classified transfers, and their distribution reflects the intrinsic ambiguity of borderline cases.

To illustrate this behaviour qualitatively, Fig. 2 shows a 2×2 panel restricted to a subset of the test set. The top-left panel displays the ground-truth reachability labels in the plane of departure true anomaly and initial orbit eccentricity, while the top-right panel shows the corresponding surrogate predictions for the same samples. The surrogate reproduces the main structure of the reachable and unreachable regions, including sharp transitions across the boundary, with visually small regions of disagreement.

The bottom-left panel depicts the signed distance of the logits to the decision threshold, i.e. the signed margin with respect to the 0.5 probability level. The distri-

bution of these margins is clearly bimodal: confidently classified reachable and unreachable transfers populate two well-separated lobes, while misclassified cases concentrate near the decision boundary where the margin is small in absolute value. This is consistent with the idea that most classification errors occur for transfers that are intrinsically close to the feasibility threshold.

Finally, the bottom-right panel shows the distribution of misclassified test samples in terms of the true maximum initial mass. As expected, these points cluster around the threshold $m_{sc} = 1500$ kg: transfers that are only marginally feasible or infeasible are the ones that are most difficult to classify, both for the surrogate. This behaviour is desirable from a mission-design standpoint, as it indicates that the surrogate rarely misclassifies clearly reachable or clearly unreachable targets, and concentrates its uncertainty in a narrow band around the operational mass limit.

B. Solar-sail surrogate performance

In the solar-sail case, the GTOC13 problem¹ is used as the simulation environment (“gym”), with Altaira as the central body. The orbital parameters follow the same bounds as in the original formulation, except for the semi-major axis, which is extended to the range 0.01–8 AU. The sail area and spacecraft mass used to compute the threshold sail strength, and thereby determine target reachability, are 15,000 m² and 500 kg, respectively. For these experiments, we reuse the same residual **Softplus** architecture identified as best in the electric low-thrust study. On the validation set, the best configuration attains loss 0.0364, precision 0.9922, accuracy 0.9945, and recall 0.9927. From these values we compute an F_1 -score of 0.9924, indicating a very high-quality surrogate that almost perfectly separates feasible from infeasible transfers when thresholded at the appropriate k_0 level.

Taken together with the electric low-thrust results, this confirms that residual networks with smooth activations provide a robust architecture choice across both reachability problems studied here. Once trained, the solar-sail surrogate can be queried at negligible cost to evaluate approximate reachability over large sets of target states, effectively turning the MIM-based dual formulation into a practical, data-driven reachability oracle for preliminary sail-mission design.

VI. CONCLUSIONS

This work introduced a dual formulation of low-thrust reachability based on the notion of maximum initial mass (MIM) for electric propulsion and an analogous minimum sail-strength parameter for solar-sail transfers. By

reformulating the classical reachability question into the evaluation of the maximum initial mass scalar field, the proposed approach replaces the construction of a high-dimensional reachable set with the computation of a well-defined, continuous, and smoother performance indicator. This reformulation retains the full feasibility information of the original problem while significantly improving its numerical structure, particularly near the boundary of the reachable set where the corresponding dual minimum time problem tends to exhibit less stability and convergence issues.

We derived indirect optimal-control formulations for both electric low-thrust and solar-sail dynamics using Pontryagin’s Minimum Principles. In the electric case, the resulting boundary-value problem is formulated via a forward-backward shooting scheme. For solar sailing, we wrote the PMP to include parameter stationarity conditions, enabling the consistent treatment of the sail-strength parameter within the same framework. These developments provide a unified viewpoint in which reachability is encoded through scalar optimality conditions.

Building on these dual formulations, we investigated data-driven surrogate models that approximate the resulting scalar reachability indicators. Across both problem classes, residual multilayer perceptrons with smooth activations, particularly Softplus-based architectures, consistently outperformed plain MLPs and SIREN networks in terms of accuracy and training stability. The learned surrogates accurately reproduce the structure of the reachable set, making them well-suited as practical reachability oracles.

Overall, the combination of indirect MIM formulations and neural surrogates provides an efficient and scalable framework for preliminary mission design under low-thrust dynamics. The resulting approach enables rapid feasibility assessment without repeated optimal-control solves, substantially reducing computational cost while preserving high accuracy.

REFERENCES

- [1] Holzinger, M. J. and Scheeres, D. J. “Reachability results for nonlinear systems with ellipsoidal initial sets.” *IEEE transactions on aerospace and electronic systems*, Vol. 48, No. 2, pp. 1583–1600, 2012.
- [2] Pellegrino, M. and Scheeres, D. J. “Reachability of a passive solar sail in earth orbit.” *Journal of Guidance, Control, and Dynamics*, Vol. 44, No. 2, pp. 360–369, 2021.
- [3] Lee, S. and Hwang, I. “Reachable set computation for spacecraft relative motion with energy-limited

¹<https://gtoc.jpl.net/gtoc13/the-problem/>

- low-thrust.” *Aerospace Science and Technology*, Vol. 77, pp. 180–188, 2018.
- [4] Chen, Q., Qiao, D., and Wen, C. “Minimum-fuel low-thrust trajectory optimization via reachability analysis and convex programming.” *Journal of Guidance, Control, and Dynamics*, Vol. 44, No. 5, pp. 1036–1043, 2021.
 - [5] Bowerfind, S. and Taheri, E. “Rapid approximation of low-thrust spacecraft reachable sets within complex two-body and cislunar dynamics.” *Aerospace*, Vol. 11, No. 5, p. 380, 2024.
 - [6] Wang, Z. and Jiang, F. “Analytical optimal solution for the reachable domain of low-thrust spacecraft.” *Journal of Spacecraft and Rockets*, Vol. 61, No. 1, pp. 274–284, 2024.
 - [7] Hennes, D., Izzo, D., and Landau, D. “Fast approximators for optimal low-thrust hops between main belt asteroids.” “2016 IEEE Symposium Series on Computational Intelligence (SSCI),” pp. 1–7. IEEE, 2016.
 - [8] Mereta, A., Izzo, D., and Wittig, A. “Machine learning of optimal low-thrust transfers between near-earth objects.” “International Conference on Hybrid Artificial Intelligence Systems,” pp. 543–553. Springer, 2017.
 - [9] Acciarini, G., Beauregard, L., and Izzo, D. “Computing Low-Thrust Transfers in the Asteroid Belt, a Comparison between Astrodynamical Manipulations and a Machine Learning Approach.” “29th International Symposium on Space Flight Dynamics (ISSFD 2024),” Darmstadt, Germany, Apr. 2024. doi:10.48550/arXiv.2405.18918.
 - [10] Izzo, D., Märten, M., Beauregard, L., Bannach, M., Acciarini, G., Blazquez, E., Hadjiivanov, A., Grover, J., Heißel, G., Shimane, Y., et al. “Asteroid mining: ACT&Friends’ results for the GTOC12 problem.” *Astrodynamics*, Vol. 9, No. 1, pp. 19–40, 2025.
 - [11] Kirk, D. E. *Optimal control theory: an introduction*. Courier Corporation, 2004.
 - [12] Bansal, S., Chen, M., Herbert, S., and Tomlin, C. J. “Hamilton-jacobi reachability: A brief overview and recent advances.” “2017 IEEE 56th annual conference on decision and control (CDC),” pp. 2242–2253. IEEE, 2017.
 - [13] Fisac, J. F., Chen, M., Tomlin, C. J., and Sastry, S. S. “Reach-avoid problems with time-varying dynamics, targets and constraints.” “Proceedings of the 18th international conference on hybrid systems: computation and control,” pp. 11–20. 2015.
 - [14] Taheri, E. and Junkins, J. L. “How many impulses redux.” *The Journal of the Astronautical Sciences*, Vol. 67, No. 2, pp. 257–334, 2020.
 - [15] Zhang, Z., Acciarini, G., Izzo, D., Baoyin, H., and Topputo, F. “Pretrained Approximators for Low-Thrust Trajectory Cost and Reachability.” *Journal of Guidance, Control, and Dynamics*, 2026. To appear.
 - [16] Pontryagin, L. S., Boltyanskii, V. G., Gamkrelidze, R. V., and Mishchenko, E. F. *The Mathematical Theory of Optimal Processes*. Interscience Publishers / John Wiley & Sons, New York; London, 1962.
 - [17] Pérez-Palau, D. and Epenoy, R. “Fuel optimization for low-thrust Earth–Moon transfer via indirect optimal control.” *Celestial Mechanics and Dynamical Astronomy*, Vol. 130, No. 2, p. 21, 2018.
 - [18] Sidhoum, Y. and Oguri, K. “Indirect forward-backward shooting for low-thrust trajectory optimization in complex dynamics.” *Journal of Guidance, Control, and Dynamics*, Vol. 47, No. 10, pp. 2164–2172, 2024.
 - [19] Guo, X., Wu, D., and Jiang, F. “Minimum-Time Rendezvous via Simplified Initial Costate Normalization and Auxiliary Orbital Transfer.” *Journal of Guidance, Control, and Dynamics*, Vol. 46, No. 8, pp. 1627–1636, 2023. doi:10.2514/1.G007268.
 - [20] Jiang, F., Baoyin, H., and Li, J. “Practical techniques for low-thrust trajectory optimization with homotopic approach.” *Journal of guidance, control, and dynamics*, Vol. 35, No. 1, pp. 245–258, 2012.
 - [21] Bertrand, R. and Epenoy, R. “New smoothing techniques for solving bang–bang optimal control problems—numerical results and statistical interpretation.” *Optimal Control Applications and Methods*, Vol. 23, No. 4, pp. 171–197, 2002.
 - [22] Izzo, D., Holt, H., Acciarini, G., Beauregard, L., and Shimane, Y. “A practical guide to implementing zero-order-hold interplanetary trajectory legs.” *30th International Symposium on Space Flight Dynamics*, 2026.

- [23] Zhang, Z., Yang, J., Izzo, D., and Topputo, F. “Forward–Backward Indirect Shooting for Low-Thrust Trajectories with Interior Events.” 30th International Symposium on Space Flight Dynamics, 2026.
- [24] Powell, M. J. “A hybrid method for nonlinear equations.” Numerical methods for nonlinear algebraic equations, pp. 87–161, 1970.
- [25] McInnes, C. R. Solar sailing: technology, dynamics and mission applications. Springer Science & Business Media, 2004.
- [26] McInnes, C. R. “Solar sail mission applications for non-Keplerian orbits.” Acta Astronautica, Vol. 45, No. 4-9, pp. 567–575, 1999.
- [27] Sitzmann, V., Martel, J., Bergman, A., Lindell, D., and Wetzstein, G. “Implicit neural representations with periodic activation functions.” Advances in neural information processing systems, Vol. 33, pp. 7462–7473, 2020.
- [28] He, K., Zhang, X., Ren, S., and Sun, J. “Deep residual learning for image recognition.” “Proceedings of the IEEE conference on computer vision and pattern recognition,” pp. 770–778. 2016.