

Tightly-Coupled Estimation and Guidance for Robust Low-Thrust Uncooperative Rendezvous via Adaptive Homotopy

Batu Candan

`dukynuke@iastate.edu`

Iowa State University

Simone Servadio

Iowa State University

Research Article

Keywords: Spacecraft rendezvous, indirect optimal control, homotopy continuation, adaptive Kalman filtering, low-thrust guidance, active debris removal

Posted Date: September 24th, 2026

DOI: <https://doi.org/10.21203/rs.3.rs-10977398/v1>

License: © ⓘ This work is licensed under a Creative Commons Attribution 4.0 International License.

[Read Full License](#)

Additional Declarations: No competing interests reported.

Tightly-Coupled Estimation and Guidance for Robust Low-Thrust Uncooperative Rendezvous via Adaptive Homotopy

Batu Candan^{1*} and Simone Servadio¹

^{1*}Department of Aerospace Engineering, Iowa State University, Ames, 50011, IA, USA.

*Corresponding author(s). E-mail(s): dukynuke@iastate.edu;
Contributing authors: servadio@iastate.edu;

Abstract

Minimum-fuel low-thrust rendezvous guidance yields bang-bang control structures that are highly sensitive to estimation errors, non-Gaussian sensor anomalies, and solver regularization, making purely aggressive closed-loop execution brittle for uncooperative proximity operations. This paper proposes a tightly-coupled estimation and guidance architecture in which navigation confidence directly modulates the homotopy parameter of a receding-horizon indirect optimal control solver. Relative motion is modeled in the Clohessy–Wiltshire frame, and the translational state is estimated with a linear Kalman filter augmented by Multiple Tuning Factors (MTF) covariance inflation mechanism that suppresses suspicious innovation directions. A composite innovation score, constructed from the normalized innovation and the MTF activity, is then mapped online to the homotopy continuation parameter, allowing the controller to relax toward a smoother and more conservative regime when confidence in the estimates degrades and to recover a more fuel-efficient regime as sensing improves. Numerical results under severe measurement degradation show that fixed bang-bang guidance remains brittle, with other methods exhibiting large terminal miss distances. In contrast, the proposed MTF-adaptive homotopy controller reduces the terminal miss, bringing the final error down to the sub-meter level, while requiring only a moderate increase in control effort relative to the open-loop fuel-optimal benchmark. A comparison study further indicates that neither mechanism is sufficient on its own: the adaptive schedule cannot meet the terminal tolerance when driven by a filter that has not rejected the corrupted measurements, and covariance inflation alone does not help a guidance problem that has

become numerically intractable. The benefit therefore comes from their interaction. The receding-horizon implementation also exhibited consistently fast and reliable solution times, supporting the practical online viability of the proposed method for the tested scenario.

Keywords: Spacecraft rendezvous, indirect optimal control, homotopy continuation, adaptive Kalman filtering, low-thrust guidance, active debris removal

1 Introduction

Autonomous rendezvous with uncooperative objects is a key enabling capability for active debris removal, on-orbit servicing, inspection, and future in-space assembly missions. In these settings, the target may be poorly characterized, tumbling, and observed through degraded sensing conditions, so estimation, planning, and control must operate robustly and with limited opportunity for human intervention [1–3]. At the same time, low-thrust propulsion is especially attractive because of its high specific impulse, fine control authority, and favorable mission-level fuel efficiency, which has motivated a broad literature on low-thrust trajectory design and spacecraft guidance [4–7]. The main difficulty is that minimum-fuel low-thrust guidance is numerically demanding. Direct transcription, convexification, and model-predictive approaches have become particularly attractive because they can accommodate path and control constraints and can be adapted to receding-horizon operation [8–16]. In parallel, learning-based surrogates and controllers are increasingly being explored to reduce onboard computational cost or to provide fast approximations of optimal guidance, including supervised surrogates, deep warm starts, reinforcement learning, and certifiably stable neural guidance laws [17–23]. Taken together, recent surveys on low-thrust optimization, convex guidance, and learning for optimal control show how rapidly this design space is expanding [4–6, 23, 24].

Indirect methods remain compelling because, once properly initialized, they can deliver highly accurate fuel-optimal solutions with a smaller decision space by exploiting Pontryagin’s Minimum Principle. Their weakness, however, is equally well known as the resulting Hamiltonian boundary-value problem has a narrow convergence basin, and the bang-bang or bang-off-bang structure of minimum-fuel low-thrust control substantially increases sensitivity to the initial co-states and to the numerical regularization used during solution [25–31]. For this reason, a large body of work has focused on smoothing, continuation, and regularization strategies that embed the original fuel-optimal problem inside a family of easier auxiliary problems [25–27, 29, 30, 32]. Related homotopic ideas have also appeared in convex and mixed-logic rendezvous optimization, where continuation improves convergence and computational reliability [13, 15]. Broader reviews of low-thrust optimization likewise emphasize the central role of homotopy and smoothing in making discontinuous minimum-fuel solutions computationally tractable [4, 31]. Despite this progress, the homotopy or regularization parameter is still usually treated as an optimizer-side continuation variable rather than as a closed-loop robustness actuator driven by onboard navigation. At the same

time, integrated rendezvous pipelines have emphasized the importance of estimation, planning, and control under uncertainty [1, 33], but the numerical aggressiveness of the low-thrust guidance law is still generally decoupled from the estimator’s instantaneous confidence. This separation is especially limiting for uncooperative rendezvous, where non-Gaussian anomalies such as optical (shadowing, glare) effects, imperfect target knowledge, or unmodeled disturbances can abruptly degrade the state estimate while the guidance law continues to pursue an aggressive fuel-optimal solution.

This paper addresses that gap by proposing a tightly coupled estimation and guidance architecture for robust low-thrust rendezvous. Relative motion is modeled in the Clohessy–Wiltshire (CW) frame, and the guidance law is generated by a receding-horizon indirect solver for the low-thrust optimal control problem. To improve robustness on the estimation side, a Multiple Tuning Factors (MTF) Kalman filtering strategy is used to adapt the covariance in the presence of non-Gaussian anomalies and measurement faults [34]. The key idea is to reinterpret the normalized innovation not only as an estimation-consistency measure, but also as a control-relevant confidence signal. Specifically, the innovation is mapped online to the homotopy continuation parameter of the indirect solver. When the innovation remains small, the homotopy parameter is reduced, and the controller recovers a fuel-optimal bang-bang-like solution. When the innovation increases, the homotopy parameter is increased automatically, relaxing the problem toward a smoother and more robust energy-optimal regime. In this way, the onboard guidance law responds not only to the estimated state, but also to the reliability of that estimate.

Therefore, the main contributions of this work are threefold. First, we formulate a tightly coupled receding-horizon architecture in which estimator confidence directly modulates the continuation process of an indirect low-thrust solver. Second, we adapt an MTF-based covariance tuning concept to translational relative motion guidance for uncooperative rendezvous. Third, we show how innovation-adaptive homotopy enables a continuous and physically interpretable trade between fuel optimality and robustness during low-thrust close-proximity operations. In short, this work demonstrates that effective guidance requires a continuous interaction between estimation and control, where the desired trajectory is not determined solely by the control objective, but is adapted according to how informative the available measurements are and how confidently the current state can be estimated. In this way, the spacecraft can select a guidance path that simultaneously accounts for mission performance and the quality of the information available to safely execute it.

A preliminary version of this study was presented at the AAS/AIAA Astrodynamics Specialist Conference [35]. The remainder of the paper is organized as follows. Section 2 presents the relative-motion model, the minimum-fuel low-thrust optimal control problem, and the homotopy-regularized indirect guidance formulation. Section 3 introduces the Kalman filtering framework, the measurement-side MTF adaptation, and the innovation-adaptive homotopy scheduling law. Section 4 presents the numerical results, including the baseline closed-loop comparison and the adaptive-only comparison study. Finally, Section 5 concludes the paper and outlines directions for future work.

2 Dynamical Model and Optimal Control

2.1 Relative Motion Dynamics

The relative motion of the chaser spacecraft with respect to the target is modeled in the CW frame. Let the state vector be

$$\mathbf{x} = [x \ y \ z \ \dot{x} \ \dot{y} \ \dot{z}]^T \quad (1)$$

and let the control acceleration be

$$\mathbf{u} = [u_x \ u_y \ u_z]^T. \quad (2)$$

Assuming the target follows a circular reference orbit with mean motion n , the linearized relative dynamics are

$$\dot{\mathbf{x}} = \mathbf{A}\mathbf{x} + \mathbf{B}\mathbf{u} \quad (3)$$

with

$$\mathbf{A} = \begin{bmatrix} 0 & 0 & 0 & 1 & 0 & 0 \\ 0 & 0 & 0 & 0 & 1 & 0 \\ 0 & 0 & 0 & 0 & 0 & 1 \\ 3n^2 & 0 & 0 & 0 & 2n & 0 \\ 0 & 0 & 0 & -2n & 0 & 0 \\ 0 & 0 & -n^2 & 0 & 0 & 0 \end{bmatrix}, \quad \mathbf{B} = \begin{bmatrix} 0 & 0 & 0 \\ 0 & 0 & 0 \\ 0 & 0 & 0 \\ 1 & 0 & 0 \\ 0 & 1 & 0 \\ 0 & 0 & 1 \end{bmatrix}. \quad (4)$$

The CW model is adopted here as a first-order translational model that is sufficiently simple for repeated onboard re-optimization while still capturing the principal relative dynamics relevant to close-proximity rendezvous, as it is assumed that the chaser spacecraft initiates the docking maneuver after completing the approaching stage of the mission, thus initiating the rendezvous below the hundreds-of-meters initial condition.

2.2 Minimum-Fuel Problem and Homotopy Regularization

The guidance objective is to steer the relative state from the current estimated state to a desired terminal condition while minimizing control effort. The nominal minimum-fuel optimal control problem is written as

$$J = \int_0^{t_f} \|\mathbf{u}(t)\| dt \quad (5)$$

subject to Eq. (3) and the control bound

$$\|\mathbf{u}(t)\| \leq u_{\max}. \quad (6)$$

Applying Pontryagin's Minimum Principle yields the Hamiltonian

$$\mathcal{H} = \|\mathbf{u}\| + \boldsymbol{\lambda}^T(\mathbf{A}\mathbf{x} + \mathbf{B}\mathbf{u}), \quad (7)$$

where $\boldsymbol{\lambda} = [\boldsymbol{\lambda}_r^T \ \boldsymbol{\lambda}_v^T]^T$ is the costate vector. The optimal thrust direction is

$$\boldsymbol{\alpha} = -\frac{\boldsymbol{\lambda}_v}{\|\boldsymbol{\lambda}_v\|}, \quad (8)$$

and the switching function is

$$S(t) = 1 - \|\boldsymbol{\lambda}_v(t)\|. \quad (9)$$

In the fuel-optimal limit, the control is bang-bang. Because this discontinuity makes the indirect boundary-value problem difficult to solve robustly, a scalar homotopy parameter $\epsilon \geq 0$ is introduced to regularize the control magnitude [36]. The smoothed control law is written as

$$\mathbf{u} = \begin{cases} 0, & S > \epsilon \\ \frac{\epsilon - S}{2\epsilon} u_{\max} \boldsymbol{\alpha}, & -\epsilon \leq S \leq \epsilon \\ u_{\max} \boldsymbol{\alpha}, & S < -\epsilon. \end{cases} \quad (10)$$

As $\epsilon \rightarrow 0$, Eq. (10) recovers the bang-bang fuel-optimal structure. For larger ϵ , the control becomes smoother and numerically more robust. In the proposed architecture, this regularization parameter is not treated purely as an offline continuation variable; instead, it is adjusted online according to estimator confidence.

This regularization has a variational origin, which is worth making explicit because it determines both why larger ϵ is easier to solve and why the classical way of exploiting it is unsuitable onboard. Equation (10) is not an ad hoc smoothing of the bang-bang law but the exact minimizer of a one-parameter family of optimal control problems. Writing the normalized throttle as $\delta = \|\mathbf{u}\|/u_{\max} \in [0, 1]$, the family is defined by

$$J_\epsilon = u_{\max} \int_0^{t_f} \left[\delta - \epsilon \delta (1 - \delta) \right] dt, \quad \epsilon \in [0, 1], \quad (11)$$

subject to the same dynamics and control bound. At $\epsilon = 0$ the integrand reduces to δ and Eq. (11) recovers the minimum-fuel index of Eq. (5); at $\epsilon = 1$ it reduces to δ^2 , giving the energy-optimal problem, whose solution is smooth and whose boundary-value problem converges readily from a coarse costate guess.

Applying Pontryagin's Minimum Principle to Eq. (11) and minimizing first over the thrust direction returns Eq. (8), after which the control-dependent part of the Hamiltonian is quadratic in the throttle,

$$\mathcal{H}_\epsilon = u_{\max} \epsilon \delta^2 + u_{\max} (S - \epsilon) \delta + \boldsymbol{\lambda}^T A \mathbf{x}, \quad (12)$$

with S given by Eq. (9). For $\epsilon > 0$ the coefficient of δ^2 is strictly positive, so $\mathcal{H}_\epsilon$ is strictly convex in δ and the stationary point is unique,

$$\delta^\star = \frac{\epsilon - S}{2\epsilon}, \quad (13)$$

which, saturated to $[0, 1]$, places the corners at $S = \epsilon$ and $S = -\epsilon$ and so reproduces Eq. (10) exactly. Every intermediate ϵ therefore returns a unique optimal solution of a defined blend of fuel and energy optimality, and the convexity supplied by the $\epsilon\delta^2$ term is what makes larger ϵ numerically benign. As $\epsilon \rightarrow 0$ that term vanishes, the minimizer degenerates to a discontinuous switch, and the shooting Jacobian becomes ill-conditioned wherever S crosses zero. Since the linear region in Eq. (10) has width 2ϵ , the parameter sets directly how far the costates may be perturbed before the predicted switching structure changes.

The classical use of this family is offline continuation in which the energy-optimal problem is solved first at $\epsilon = 1$, and ϵ is then reduced in steps, each solve warm-started from the converged costates of the previous one, until $\epsilon = 0$ recovers the fuel-optimal solution [36]. Consecutive problems differ only slightly, so the previous solution is a good initial guess for the next, and the chain succeeds where a direct attack on the bang-bang problem would not. That procedure suits the computation of a single trajectory on the ground, but transfers poorly to a receding-horizon implementation. It is inherently sequential, so one guidance update costs many solves rather than one, with no bound established in advance on how many steps the chain requires or whether it completes. More importantly, the endpoint it is designed to reach is the one that becomes least reliable when navigation degrades: at $\epsilon = 0$ the switching structure is maximally sensitive to costate error, which is precisely the error that grows when measurements are corrupted. Marching to $\epsilon = 0$ on a fixed schedule thus drives the solver toward its most fragile configuration on a timetable unrelated to whether the state estimate can support it. The architecture proposed here retains the family of Eq. (11) but abandons the monotone increase. ϵ is set once per update from the current estimator confidence, so that each update requires a single solve at the selected value.

2.3 Receding-Horizon Indirect Guidance

At each guidance update, the current state estimate $\hat{\mathbf{x}}_k$ is used as the initial condition of a finite-horizon indirect optimal control problem. The corresponding costate boundary-value problem is solved numerically, and only the first segment of the resulting control is applied before the problem is re-solved at the next update. This produces a receding-horizon indirect guidance strategy that retains the structure and efficiency of the indirect formulation while allowing the controller to react to updated navigation information and changing confidence levels.

3 Estimation and Innovation-Adaptive Homotopy

3.1 Measurement Model and Linear Kalman Filter

The translational state is estimated with a linear Kalman filter consistent with the CW dynamics. The discrete-time propagation model is

$$\mathbf{x}_{k+1}^- = \Phi \mathbf{x}_k^+ + \Gamma \mathbf{u}_k, \quad (14)$$

$$P_{k+1}^- = \Phi P_k^+ \Phi^T + Q, \quad (15)$$

where $\Phi = e^{A\Delta t}$ and $\Gamma = \int_0^{\Delta t} e^{A\tau} B d\tau$ are the exact discrete-time state-transition and control matrices over the sampling interval Δt . Q is the process noise covariance matrix that represents the uncertainty, disturbances, and unmodeled errors in the dynamical model.

The measurement model is taken as

$$\mathbf{y}_k = H\mathbf{x}_k^+ + \mathbf{v}_k, \quad H = \begin{bmatrix} I_{3 \times 3} & 0_{3 \times 3} \end{bmatrix}, \quad (16)$$

so that the relative position is measured directly. The nominal innovation and innovation covariance are

$$\boldsymbol{\nu}_k = \mathbf{y}_k - H\mathbf{x}_k^-, \quad (17)$$

$$S_k = HP_k^- H^T + R_{\text{nominal}}. \quad (18)$$

This selection has been purposely made. The aim of the proposed homotopy approach is to improve the mission robustness by letting guidance and navigation share confidence levels and information. By selecting a linear Kalman filter, any improvements reported in the numerical application are due to the proposed approach rather than the selection of the filtering algorithm or the moments prediction technique selected.

3.2 Measurement-Side MTF Adaptation

To improve robustness against non-Gaussian measurement anomalies, an MTF correction is constructed from the mismatch between the instantaneous innovation and the nominal innovation covariance, inspired by the prior works of the authors [34, 37–39]. First, define

$$S_{k,\text{raw}} = \boldsymbol{\nu}_k \boldsymbol{\nu}_k^T - S_k. \quad (19)$$

The MTF matrix is then formed by retaining only positive diagonal excess terms,

$$S_{k,\text{MTF}} = \text{diag}(\max(\text{diag}(S_{k,\text{raw}}), 0)). \quad (20)$$

This construction identifies measurement directions whose observed innovation energy exceeds the nominal expectation and inflates only those axes. The effective measurement covariance becomes

$$R_{k,\text{eff}} = R_{\text{nominal}} + S_{k,\text{MTF}}, \quad (21)$$

and the Kalman gain is computed as

$$K_k = P_k^- H^T (HP_k^- H^T + R_{k,\text{eff}})^{-1}. \quad (22)$$

The state and covariance updates are then

$$\mathbf{x}_k^+ = \mathbf{x}_k^- + K_k \boldsymbol{\nu}_k, \quad (23)$$

$$P_k^+ = (I - K_k H) P_k^- (I - K_k H)^T + K_k R_{k,\text{eff}} K_k^T. \quad (24)$$

In the present implementation, the MTF term is used as a measurement-side innovation-covariance inflation. Its role is to reduce the filter's trust in suspicious measurements while preserving sensitivity along well-behaved axes.

3.3 Innovation-Based Confidence Score

A raw normalized innovation score is computed as

$$s_k^{\text{NIS}} = \boldsymbol{\nu}_k^T S_k^{-1} \boldsymbol{\nu}_k. \quad (25)$$

To capture not only the instantaneous innovation but also the severity of the MTF correction, a relative MTF magnitude is introduced:

$$\rho_k^{\text{MTF}} = \frac{\text{tr}(S_{k,\text{MTF}})}{\text{tr}(S_k)}. \quad (26)$$

These quantities are combined into a scalar adaptation score,

$$s_k^{\text{comp}} = \max(s_k^{\text{NIS}} - m, 0) + \kappa_{\text{MTF}} \rho_k^{\text{MTF}}, \quad (27)$$

where m is the measurement dimension and κ_{MTF} is a tuning parameter. In practice, a first-order low-pass filter is applied to reduce sensitivity to isolated spikes:

$$\bar{s}_k = \rho_s \bar{s}_{k-1} + (1 - \rho_s) s_k^{\text{comp}}. \quad (28)$$

3.4 Innovation-Adaptive Homotopy Scheduling

The filtered confidence score is mapped online to the homotopy parameter. A generic exponential map is

$$\epsilon_k = \epsilon_{\min} + (\epsilon_{\max} - \epsilon_{\min}) (1 - e^{-\beta \bar{s}_k}), \quad (29)$$

with $\beta > 0$ a sensitivity coefficient. In the implementation used for the numerical study, this mapping may also include a deadband and saturation to prevent small nominal innovation fluctuations from triggering unnecessary smoothing. Regardless of the exact tuning, the interpretation is the same when the estimator is confident, i.e., ϵ_k remains near the aggressive low value, and the controller behaves close to the fuel-optimal bang-bang regime. When the innovation and MTF activity increase, ϵ_k rises, informing that the filter is experiencing degraded measurements with noise influence above the assumed nominal level R_{nominal} . Consequently, the control law is automatically relaxed toward a smoother and more conservative regime.

The proposed coupling creates two distinct robustness mechanisms. The first is the estimator-side mechanism. The MTF term inflates the effective measurement covariance when anomalous measurements are detected, preventing the filter from overreacting to corrupted data. The second is the guidance-side mechanism. The innovation-based homotopy scheduling changes the numerical structure of the indirect optimal control problem itself, making the commanded trajectory less aggressive

when state confidence degrades. Together, these two loops produce a tightly-coupled estimation and guidance architecture in which both the state estimate and the control regularization respond to the same confidence information. Figure 1 summarizes the logic of the proposed tightly-coupled architecture. In the conventional arrangement, the estimator passes a state estimate to the guidance law and nothing else; the guidance law has no knowledge of how much that estimate can be trusted, and solves the same optimization problem whether the navigation solution is clean or corrupted. The architecture proposed here adds a second channel carrying that missing information, and uses it to reshape the guidance problem rather than merely to filter its input.




Fig. 1 System-level schematic of the proposed tightly-coupled adaptive homotopy estimation and guidance architecture.

Reading the loop clockwise from the lower left, the *Estimation and Sensing* block receives the degraded measurement $\mathbf{y}_k$ from the relative motion dynamics together with its own prior $\hat{\mathbf{x}}_{k-1}$, and performs a Kalman update with measurement-side MTF covariance inflation. It emits two distinct products. The first is the state estimate $\hat{\mathbf{x}}_k$, which follows the conventional path and supplies the initial condition of the guidance problem. The second is a confidence level, formed from the normalized innovation together with the MTF inflation ratio, which quantifies the degree to which the incoming measurements are consistent with the filter’s own predictions. It is this second output that a conventional architecture ignores.

The confidence level is consumed by the *Innovation-Adaptive Homotopy Scheduler*, which maps it to a single scalar, the homotopy parameter ϵ_k , once per guidance update. The scheduler does not alter the state estimate, the dynamics, the cost functional, or the terminal conditions. Its sole authority is over which member of the problem family of Eq. (11) the solver is asked to solve at this update.

The *Receding-Horizon Indirect Solver* receives both signals: $\hat{\mathbf{x}}_k$ as the initial condition and ϵ_k as the smoothing level. The inset illustrates what ϵ_k controls. At low

ϵ_k the solution approaches the discontinuous bang-bang profile of the fuel-optimal limit (dashed); at high ϵ_k the same solver returns a smooth, saturated-linear profile (solid). These are not two different controllers, but two members of one family, and the solver moves between them by solving the boundary-value problem at the requested ϵ_k . In implementation, this is a single-shooting problem in the six initial costates that has been warm-started from the previous update, so that one solve is required per cycle rather than the continuation sequence the homotopy family would classically demand. Because the control bound is already enforced analytically by the saturation in Eq. (10), the solver sees no inequality constraints. The resulting commanded control $\mathbf{u}_k$ is applied to the chaser, closing the loop through the relative motion of the chaser and the uncooperative target and producing the next degraded measurement.

The lower panel states the sign convention that gives the loop its behaviour. A high ϵ_k yields conservative, smooth, numerically robust control that costs additional propellant; a low ϵ_k yields aggressive, fuel-optimal, bang-bang control that is sensitive to costate error. Because ϵ_k is driven by estimator confidence, the loop retreats toward the robust end of this family precisely when the navigation solution degrades, and relaxes toward the efficient end as confidence recovers. The mechanism is therefore self-regulating in a specific and narrow sense: the quantity that makes the guidance problem hard to solve—costate sensitivity near the switching surface—is the same quantity that grows when measurements are corrupted, and ϵ_k is the single parameter that bounds it.

Two properties of this arrangement are worth making explicit. First, the additional channel is scalar. Only one number crosses from the estimator to the guidance layer beyond the state estimate itself, so the coupling adds no state, no tuning beyond the mapping itself, and negligible computational cost relative to the shooting solve. Second, the coupling is unidirectional at each update: confidence sets ϵ_k , and ϵ_k conditions the solve, but the guidance layer does not feed back into the filter. The estimator remains a standard Kalman filter with MTF inflation and can be replaced independently, which is what permits the estimator ablation reported in Section 4.

4 Numerical Results and Discussion

The terminal rendezvous case considered in this study is not an abstract, standalone benchmark, but is derived from the proximity operations environment developed in our prior vision-based relative navigation work. In that earlier effort, a tumbling target spacecraft was modeled in Blender and observed by a virtual camera during a fly-around sequence, and the relative pose was estimated using an adaptive unscented Kalman filtering framework [38, 39]. The present work builds on that same environment by treating the resulting relative geometry and terminal approach segment as the guidance and control testbed. The scenario is motivated by autonomous rendezvous with large uncooperative spacecraft and debris objects, with particular emphasis on the ENVISAT target problem that has become a representative benchmark for active debris removal and on-orbit servicing studies [40–43]. For the current study, the relative motion is modeled in the Clohessy-Wiltshire frame, and the chaser is initialized

at

$$\mathbf{x}_0 = [0.03031809, 0, 31.16639, -0.02963377, 0.04570523, 0]^T,$$

with position expressed in meters and velocity in meters per second. The terminal target state is taken as the origin, corresponding to rendezvous with the target-centered frame over an 800 s horizon.

Measurement degradation is introduced in two overlapping stages. Outside the anomaly, the position measurement follows Eq. (16) with the nominal covariance R_{nominal} . Over the interval $t = [250, 400]$ s, a constant factor inflates measurement-noise covariance,

$$R_k = \kappa R_{\text{nominal}}, \quad \kappa = 300, \quad (30)$$

raising the per-axis standard deviation from 1 m to ≈ 17.3 m. From $t_b = 325$ s until the end of the same window, a constant additive bias is superimposed, so that the measurement model becomes

$$\mathbf{y}_k = H\mathbf{x}_k + \mathbf{b} + \mathbf{v}_k, \quad \mathbf{v}_k \sim \mathcal{N}(\mathbf{0}, R_k), \quad \mathbf{b} = [10, -8, 12]^T \text{ m}. \quad (31)$$

The degradation parameters, together with the simulation, scheduler, and solver settings, are collected in Table 1.

The control input is bounded by a defined maximum acceleration, and open-loop energy-optimal and fuel-optimal solutions are first computed as reference trajectories. Then, three closed-loop architectures were compared under the same measurement degradation. A plain Kalman filter with fixed bang-bang control, an MTF-based Kalman filter with fixed bang-bang control, and the proposed MTF-based filter coupled with innovation-adaptive homotopy scheduling. For reference, the open-loop energy-optimal and fuel-optimal trajectories were also computed offline.

The open-loop baselines achieved

$$\Delta v_{\text{energy}} = 0.150239 \text{ m/s}, \quad \Delta v_{\text{fuel}} = 0.112800 \text{ m/s},$$

with negligible terminal miss. In the closed loop, however, both fixed- ϵ controllers failed to satisfy the terminal objective under degraded sensing. Specifically, the plain-KF and MTF-KF bang-bang controllers each produced a terminal miss of 4.4040 and 7.6134 meters, respectively, despite maintaining a control effort around the open-loop fuel-optimal benchmark. By contrast, the proposed MTF-adaptive homotopy controller reduced the terminal miss to 0.3350 m while using 0.14668 m/s of cumulative control effort, corresponding to a fuel penalty of 30.039% relative to the open-loop fuel-optimal solution. The improvement in terminal accuracy is roughly an order of magnitude, a factor of 13 against the plain filter and 23 against the MTF filter at $\epsilon = 0$.

The visual results in Figs. 2 and 3 are consistent with the quantitative trends discussed above. Figure 2 compares the resulting relative trajectories and the corresponding position and estimation errors. In the 3D and x - z views, the proposed MTF-adaptive homotopy trajectory remains close to the reference fuel-optimal path while preserving a smoother and more robust terminal approach under degraded sensing. The position error history further shows that, although all controllers begin from the same initial condition, only the adaptive architecture maintains reliable convergence to the target during the degradation interval. The estimation error plot also

Table 1 Simulation, degradation, scheduler, and solver parameters used in the study.

Parameter	Value	Description
<i>Scenario and Dynamics</i>		
T_f	800 s	Terminal rendezvous horizon
Δt	1 s	Simulation / filter update step
n	1.04218×10^{-3} rad/s	Target mean motion (LEO, ≈ 100.5 min)
$u_{\max}$	8×10^{-4} m/s ²	Maximum control acceleration
$\ r_0\ $	31.17 m	Initial relative distance
x_f	0	Terminal target state
<i>Estimation</i>		
R_{nominal}	$\text{diag}(1, 1, 1)$ m ²	Nominal measurement covariance
Q	$\text{diag}(10^{-6} \mathbf{I}_3, 10^{-8} \mathbf{I}_3)$	Process covariance (m ² , m ² /s ²)
P_0	$\text{diag}(1 \cdot \mathbf{I}_3, 10^{-4} \mathbf{I}_3)$	Initial covariance (m ² , m ² /s ²)
m_{meas}	3	Measurement dimension
<i>Measurement Degradation</i>		
$[t_{d,0}, t_{d,1}]$	[250, 400] s	Anomaly window
κ	300	Measurement-noise inflation factor
$\mathbf{b}$	[10, -8, 12] m	Sustained measurement bias
<i>Adaptive Scheduler</i>		
β	0.25	Innovation-to-homotopy sensitivity
ρ_{score}	0.85	Low-pass filter coefficient for score
k_{MTF}	12.0	Weight on MTF severity in composite score
$\text{score}_{\max}$	50.0	Saturation limit for composite score
$\epsilon_{\min}, \epsilon_{\max}$	0.0, 1.0	Homotopy parameter bounds
α_{ϵ}	0.02	Homotopy smoothing/blending factor
<i>Guidance Solver</i>		
t_{resolve}	10 s	Receding-horizon re-solve interval
$t_{\text{min,rem}}$	20 s	Minimum remaining horizon for re-solve
Algorithm	Levenberg–Marquardt	fsolve root-finding algorithm
FunctionTolerance	10^{-10}	fsolve termination on residual
StepTolerance	10^{-10}	fsolve termination on step
MaxIterations	300	fsolve iteration limit
δ_{tol}	10^{-4}	Shooting residual gate for an accepted solve
RelTol, AbsTol	$10^{-9}, 10^{-11}$	Integration tolerances, Hamiltonian system
<i>Terminal Success Criteria</i>		
ϵ_r	1 m	Terminal position tolerance
ϵ_v	0.01 m/s	Terminal velocity tolerance

illustrates that the MTF-based filtering framework suppresses the growth of state estimation error during the imposed anomaly window, with the adaptive case exhibiting the smallest sustained error and the cleanest recovery behavior.

Figure 3 highlights the corresponding control-level and adaptation-level behavior. The instantaneous thrust and cumulative- Δv plots show that the fixed- ϵ controllers retain a more brittle bang-bang structure, whereas the adaptive controller deliberately departs from the purely fuel-optimal policy and expends additional control effort in exchange for substantially improved terminal robustness. The scheduler input plot clarifies the mechanism behind this behavior. Once the innovation statistics increase during the degradation interval, the composite confidence metric rises sharply, triggering an increase in the homotopy parameter. This behavior is shown explicitly in the homotopy parameter history, where ϵ is elevated during periods of degraded confidence and then relaxes as the sensing condition recovers. Together, these figures confirm that

the proposed method does not merely improve estimation in isolation, but actively reshapes the closed-loop guidance law in response to estimator confidence.




Fig. 2 Trajectory and state-level comparison of the open-loop and closed-loop rendezvous solutions.




Fig. 3 Control effort and adaptation behavior for the closed-loop rendezvous case. The upper panels compare the instantaneous thrust magnitude and cumulative Δv , showing the robustness versus fuel trade introduced by the adaptive controller. The lower panels show the innovation-based scheduler inputs and the resulting homotopy parameter history, demonstrating how increased innovation activity during sensor degradation drives the controller toward a smoother and more conservative regime.

Figure 4 focuses on the final portion of the position error history and makes the terminal performance gap especially clear. While the fixed ϵ architectures flatten out at a nonzero miss distance, the MTF-adaptive homotopy controller continues reducing the terminal error and converges much closer to the target. Figure 5 zooms in on the scheduler inputs during the degradation and recovery window, showing more clearly how the raw innovation, MTF activity, and composite score evolve together. Close-up views reinforce the interpretation that the proposed framework improves both the internal adaptation logic and the final rendezvous accuracy.




Fig. 4 Close-up view of the terminal position error near the end of the rendezvous.




Fig. 5 Close-up view of the innovation-based scheduler during the sensor degradation and recovery.

4.1 Role of MTF and Adaptive Homotopy

Although the MTF mechanism improved the estimator’s robustness to faulty measurements, the controller itself remained fixed in the aggressive bang-bang regime, so the overall trajectory remained brittle. To isolate the role of MTF within the adaptive architecture, an additional comparison study was performed in which the homotopy scheduler remained active, but the MTF contribution was disabled. In that case, the plain adaptive controller couldn’t succeed, but achieving a lower terminal miss of

$$2.3309 \text{ m},$$

but at a substantially higher control cost of

$$0.35181 \text{ m/s},$$

corresponding to a 211.890% fuel penalty relative to the open-loop fuel-optimal baseline. When MTF was restored within the same adaptive architecture, the terminal miss was reduced to

$$0.3350 \text{ m},$$

while the cumulative control effort decreased to

$$0.14668 \text{ m/s},$$

corresponding to a reduced fuel penalty of 30.039%. Thus, adding MTF to the adaptive controller reduced the terminal miss from approximately 2.3309 m to 0.3350 m while also lowering the control effort by almost 58%. These results clarify the relative roles of the two mechanisms. The adaptive schedule supplies most of the accuracy recovery, cutting the miss from 4.4040 m to 2.3309 m even on the plain filter, but it cannot cross the tolerance on its own and pays a 212% fuel penalty in the attempt. On the other hand, MTF is what makes that schedule both affordable and sufficient. By cleaning the innovation sequence from which the score is computed, it cuts the miss a further sevenfold and the propellant by 58%. Therefore, it can be stated that neither element meets the requirement alone.

Figures 6 and 7 provide a visual interpretation of this comparative result. Figure 6 compares the control effort, cumulative Δv , scheduler inputs, and homotopy histories for the two adaptive controllers. The plain adaptive case reacts strongly to the degradation interval and remains significantly more conservative even after recovery, whereas the MTF-adaptive case exhibits a more structured confidence response and a less wasteful control history. Figure 7 further emphasizes this distinction through close-up views of the composite score and the terminal position error. In particular, the zoomed terminal-error plot makes clear that adaptive homotopy alone leaves the trajectory above the tolerance line, and that the addition of MTF is what carries it below.




Fig. 6 Adaptive-only comparison comparing the plain adaptive-homotopy controller and the MTF-adaptive controller.




Fig. 7 Close-up view of the adaptive-only comparison study.

Taken together, these results support the main claims of the paper. The central benefit comes from tightly coupling estimator confidence to the homotopy regularization of the indirect solver. The gain arises from allowing the controller to become less aggressive when confidence deteriorates, and from the MTF mechanism that makes

the confidence signal trustworthy enough to act on, by refining the innovation statistics, reducing sensitivity to axis-specific anomalous measurements, and improving the computational reliability of the adaptive closed-loop guidance process.

4.2 Benchmark Controllers

Eleven receding-horizon controllers are compared on a common degraded-measurement realization. They fall into four groups. The first isolates the estimator: a plain Kalman filter and the MTF filter are each run with the homotopy disabled, $\epsilon = 0$. The second is a fixed- ϵ ladder, $\epsilon \in \{0, 0.1, 0.25, 0.5, 0.75, 1\}$, in which the smoothing level is chosen offline and held for the whole mission. The third group contains two alternative schedules that vary ϵ online but drive it from something other than filter confidence, described below. The fourth is the proposed method, in which the MTF filter supplies a confidence score that sets ϵ at each replan.

4.2.1 Deterministic Continuation

One of the two schedulers from the third group is deterministic continuation. It follows a fixed clock and the requested smoothing level descends linearly, $\epsilon_{\text{req}}(t) = \max(0, 1 - t/T_d)$ with $T_d = 400$ s, so the solver is asked to return to the fuel-optimal end of the homotopy on a schedule fixed before the mission begins and independent of what the measurements subsequently do. This is the natural continuation strategy when the problem is expected to become easier over time, and it is included precisely because that expectation fails here. It is expected the descent would drive ϵ_{req} through its lowest values during the 250–400 s anomaly window, when the estimate is least trustworthy, and the shooting problem is hardest.

4.2.2 Failure-Triggered Homotopy

This method is reactive rather than scheduled. Each replan first requests $\epsilon = 0$; if the shooting solver fails to converge, the request is raised to the next step of a fixed ladder $\{0, 0.05, 0.1, 0.2, 0.4, 0.6, 0.8, 1\}$, and the solve is retried, repeating until a step converges. The applied ϵ is therefore whichever step happened to work, and the strategy is guaranteed to produce a solution provided the top rung is tractable.

4.3 Closed-Loop Comparison

Figures 8 and 9 plot the four outcome measures against the fixed- ϵ ladder. Together they show that the smoothing level does not trade one quantity against another so much as govern four quantities that saturate at different points, and that no single fixed value sits at the best point of all four. The remaining controllers, which have no position on the ϵ axis, are reported in Tables 2 and 3.

Figure 8 pairs replan reliability, which represents the percentage of the converged runs, with fuel consumption. Reliability is strongly nonlinear in ϵ , essentially flat at roughly 5% below $\epsilon = 0.25$, rising steeply through the middle of the range, and reaching 100% only at $\epsilon = 1$, with $\epsilon = 0.75$ still short at 94.8%. Fuel consumption, by contrast, climbs steadily across the whole range, from 0.11320 to 0.15312 m/s. A fixed choice

must therefore buy its last few percent of reliability on the steepest part of the fuel curve. On the other hand, the proposed schedule reaches 100% at 0.14668 m/s, 4.2% below $\epsilon = 1$, the only fixed setting that also converges on every replan.

Figure 9 pairs solver latency with terminal accuracy, and the two stop improving at different places. Terminal error falls by more than an order of magnitude between $\epsilon = 0.1$ and $\epsilon = 0.5$, crossing the tolerance, and is then flat. Every setting from $\epsilon = 0.5$ upward lands within the same terminal miss, so further smoothing buys no additional accuracy. Latency keeps improving above that point, falling from 1386 ms at $\epsilon = 0.5$ to 65 ms at $\epsilon = 1$ in the 95th percentile. Above the accuracy floor, the ladder is thus trading fuel for latency alone. The proposed schedule attains the floor at 97.6 ms, against roughly 274 ms for the ladder interpolated at the same mean ϵ .

For deterministic continuation, the accepted smoothing level consequently tracks the request down to 0.4 and then freezes, because below that level the solver stops converging and the last successful solution is retained; only 27.3% of replans succeed thereafter, and the resulting 1.72 m terminal error misses the tolerance. For failure-triggered homotopy, every replan eventually converges, and the mission succeeds. However, the cost is paid in attempts. Convergence requires between one and eight solves per replan, 4.026 on average, and because the retries are serial, the time cost deteriorates. A median of 2542.7 ms, a 95th percentile of 12760 ms, and a maximum of 19634 ms against a 10 s replan period. The method is reliable in the sense that it always answers, and infeasible in the sense that it cannot be relied upon to answer in time. It is the clearest illustration of why reliability and latency must be reported together.

Two properties of the proposed method’s placement in these figures are particularly important. It has no fixed ϵ , so it is drawn at the mean of its accepted smoothing level, $\epsilon = 0.83$, with the shaded band spanning the full range that the schedule visited, $\epsilon = 0.69$ to $\epsilon = 1.00$. That mean summarizes the schedule but does not predict its cost. The adaptive law runs at $\epsilon = 0.91$ on average inside the anomaly window and relaxes outside it, so the same mean is reached by a path that concentrates smoothing where the solver needs it. This is why the fuel star in Fig. 8 sits slightly above the ladder read at the same abscissa while the latency star in Fig. 9 sits well below it. In both figures the intended comparison is the horizontal one, against the fixed setting that achieves the same reliability.

Moreover, two qualifications follow from the tables. First, the proposed method is not the cheapest successful controller on this realization: fixed $\epsilon = 0.5$ and $\epsilon = 0.75$ both meet the tolerance at 0.13214 and 0.13914 m/s, below the proposed method’s 0.14668 m/s. What separates them is the tail, which is their maximum solve times reach 4902.8 and 3719.4 ms against 150.83 ms and their sub-unity solve rates. Second, the estimator alone is not sufficient. The MTF at $\epsilon = 0$ fails as badly as the plain filter, while the plain filter driving the adaptive schedule fails differently, spending 0.35181 m/s, a 212% fuel penalty, and still missing the tolerance, because a confidence signal computed from a filter that has not rejected the corrupted measurements schedules the wrong ϵ . The claim supported here is therefore narrower than superiority on every axis. The adaptive schedule attains the reliability and bounded latency





Fig. 8 Successful-replan rate (left axis) and control expenditure (right axis) across the fixed- ϵ ladder, under the common degraded-measurement realization. Reliability saturates only at $\epsilon = 1$ while fuel rises throughout, so a fixed setting pays for its last increment of reliability at the steepest part of the fuel curve. The proposed method has no fixed ϵ and is plotted at the mean of its accepted smoothing level, with the shaded band spanning the range that schedule visited.

of the most conservative fixed setting at lower fuel cost, without requiring that the appropriate ϵ be identified in advance.

Table 2 Closed-loop performance under the common degraded-measurement realization.

Method	Δv [m/s]	Fuel penalty [%]	Position error [m]	Velocity error [m/s]	Success
Plain KF, fixed $\epsilon = 0$	0.11240	-0.355	4.4040	6.1321×10^{-3}	No
MTF KF, fixed $\epsilon = 0$	0.11320	0.355	7.6134	1.5459×10^{-2}	No
Plain KF, adaptive ϵ	0.35181	211.890	2.3309	1.4499×10^{-2}	No
MTF KF, adaptive ϵ (proposed)	0.14668	30.039	0.3350	1.0545×10^{-3}	Yes
MTF KF, fixed $\epsilon = 0.1$	0.11837	4.942	8.8031	1.5831×10^{-2}	No
MTF KF, fixed $\epsilon = 0.25$	0.11984	6.239	2.2066	2.4128×10^{-3}	No
MTF KF, fixed $\epsilon = 0.5$	0.13214	17.145	0.33765	1.0276×10^{-3}	Yes
MTF KF, fixed $\epsilon = 0.75$	0.13914	23.349	0.33549	1.0489×10^{-3}	Yes
MTF KF, fixed $\epsilon = 1$	0.15312	35.743	0.33468	1.0440×10^{-3}	Yes
MTF KF, deterministic continuation	0.13073	15.898	1.7211	2.6383×10^{-3}	No
MTF KF, failure-triggered homotopy	0.13917	23.381	0.33751	1.0543×10^{-3}	Yes

**Fig. 9** Solver latency (left axis) and terminal position error (right axis) across the same ladder, both on logarithmic scales. The dash-dotted tolerance line refers to the right axis only. Accuracy reaches a floor at $\epsilon = 0.5$ while latency continues to improve to $\epsilon = 1$, so above the floor the ladder trades fuel for latency alone. The proposed method is plotted as in Fig. 8.

4.4 Code and Media Availability

Rendered Blender videos of the terminal rendezvous scenario are provided as supplementary media and are available together with the implementation repository.

- Code repository: [GitHub repository](#).
- Supplementary rendezvous video: [Unlisted video link](#).

Table 3 Receding-horizon shooting-solver reliability and computational burden. Mean attempts is the number of solver calls per replan; only the failure-triggered schedule retries.

Method	Solve rate [%]	Median [ms]	95th percentile [ms]	Maximum [ms]	Mean attempts
Plain KF, fixed $\epsilon = 0$	2.597	3235.4	5401.3	7007.3	1.000
MTF KF, fixed $\epsilon = 0$	5.195	3459.6	4355.0	5648.7	1.000
Plain KF, adaptive ϵ	75.325	140.56	1173.4	3449.3	1.000
MTF KF, adaptive ϵ (proposed)	100.000	65.04	97.632	150.83	1.000
MTF KF, fixed $\epsilon = 0.1$	5.195	276.23	4544.7	5226.6	1.000
MTF KF, fixed $\epsilon = 0.25$	14.286	521.05	3311.0	3456.8	1.000
MTF KF, fixed $\epsilon = 0.5$	53.247	178.98	1386.3	4902.8	1.000
MTF KF, fixed $\epsilon = 0.75$	94.805	81.247	512.42	3719.4	1.000
MTF KF, fixed $\epsilon = 1$	100.000	50.534	65.455	80.495	1.000
MTF KF, deterministic continuation	27.273	359.26	3545.9	3911.0	1.000
MTF KF, failure-triggered homotopy	100.000	2542.7	12760	19634	4.026

5 Conclusion

This paper presented a tightly coupled estimation and guidance architecture for robust low-thrust rendezvous in which estimator confidence is used not only to improve the state estimate, but also to reshape the numerical structure of the onboard indirect optimal control problem. The proposed framework combines a receding-horizon indirect low-thrust guidance law in the Clohessy–Wiltshire frame with a linear Kalman filter employing measurement-side MTF covariance inflation. A composite innovation metric, formed from the normalized innovation and the MTF activity, is mapped online to the homotopy continuation parameter, so that the smoothing applied to the switching function increases when sensing degrades and relaxes as confidence recovers.

Under a degraded-measurement realization combining inflated noise and a sustained measurement bias, the proposed controller was the only configuration to meet the terminal tolerance while converging on every replan. The ablation is more informative than either result on its own, and it does not support attributing robustness to a single mechanism. Neither component is sufficient alone. The MTF filter with the homotopy disabled failed to meet tolerances, worse than the plain filter under the same setting, because rejecting corrupted measurements does nothing for a guidance problem that has become numerically intractable. The plain filter driving the adaptive schedule also failed, because a confidence signal computed from a filter that has not rejected the corrupted measurements schedules the wrong smoothing level. Only the two together met the tolerance. The contribution is therefore properly described as an interaction rather than as a filtering improvement with an adaptive refinement layered on top. The schedule requires a trustworthy confidence signal, and the filter requires a consumer that acts on it.

Moreover, the computational results support the practical feasibility of the proposed method. Every replan converged on the first attempt, and the MTF update contributes negligibly to the per-update cost, since a covariance inflation costs a small multiple of a filter update against a shooting-based re-optimization. The comparison with alternative schedules is where the latency argument becomes sharp.

Beyond it, the method should be extended past the linear CW setting to higher-fidelity nonlinear relative dynamics, richer measurement models, and hardware-targeted timing studies, and the regime in which MTF contributes most strongly structured, axis-dependent outliers and model mismatch deserves separate treatment. A further open item is that the residual fuel penalty and a recovery law with a stronger

downward drive, or an explicit fuel-penalty term in the mapping, may close part of that gap. Finally, a particularly important extension is to embed explicit keep-out and no-go regions into the adaptive guidance problem, including exclusion cones generated by the tumbling motion of solar panels and other appendages of derelict spacecraft during close-proximity operations.

Declarations

- Funding: [To be completed.]
- Conflict of interest: The authors declare no competing interests.
- Author contributions: [To be completed.]

References

- [1] Albee, K., Oestreich, C., Specht, C., Teran Espinoza, A., Todd, J., Hokaj, I., Lampariello, R., Linares, R.: A robust observation, planning, and control pipeline for autonomous rendezvous with tumbling targets. *Frontiers in Robotics and AI* **8**, 641338 (2021) <https://doi.org/10.3389/frobt.2021.641338>
- [2] Servadio, S.: High order filters for relative pose estimation of an uncooperative target. PhD thesis, Politecnico di Milano (2017)
- [3] Servadio, S., Cavenago, F., Di Lizia, P., Massari, M.: Nonlinear prediction in marker-based spacecraft pose estimation with polynomial transition maps. *Journal of Spacecraft and Rockets* **59**(2), 511–523 (2022) <https://doi.org/10.2514/1.A35068> <https://doi.org/10.2514/1.A35068>
- [4] Morante, D., Sanjurjo Rivo, M., Soler, M.: A survey on low-thrust trajectory optimization approaches. *Aerospace* **8**(3), 88 (2021) <https://doi.org/10.3390/aerospace8030088>
- [5] Wang, Z.: A survey on convex optimization for guidance and control of vehicular systems. *Annual Reviews in Control* **57**, 100957 (2024) <https://doi.org/10.1016/j.arcontrol.2024.100957>
- [6] Malyuta, D., Yu, Y., Elango, P., Acikmese, B.: Advances in trajectory optimization for space vehicle control. *Annual Reviews in Control* **52**, 282–315 (2021) <https://doi.org/10.1016/j.arcontrol.2021.04.013>
- [7] Gurfil, P.: Spacecraft rendezvous using constant-magnitude low thrust. *Journal of Guidance, Control, and Dynamics* **46**(11), 2183–2191 (2023) <https://doi.org/10.2514/1.G007472>
- [8] Malyuta, D., Reynolds, T.P., Szmuk, M., Lew, T., Bonalli, R., Pavone, M., Acikmese, B.: Convex optimization for trajectory generation: A tutorial on

- generating dynamically feasible trajectories reliably and efficiently. *IEEE Control Systems Magazine* **42**(5), 40–113 (2022) <https://doi.org/10.1109/MCS.2022.3187542>
- [9] Leomanni, M., Rogers, E., Gabriel, S.B.: Explicit model predictive control approach for low-thrust spacecraft proximity operations. *Journal of Guidance, Control, and Dynamics* **37**(6), 1780–1789 (2014) <https://doi.org/10.2514/1.G000477>
 - [10] Bhagat, M.R.: Convex guidance for envisat rendezvous. Msc thesis, Delft University of Technology (2016)
 - [11] Hofmann, C., Topputo, F.: Closed-loop guidance for low-thrust interplanetary trajectories using convex programming. In: *Proceedings of the 11th International ESA Conference on Guidance, Navigation and Control Systems* (2021). Conference paper.
 - [12] Hofmann, C., Topputo, F.: Rapid low-thrust trajectory optimization in deep space based on convex programming. *Journal of Guidance, Control, and Dynamics* **44** (2021) <https://doi.org/10.2514/1.G005839> . Please verify final issue number and page range.
 - [13] Hofmann, C., Topputo, F.: Homotopic approach for high-fidelity convex low-thrust trajectory optimization. *IEEE Transactions on Aerospace and Electronic Systems* (2025) <https://doi.org/10.1109/TAES.2025.3572466> . Early access / online-first article; verify final volume and page range.
 - [14] Pavanello, Z., Pirovano, L., Armellin, R., De Vittori, A., Di Lizia, P.: Collision avoidance maneuver optimization during low-thrust propelled trajectories. *Astrodynamics* **9**(2), 247–271 (2025) <https://doi.org/10.1007/s42064-024-0227-3>
 - [15] Malyuta, D., Acikmese, B.: Fast homotopy for spacecraft rendezvous trajectory optimization with discrete logic. *Journal of Guidance, Control, and Dynamics* **46**, 1262–1279 (2023) <https://doi.org/10.2514/1.G006295>
 - [16] Servadio, S., Armellin, R., Linares, R.: Koopman-operator control optimization for relative motion in space. *Journal of Guidance, Control, and Dynamics* **46**(1) (2023) <https://doi.org/10.2514/1.G007217> . Please verify final page range.
 - [17] Li, H., Chen, S., Izzo, D., Baoyin, H.: Deep networks as approximators of optimal low-thrust and multi-impulse cost in multitarget missions. *Acta Astronautica* **166**, 469–481 (2020) <https://doi.org/10.1016/j.actaastro.2019.09.023>
 - [18] Cheng, L., Wang, Z., Jiang, F., Zhou, C.: Real-time optimal control for spacecraft orbit transfer via multiscale deep neural networks. *IEEE Transactions on Aerospace and Electronic Systems* **55**(5), 2436–2450 (2019) <https://doi.org/10.1109/TAES.2018.2889571>

- [19] Zavoli, A., Federici, L.: Reinforcement learning for robust trajectory design of interplanetary missions. *Journal of Guidance, Control, and Dynamics* (2021) <https://doi.org/10.2514/1.G005794> . Please verify final volume, issue, and page range.
- [20] Yang, B., Yao, Z., Li, S.: Low-thrust trajectory optimization via knowledge-enhanced deep neural network. *IEEE Transactions on Aerospace and Electronic Systems* (2026) <https://doi.org/10.1109/TAES.2026.3652990> . Early access / online-first article; verify final volume and page range.
- [21] Wang, K., Armellin, R., Evans, A., Holt, H., Chen, Z.: Learning-Based Stable Optimal Guidance for Spacecraft Close-Proximity Operations. *arXiv preprint. arXiv:2501.01096* (2025)
- [22] Zhang, Z., Topputo, F.: Neural Approximators for Low-Thrust Trajectory Transfer Cost and Reachability. *arXiv preprint. arXiv:2508.02911; AAS 25-542* (2025)
- [23] Wang, X., Dou, Y., Yi, X., Zhang, Y., Li, X., Li, B., Peng, H., Wang, L., Teo, K.L.: When optimal control meets neural network: A comprehensive survey. *Archives of Computational Methods in Engineering* (2026) <https://doi.org/10.1007/s11831-026-10499-6>
- [24] Servadio, S., Arnas, D., Linares, R.: Dynamics near the three-body libration points via koopman operator theory. *Journal of Guidance, Control, and Dynamics* **45**(10), 1800–1814 (2022). <https://doi.org/10.2514/1.G006519>
- [25] Taheri, E., Kolmanovsky, I., Atkins, E.: Enhanced smoothing technique for indirect optimization of minimum-fuel low-thrust trajectories. *Journal of Guidance, Control, and Dynamics* **39**(11), 2500–2511 (2016) <https://doi.org/10.2514/1.G000379>
- [26] Pan, B., Pan, X., Ma, Y.: A quadratic homotopy method for fuel-optimal low-thrust trajectory design. *Proceedings of the Institution of Mechanical Engineers, Part G: Journal of Aerospace Engineering* (2018) <https://doi.org/10.1177/0954410018761965> . Please verify publisher-formatted page range in the final bibliography.
- [27] Huang, Y., Zhang, L., Zhang, X.: A modified robust homotopic method for the low-thrust fuel-optimal orbital control problem. *Aerospace* **11**(9), 766 (2024) <https://doi.org/10.3390/aerospace11090766>
- [28] Servadio, S., Zanetti, R., Jones, B.A.: Nonlinear filtering with a polynomial series of gaussian random variables. *IEEE Transactions on Aerospace and Electronic Systems* **57**(1), 647–658 (2020). <https://doi.org/10.1109/TAES.2020.3028487>
- [29] Tafazzol, S., Taheri, E., Li, N.: Comparison of Control Regularization Techniques

for Minimum-Fuel Low-Thrust Trajectory Design Using Indirect Methods. arXiv preprint. arXiv:2409.01490; AAS 24-302 (2024)

- [30] Nurre, N.P., Taheri, E.: Comparison of indirect and direct methods for operationally constrained low-thrust trajectories. *Journal of Guidance, Control, and Dynamics* **48**(4) (2025) <https://doi.org/10.2514/1.G008559> . Please verify final page range in the published issue.
- [31] Parrish, N.L.O.: Low thrust trajectory optimization in cislunar and translunar space. PhD thesis, University of Colorado Boulder (2018)
- [32] Mall, K., Grant, M.J., Taheri, E.: Uniform trigonometrization method for optimal control problems with control and state constraints. *Journal of Spacecraft and Rockets* **57**, 995–1007 (2020) <https://doi.org/10.2514/1.A34624>
- [33] Servadio, S., Jang, D., Linares, R.: Operator-based optimal transfers and debris collision assessment in leo. *Journal of Spacecraft and Rockets* **62**(6), 2199–2211 (2025). <https://doi.org/10.2514/1.A36017>
- [34] Candan, B., Soken, H.E.: Robust attitude estimation using imu-only measurements. *IEEE Transactions on Instrumentation and Measurement* **70**, 1–9 (2021) <https://doi.org/10.1109/TIM.2021.3104042>
- [35] Candan, B., Servadio, S.: Tightly-coupled estimation and guidance for robust low-thrust rendezvous via adaptive homotopy. In: AAS/AIAA Astrodynamics Specialist Conference, Whistler, BC, Canada (2026). <https://doi.org/10.48550/arXiv.2605.04481>
- [36] 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**(8), 1501–1510 (2015) <https://doi.org/10.2514/1.G001080> <https://doi.org/10.2514/1.G001080>
- [37] Candan, B., Soken, H.E.: Robust attitude estimation using magnetic and inertial sensors. *IFAC-PapersOnLine* **56**(2), 4502–4507 (2023) <https://doi.org/10.1016/j.ifacol.2023.10.941>
- [38] Candan, B., Servadio, S.: Relative pose estimation of an uncooperative target with camera marker detection. *Aerospace* **12**(5) (2025) <https://doi.org/10.3390/aerospace12050425>
- [39] Candan, B., Servadio, S.: Adaptive relative pose estimation framework with dual noise tuning for safe approaching maneuvers. In: AAS/AIAA Astrodynamics Specialist Conference, Boston, MA (2025). <https://doi.org/10.48550/arXiv.2507.16214>
- [40] Estable, S., Pruvost, C., Ferreira, E., Telaar, J., Fruhnert, M., Imhof, C., Rybus,

- T., Peckover, G., Lucas, R., Ahmed, R., Oki, T., Wygachiewicz, M., Kicman, P., Lukasik, A., Santos, N., Milhano, T., Arroz, P., Biesbroek, R., Wolahan, A.: Capturing and deorbiting envisat with an airbus spacetug. results from the esa e.deorbit consolidation phase study. *Journal of Space Safety Engineering* **7**(1), 52–66 (2020). DOI:<https://doi.org/10.1016/j.jsse.2020.01.003>
- [41] Pensado, E.A., Santos, L.M.G., Jorge, H.G., Sanjurjo-Rivo, M.: Deep learning-based target pose estimation using lidar measurements in active debris removal operations. *IEEE Transactions on Aerospace and Electronic Systems* **59**(5), 5658–5670 (2023). DOI:10.1109/TAES.2023.3262505
- [42] Jaekel, S., Lampariello, R., Rackl, W., De Stefano, M., Oumer, N., Giordano, A.M., Porges, O., Pietras, M., Brunner, B., Ratti, J., Muehlbauer, Q., Thiel, M., Estable, S., Biesbroek, R., Albu-Schaeffer, A.: Design and operational elements of the robotic subsystem for the e.deorbit debris removal mission. *Frontiers in Robotics and AI* **Volume 5 - 2018** (2018). DOI:10.3389/frobt.2018.00100
- [43] European Space Agency (ed.): *Envisat Mission: Opportunities for Science and Applications*. ESA SP, vol. 1218. European Space Agency, Noordwijk, The Netherlands (1998)