

Research Paper

Sensitivity analysis of key modeling factors in angles-only relative orbit navigation for LEO applications

Alessandro Scalvini^{ID}*, Giacomo Borelli^{ID}, Gabriella Gaias^{ID}

Department of Aerospace Science and Technology, Politecnico di Milano, Via La Masa 34, Milano, 20156, Italy

ARTICLE INFO

Keywords:

Angles-only navigation
Relative orbit determination
Sensitivity analysis
Analytical batch least-squares filtering

ABSTRACT

Angles-only relative navigation to non-cooperative targets is constrained by weak observability, placing stringent requirements on filter model fidelity. This study presents a systematic sensitivity analysis of the four key components of a batch least-squares filter architecture: analytical propagation dynamics, maneuver modeling, measurement mapping and batch configuration. Representative near-circular low-Earth-orbit scenarios are examined, including bounded and drifting relative orbits at altitudes where differential drag is significant. The analysis yields quantitative validity bounds and actionable design guidelines. Empirical differential drag modeling is shown to be essential at low altitudes, providing substantial improvement over models ignoring this perturbation. A burn-duration threshold is identified beyond which continuous maneuver modeling must replace the impulsive assumption. Exact and simplified-exact mappings from relative orbital elements to relative position maintain accuracy across the full tested range, whereas first-order rectilinear mappings are acceptable only at close range. Batch length is proven to govern fundamental observability while batch size affects noise rejection. Residual standard deviation and bias are confirmed as a practical filter health diagnostic and the elevation bias is shown to provide a geometric explanation of the estimation error. While developed for batch architectures, the propagation and measurement model recommendations apply directly to sequential filters, providing a unified reference for angles-only navigation system design.

1. Introduction

Distributed space systems, satellite servicing and assembly and active debris removal often share the operational requirement that at least one target satellite be non-cooperative and thus incapable of transmitting navigation-relevant data to the observer(s). Estimating the relative state without support from the target itself or the ground is therefore a core challenge in spacecraft guidance, navigation and control (GNC). Within this context, autonomous angles-only (AO) relative orbit determination has emerged as a viable solution for far-to mid-range scenarios (*i.e.*, from tens of kilometers to hundreds of meters) since it relies exclusively on information acquired onboard the observing spacecraft [1,2] or by a set of satellites building a distributed observer [3,4]. Target detection is performed using a passive vision-based sensor (VBS) suite, which offers low power consumption, limited mass and volumes and broad compatibility with modern spacecraft platforms [5]. These operational and economic advantages are partially offset by the fundamental problem of weak observability [6], which imposes significant constraints on the navigation algorithms and architectures needed for reliable state estimation [1].

The technical evolution of AO navigation systems across demonstrated missions traces a clear progression from ground-based to fully space-borne autonomous implementations. ARGON, conducted within the PRISMA mission [7], employed a ground-based non-linear batch least-squares (LSQ) relative orbit determination (ROD) filter [8] with ground-in-the-loop support [2]. The AVANTI experiment [1], conducted within the FIREBIRD mission [9], advanced this architecture by migrating the navigation task onboard the chaser and replacing the batch LSQ filter with a fully analytical extended Kalman filter (EKF). The StarFOX experiment [10,11], part of the NASA Starling mission [12], represents a further step toward increased autonomy, introducing the Angles-only Relative Trajectory Measurement System [13, 14] for a distributed CubeSat network in which each observer runs an adaptive unscented Kalman filter (UKF) [15] and a multi-hypothesis tracking algorithm [16]. Compared to single-observer architectures [3, 17], this distributed configuration substantially improves observability and ROD accuracy. Beyond the past and present formation flying missions, satellite servicing and assembly [18,19] and active debris removal [20,21] represent future activities, which will certainly require robust autonomous navigation to non-cooperative targets.

* Corresponding author.

E-mail addresses: alessandro.scalvini@polimi.it (A. Scalvini), giacomo.borelli@polimi.it (G. Borelli), gabriella.gaias@polimi.it (G. Gaias).




Fig. 1. Batch LSQ estimation pipeline showing the four key components evaluated in this study: unforced dynamics, maneuver modeling, measurement model and batch configuration (batch length and size).

Sequential filters such as the EKF and UKF are standard techniques for AO state estimation, with recent work combining both filters in an adaptive formulation [22]. However, in-flight experience has shown that even a few measurement outliers can trigger filter divergence [23, 24]. Batch least-squares filtering processes a fixed observation window simultaneously, making it structurally more robust to sporadic outlier corruptions and well suited to serve as a real-time cross-check or re-initialization fallback for sequential filters [24]. Its performance, however, is sensitive to model fidelity and to the number and distribution of observations composing the batch. In fact, Gaias et al. [8] demonstrated that modeling the perturbed environment leads to weakly observable (rather than unobservable) state. On the other hand, calibrated maneuvers, which are the primary additional observability source [25], can be detrimental to filter convergence if not properly propagated. Similarly, a well propagated state is insufficient if we cannot accurately map it to the corresponding bearing angles. Finally, the influence of batch configuration is not well explored in the present literature. Therefore, a rigorous sensitivity analysis of the four modeling blocks illustrated in Fig. 1 is a prerequisite for reliable filter design and deployment. These blocks are: the *unforced dynamics*, which propagates the state under the effects of the perturbed environment, the *maneuver modeling*, which incorporates impulsive or continuous thrust arcs into the propagation; the *measurement model*, which computes azimuth and elevation angles from the propagated state; the *batch configuration*, defined by the number of measurements processed (batch size) and the time span they cover (batch length).

Regarding the propagation model, Sullivan et al. [26] showed that OE-based dynamics models generally outperform Cartesian formulations at increasing inter-satellite separations, and a large variety of equivalent analytical OE-based relative motion models exist in the literature. However, that analysis was restricted to high-altitude low Earth orbit (LEO), did not consider drifting orbits with non-null relative semi-major axis and applied the same osculating-to-mean transformation to all models regardless of fidelity. Although Gaias et al. [27] extended

the ROE analytical framework to include contributions from all zonal harmonics and second-order effects, a structured comparison of recent high-fidelity ROE-based formulations at lower altitudes and across drifting scenarios is still absent from the literature. Regarding maneuver modeling, Grzymisch and Fichter [28] exposed the distinct observability behavior of impulsive and continuous thrust profiles, while Chernick and D'Amico [29] extended the input sensitivity matrix to eccentric orbits. Neither work explicitly characterizes the burn-duration threshold beyond which the impulsive approximation becomes inaccurate, nor how this threshold scales with maneuver magnitude. Regarding the measurement model, Gaias et al. [30] developed a second-order OE-to-rectilinear-Cartesian pipeline demonstrating arcsecond-level accuracy against PRISMA precise orbit determination products. Willis [31] derived an independent second-order mapping, confirming the accuracy benefits of a ROE-based framework over Cartesian formulations. In both works, mapping accuracy was assessed within a combined propagation-mapping pipeline, preventing fully unambiguous error assessment of the mapping stage alone. Regarding batch configuration, the influence of batch length and size on AO estimation accuracy has received little systematic attention. However, the literature provides the building blocks upon which this analysis can be constructed. Wang et al. [32] showed that higher-order dynamic models yield measurable observability gains, while Ardaens [24] established measurement residual analysis as a standard quality indicator, noting that residuals alone cannot detect convergence to an incorrect local minimum and should be complemented by the analysis of the mean relative longitude, a metric whose role as a convergence indicator is further confirmed in Sullivan [3]. Despite this, principled selection guidelines for batch length and size remain absent from the literature.

The main novelty of this work is a systematic, component-by-component sensitivity analysis of the four modeling blocks constituting an AO-based batch LSQ filter, conducted to identify optimal combinations across representative near-circular LEO scenarios. Specific contributions include:

- **Unforced dynamics:** A structured comparison of analytical propagation models extended to include drifting orbits, carried out at lower orbital altitudes where differential drag (DD) is significant, with each model paired to an osculating-to-mean transformation of matching fidelity.
- **Maneuver modeling:** A dedicated assessment of impulsive versus continuous-thrust maneuver modeling across a range of burn durations and magnitudes.
- **Measurement model:** A novel testing strategy that evaluates ROE-to-Cartesian position mapping accuracy independently of the propagation module, using geometrically interpretable scalar metrics swept over inter-satellite range and complemented by an assessment of observability.
- **Batch configuration:** A systematic sensitivity analysis of batch LSQ performance with respect to batch size and length, both in terms of estimation accuracy and observability.
- **Actionable guidelines:** The results are synthesized in simple and immediately applicable guidelines for the preliminary design of relative navigation/estimation frameworks.

The remainder of this paper is organized as follows. Section 2 introduces the theoretical background of propagation, measurement models and LSQ filtering framework. Section 3 details the methodology of the sensitivity analysis, including the definitions of metrics to evaluate the accuracy of propagation, mapping from ROE to measurements and estimation. In Section 4, results are presented and discussed for representative near-circular LEO scenarios and brief actionable guidelines are defined. Section 5 summarizes the key findings and presents concluding remarks.

2. Background theory

This section outlines the mathematical foundations of all components of the filtering pipeline depicted in Fig. 1.

2.1. Propagation model

The state propagated within the estimation pipeline comprises two coupled components: the relative trajectory of the deputy (D) with respect to the chief (C) and the reference orbit of the chief itself. Both must be advanced under the same perturbation assumptions, so that the dynamical model for the relative state and the one for the chief absolute orbit remain mutually consistent. For instance, propagating the relative orbit under Keplerian dynamics requires the same hypothesis for the chief absolute orbit; the same principle applies when higher-fidelity geo-potential models are employed. The sole exception concerns atmospheric drag, whose differential effect is accounted for in the relative dynamics but absent in the propagation of the chief absolute state. In practice, the chief state is typically available at discrete epochs from an onboard GNSS-based navigation solution, which may render its internal propagation unnecessary. However, when one or more components of the chief absolute state are included among the estimated parameters, internal propagation is required to preserve the correlation between estimate and modeled measurements [5].

In this work, the state of the chief (reference) spacecraft is parameterized with the classical Keplerian OE [33], defined as:

$$\alpha_C = (a_C, e_C, i_C, \Omega_C, \omega_C, M_C)^T, \quad (1)$$

where a_C is the semi-major axis, e_C is the eccentricity, i_C is the inclination, Ω_C is the right ascension of the ascending node (RAAN), ω_C is the argument of perigee and M_C is the mean anomaly of the chief spacecraft. For what concerns the representation of the state of the deputy relative to the chief, the quasi-non-singular ROE are used. They are formulated as a (non-linear) combination of the absolute orbital elements of the two spacecraft [3]. By introducing the mean argument of latitude $u_C = \omega_C + M_C$, we can recall the definition of the quasi-non-singular ROE of the deputy relative to chief:

$$a\delta\alpha = \begin{pmatrix} a\delta a \\ a\delta\lambda \\ a\delta e_x \\ a\delta e_y \\ a\delta i_x \\ a\delta i_y \end{pmatrix} = a_C \begin{pmatrix} (a_D - a_C)/a_C \\ (u_D - u_C) + (\Omega_D - \Omega_C) \cos(i_C) \\ e_D \cos(\omega_D) - e_C \cos(\omega_C) \\ e_D \sin(\omega_D) - e_C \sin(\omega_C) \\ i_D - i_C \\ (\Omega_D - \Omega_C) \sin(i_C) \end{pmatrix}. \quad (2)$$

The quantity $a\delta a$ is the relative semi-major axis, $a\delta\lambda$ is the relative mean longitude, $a\delta\mathbf{e} = (a\delta e_x, a\delta e_y)$ is the relative eccentricity vector and $a\delta\mathbf{i} = (a\delta i_x, a\delta i_y)$ is the relative inclination vector. The formulation in Eq. (2) is dimensional [30], meaning that each component has units of length. By inverting Eq. (2), we can also derive the absolute orbital elements of the deputy from the ROE and OE of the chief in an exact way, provided we avoid the singularity at $i_C = 0$ deg.

While the relative state is parameterized via the quasi-non-singular ROE, other viable options exist, as surveyed by Sullivan et al. [26]. Importantly, the selection of a framework based on ROE is a design choice motivated by several desirable properties that make them a particularly suitable formulation in AO-based relative orbit determination / navigation [3,24,34].

In fact, ROE provide a geometrically meaningful description of the deputy motion while retaining favorable analytical properties for propagation and estimation [3,24,34]. The elements directly describe the shape and orientation of the relative trajectory in the Radial-Transversal-Normal (RTN) frame of the chief, indicated by R in this work. This frame is composed by the radial unit vector $\hat{\mathbf{R}}$, pointing from the central body to the chief center of mass (COM), the normal unit vector $\hat{\mathbf{N}}$, which is aligned with the orbital angular momentum vector and the transversal unit vector $\hat{\mathbf{T}}$, which completes the right-handed

orthonormal triad. In addition, the ROE formulation naturally isolates the weakly observable components of the relative state ($a\delta\lambda$), thereby supporting a more effective treatment in angles-only navigation [3,8].

The “FULL” state vector, collecting the information associated with the *mean* relative orbit and the *mean* chief orbit at a generic epoch, is defined as:

$$\mathbf{x}_{\text{FULL}}(t) = [a\delta\alpha(t)^T, \alpha_C(t)^T, a\delta\dot{\alpha}^T]^T. \quad (3)$$

The third component, $a\delta\dot{\alpha}$, gathers the additional parameters required by the adopted propagation model, namely the constant empirical differential drag coefficients. Introduced by Gaias et al. [35] and subsequently employed by Koenig et al. [36] and Gaias and Colombo [37], these coefficients model the averaged secular effect of differential drag on the relative motion. The dynamics models discussed in this work that do not account for differential drag will simply neglect this component during state propagation.

Throughout this work, both the *mean* relative trajectory and the *mean* chief absolute orbit are propagated via analytical models. Compared to numerical integration, analytical propagators offer a substantial reduction in computational cost, which is a critical advantage, especially if onboard implementation is targeted. The general *propagation model* is formulated as a function that takes the full state at an initial epoch t_0 and returns the propagated state at a generic time t . Formally,

$$\mathbf{x}_{\text{FULL}}(t) = \mathbf{f}_{\text{STM}}(t_0, t, \mathbf{x}_{\text{FULL}}(t_0), \{^R\mathbf{u}(t)\}). \quad (4)$$

The set $\{^R\mathbf{u}(t)\}$ denotes the time series of thrust commands executed over the interval $[t_0, t]$, expressed in frame R . The subscript “STM” is used here in a broad sense to denote any analytical propagation scheme even though some models improve the propagation with a state transition tensor (STT) that captures second-order effects.

2.1.1. Unforced dynamics

Regardless of the specific unforced propagation model used in the following analysis, the dynamics are decomposed into independently modeled contributions whose effects on the state evolution are linearly superimposed. For the ROE portion of the full state, this reads:

$$\begin{pmatrix} a\delta\alpha(t) \\ a\delta\dot{\alpha} \end{pmatrix} = \Phi(\tau) \begin{pmatrix} a\delta\alpha(t_0) \\ a\delta\dot{\alpha} \end{pmatrix} + \begin{pmatrix} \frac{1}{2}\Psi(\tau) \otimes a\delta\alpha(t_0) \otimes a\delta\alpha(t_0) \\ \mathbf{0} \end{pmatrix} \quad (5)$$

where t is a generic time instant, t_0 the initial time, $\tau = t - t_0$ the propagation interval, $\Phi(\tau)$ the first-order STM, $\Psi(\tau)$ the second-order STT and ‘ $\otimes$ ’ the dyadic product. The STM is structured as:

$$\Phi(\tau) = \begin{bmatrix} \Phi_{\text{KEP}+J_i}(\tau) & \Phi_{(\text{KEP}+J_i)+\text{DD}}(\tau) \\ \mathbf{0}_{3 \times 6} & \mathbf{I}_{3 \times 3} \end{bmatrix}. \quad (6)$$

The upper-left block captures the Keplerian contribution and the secular effects of the even zonal harmonics up to i th term (J_i). The Keplerian contribution, here referred to as KEP-STM, maintains constant all the ROE aside from the mean relative longitude [31] that varies accordingly to: $a\delta\lambda = -\frac{3}{2}n_C\tau a\delta a$, where n_C is the mean motion of the chief orbit. The upper-right block of Eq. (6) models the one-directional coupling between the Keplerian-plus-zonal dynamics and differential drag, treated here in a simplified manner with a one-way-coupling. The corresponding lower-left block is therefore zero. The lower-right identity matrix reflects the assumption that the differential drag coefficients $a\delta\dot{\alpha}$ remain constant throughout propagation. Unless otherwise stated, a pre-fitted value of these parameters is provided to the filter in order to limit the degrees of freedom of the sensitivity analyses. The STT $\Psi(\tau)$ captures second-order effects arising from Keplerian dynamics and the J_2 perturbation. In general, the models compared in this work will implement only a subset of the blocks in Eqs. (5) and (6); the remaining blocks are set to zero, recovering a simpler propagation structure from the general formulation.

When accounting for the geopotential perturbation, it is convenient to consider the *mean* OE, intended as *secular* components, where short-

Table 1
Summary of analytical relative motion models.

Model	Orbit validity	Geo-potential	Differential drag	Mean ↔ Osculating	Solution form
KEP-STM [31]	$e \approx 0$	No	No	None	STM
GDA-STM [8]	$e \approx 0$	J_2	No	1st order [38]	STM
GA-STM [39]	$0 \leq e < 1$	J_2	No	1st order [38]	STM
KGD-STM [36]	$0 \leq e < 1$	J_2	No	1st order [38]	STM
GCL-STM [38]	$0 \leq e < 1$	J_2, J_4, J_6	No	2nd order [38]	STM, STT
KGD-STM-dd [36]	$0 \leq e < 1$	J_2	Empirical	1st order [38]	STM
GCL-STM-dd [38]	$0 \leq e < 1$	J_2, J_4, J_6	Empirical	2nd order [38]	STM, STT

and long-term oscillating terms have been removed. In this framework, only the Ω_C , ω_C , and M_C secular elements vary over time, as function of the even zonal terms of the potential, up to a given order, retained in the averaging process. Accordingly, the absolute orbit of the chief can be propagated analytically as:

$$\alpha_C(t) = \alpha_C(t_0) + \dot{\alpha}_{C, \text{KEP}+J_i}(t_0) \tau, \quad (7)$$

where $\dot{\alpha}_{C, \text{KEP}+J_i}$ is the rate change of the secular chief OE due to Keplerian dynamics and even zonal harmonics up to the i th term, both known as explicit functions of the constant $\alpha_C(t_0)$ components. Hence, by combining Eqs. (5) and (7) we obtain the generic analytical dynamics model capable of propagating the mean relative orbital elements and mean absolute OE of the chief under a coherent set of accelerations. When introducing also the DD perturbation through density-model-free approaches [35,36], only the time-variation of the ROE due to DD is regarded, since these are specifically designed to avoid any dependence on atmospheric density models in the propagation phase. As such, the effect of the atmospheric drag on the chief orbit has to be retrieved from an independent information source, typically from the GNSS-based absolute navigation filter. For further details on the analytical propagation of the relative motion and on the effects of perturbations, the reader is encouraged to refer to [3,26,36,38].

The set of models considered in this work is listed in Table 1 in order of increasing fidelity. The KEP-STM models first-order Keplerian motion only. The GDA-STM [8] augments it with simplified secular J_2 effects, neglecting its coupling with non-zero $a\delta a$. The GA-STM [39] also captures secular J_2 dynamics using orbital element differences¹ and extends applicability to eccentric orbits. The KGD-STM [36] is also a first-order motion accounting for all secular J_2 effects, though parameterized in ROE and valid for arbitrary eccentricity. Finally, the GCL-STM [38] includes zonal geo-potential terms up to order J_6 and J_2^2 and provides second-order expansions (via STT) for Keplerian and J_2 contributions at arbitrary eccentricity. Within the considered subset of models, GCL-STM is the only one including the STT of Eq. (5).

Even though both KGD-STM-dd and GCL-STM-dd include J_2 and differential drag, they are derived through different approaches. Koenig et al. [36] first linearized the perturbed OE rates and then obtained the STM in closed form, introducing a rotation of the relative eccentricity vector by the chief argument of perigee to render the plant matrix time invariant. This formulation is formally valid for arbitrary eccentricity, although practical limitations arise from the first-order treatment of δe . In contrast, GCL-STM follows the geometric approach of Gim and Alfriend [39]: it first solves the perturbed chief orbit analytically and then expands the relative motion geometrically, which allows higher-order derivatives for whatever eccentricity value. For differential drag, Koenig et al. [36] cast the average OE dd-variation in the same parameters introduced in Gaias et al. [35]. Whereas, GCL-STM-dd adopts straightaway the empirical acceleration approach of [35] to augment its state by the required additional parameters.

¹ When alternative OE-based parameterizations of the relative motion are adopted, the overall framework based on ROE and Keplerian OE remains unchanged. Dedicated interface layers map the internal state representation to and from this standard format, enabling seamless input/output compatibility with Eq. (5).

A critical factor for accurate propagation is the consistent initialization of the propagation, i.e. the computation of the *mean* ROE and chief *mean* OE ($a\delta\alpha(t_0), \alpha_C(t_0)$). All models including geo-potential perturbations therefore require mean to/from osculating transformations, used both at initialization and prior to the measurement model (see Section 2.2). Their accuracy must be consistent with the propagation fidelity. Refs. [8,36,39] adopt first-order J_2 -only transformations, whereas GCL-STM employs a second-order analytical mapping including zonal, sectorial and tesseral terms up to degree and order six, valid for near-circular orbits [27,38]. This consistency principle ensures that a model is not initialized with transformations that account for perturbations beyond its own modeling capability. Applying a transformation that includes higher-order effects to a model that neglects those same effects would introduce inconsistency in the framework.

2.1.2. Maneuver modeling

Proximity operations often require frequent maneuvers. To achieve accurate propagation, it is thus necessary to include their effect on the dynamics, both from the perspective of the relative motion and of the reference spacecraft. By linearizing the non-linear equations of motion about the chief absolute orbit, it is possible to obtain a closed form solution to propagate the ROE-state from a generic time t_i to time t_j [3]:

$$a\delta\alpha(t_j) = \Phi(t_j, t_i) a\delta\alpha(t_i) + \int_{t_i}^{t_j} \Phi(t_j, t) \mathbf{B}(u_C(t)) {}^R \mathbf{u}(t) dt, \quad (8)$$

where ${}^R \mathbf{u}(t)$ is a generic control profile expressed in the RTN frame of the chief and $\mathbf{B}(t)$ is the time-variant input sensitivity matrix that maps the control input into ROE-space. The above equation is thus composed by a zero-input response and a zero-state response. The former has been analyzed in detail in Section 2.1.1. The latter component is obtained by solving the definite integral, which can be done analytically or numerically by means of some discretizations. A particularly simple case is obtained when the maneuvers can be considered impulsive and the instantaneous “jump” in ROE simply reduces to:

$$\Delta(a\delta\alpha) = \mathbf{B}(u_C^{\text{MAN}}) {}^R \Delta \mathbf{v}^{\text{MAN}}, \quad (9)$$

where the ${}^R \Delta \mathbf{v}^{\text{MAN}}$ represents the change in velocity of the maneuvering spacecraft (the chief) expressed in the RTN frame due to an impulsive maneuver executed at the mean latitude u_C^{MAN} . In the case of quasi-circular orbit it is possible to derive a simple form of the input-sensitivity matrix from the Gauss variational equations (GVE) [34,40]:

$$\mathbf{B}(u_C^{\text{MAN}}) = -\frac{1}{n_C} \begin{bmatrix} 0 & 2 & 0 \\ -2 & 0 & 0 \\ \sin u_C^{\text{MAN}} & 2 \cos u_C^{\text{MAN}} & 0 \\ -\cos u_C^{\text{MAN}} & 2 \sin u_C^{\text{MAN}} & 0 \\ 0 & 0 & \cos u_C^{\text{MAN}} \\ 0 & 0 & \sin u_C^{\text{MAN}} \end{bmatrix}. \quad (10)$$

The above matrix depends only on the maneuver execution point along the orbit, via u_C^{MAN} and the orbit energy, via the mean motion of the chief n_C . During formation flying operations, the maneuvers are usually quite small (from mm/s to few cm/s) and in general the chief (which is the one maneuvering) orbit is not affected much by the them. However,

Table 2
Mapping to relative Cartesian State/Position models.

Model	Orbit validity	Mapping order	Orbit curvature	Peculiarity
GDA-MAP-r [8]	$e \approx 0$	First-order linear	No	None
GDA-MAP-c [8]	$e \approx 0$	First-order linear	Yes	Curvilinear-to-rectilinear transformation [39]
GA-MAP-r [39]	$e \approx 0$	First-order linear	No	Approximate solution of Kepler's equation [42]
GA-MAP-c [39]	$e \approx 0$	First-order linear	Yes	Curvilinear-to-rectilinear transformation [39]
GAC-MAP [30]	$e \approx 0$	Second-order	Yes	Curvilinear-to-rectilinear transformation [39]
WD-MAP [31]	$e \approx 0$	Second-order	Yes	Approximate solution of Kepler's equation [42]
GAC-MAP-e [30]	$e \approx 0$	Non-linear exact	Yes	Approximate solution of Kepler's equation [42]

for the sake of consistency, also the mean absolute OE of the chief are updated with the same strategy outlined in this section, but applied to the OE of the chief instead of the ROE. In particular, the input sensitivity matrix to map the ${}^R\Delta\mathbf{v}^{\text{MAN}}$ to the “jump” in α_C is now obtained from Eqs. (9)–(24) of Vallado [33] and is not repeated here for conciseness.

Sometimes, the maneuvers performed during the operations have duration comparable to the chief orbit period, making the assumption leading to Eq. (9) inadequate. We can however still assume that the maneuver has a constant thrust or, at least, that can be divided into constant-thrust segments. With this assumption, it could be possible to follow the work of Grzymisch and Fichter [28] to analytically solve the convolution in Eq. (8). However, Grzymisch and Fichter [28] placed the reference on the non-maneuvering spacecraft, here the deputy. This assumption is necessary to obtain a constant STM and thus solve the integral analytically. This is not convenient in an AO setting since, in general, the deputy state might not be known precisely (with respect to a differential GNSS based estimation) due to intrinsic limitations of the AO framework. Thus, we are compelled to set the reference on the maneuvering spacecraft, whose orbital elements are affected by the maneuver as well, making the STM inside the integral of Eq. (8) time-varying. Consequently, the convolution should be integrated numerically. Alternatively, the quasi-continuous modeling approach introduced in our previous work [41] can be used. Here, the constant continuous maneuver (segment) is approximated via a set of equivalent impulses. A continuous maneuver of duration Δt^{MAN} is subdivided into N^{IMP} time intervals of equal length $\Delta t_{\text{max}}^{\text{MAN}}$, each approximated by an equivalent impulse applied at its midpoint via Eqs. (9)–(10), with free-propagation arcs inserted between consecutive midpoints. This approach, accurately approximates continuous thrust without requiring an analytical solution of the forced dynamics.

2.2. Measurements

The angles-only measurement model describes the relationship between the ROE and chief OE and the corresponding modeled measurement which consist of a pair of angles, the azimuth (α) and elevation (ϵ). The model can be defined as:

$$\mathbf{y} = \begin{pmatrix} \alpha \\ \epsilon \end{pmatrix} = \begin{pmatrix} \arcsin({}^V l_y) \\ \arctan({}^V l_x / {}^V l_z) \end{pmatrix}. \quad (11)$$

Where the quantities ${}^V l_x$, ${}^V l_y$ and ${}^V l_z$ are the component of the line-of-sight (LOS) unit vector ${}^V \hat{\mathbf{l}}$ expressed in the VBS frame V . Azimuth and elevation subtend this vector. The LOS vector points from the COM of the chief spacecraft to the deputy COM. By expressing it in the frame of the VBS we can write:

$${}^V \hat{\mathbf{l}} = \frac{{}^V \delta \mathbf{r}}{\|{}^V \delta \mathbf{r}\|_2}, \quad (12)$$

where ${}^V \delta \mathbf{r}$ is the relative position vector of deputy with respect to the chief, expressed in the VBS frame.

In order to reconcile the measurement model with our OE-based framework, it remains to be discussed the process to obtain the relative state expressed in R from the *mean* ROE and *mean* absolute OE of the

chief. The complete pipeline is expressed by the following chain of non-linear transformations:

$$\begin{aligned} (\alpha_C, a\delta\alpha) &\xrightarrow{\text{Eq. (2)}} (\alpha_C, \alpha_D) \xrightarrow{\text{MEAN} \rightarrow \text{OSC.}} (\tilde{\alpha}_C, \tilde{\alpha}_D) \rightarrow \\ &\xrightarrow{\text{KEP.} \rightarrow \text{CART.}} ({}^I \mathbf{y}_C, {}^I \mathbf{y}_D) \xrightarrow{{}^R \mathbf{R}_I} ({}^R \mathbf{y}_C, {}^R \mathbf{y}_D) \xrightarrow{\delta(\cdot)} {}^R \delta \mathbf{x} \end{aligned} \quad (13)$$

After the *mean* element of the deputy are computed solving Eq. (2) for the deputy OE, the mean-to-osculating transformation converts (α_C, α_D) to their osculating counterparts, denoted by $(\tilde{\cdot})$. The mean-to-osculating transformation must be consistent with the one used to initialize the propagation model, as discussed in Section 2.1.1. Then, the osculating Keplerian elements are transformed to Cartesian states ${}^I \mathbf{y}$ in the ECI frame with a Keplerian-to-Cartesian transformation [43] and ${}^R \mathbf{R}_I$ rotates both the position and velocity into the RTN frame of the chief. This matrix is immediately obtained from the chief OE [33]. Finally, the relative Cartesian state ${}^R \delta \mathbf{x}$ is obtained by differencing $\delta(\cdot)$ the two absolute Cartesian states. Because RTN is non-inertial, the relative velocity includes the usual transport term associated with the frame angular velocity.

The first three components of ${}^R \delta \mathbf{x}$ form the relative position vector ${}^R \delta \mathbf{r}$ that can be rotated to the VBS-frame using the VBS attitude matrix (${}^V \mathbf{R}_R$). Then, the LOS vector is found using Eqs. (12). Consequently, the bearing angles are immediately obtained applying Eq. (11). The non-linear set of operations in this expression represents the exact transformation to obtain ${}^R \delta \mathbf{x}$ and thus the LOS vector. However, this chain of transformations can be quite computationally expensive since it involves the numerical solution of Kepler's equation to compute the true anomaly from the mean anomaly. Consequently, several attempts have been made to derive simplified mappings to obtain the relative position in a more computationally efficient manner. A subset of these models are reported in Table 2. When using these mappings, the exact non-linear transformation procedure expressed by Eq. (13) is modified to explicitly make clear the usage of one of the models in Table 2:

$$(\alpha_C, a\delta\alpha) \xrightarrow{\square\text{-MAP}} {}^R \delta \mathbf{x} \quad (14)$$

Notice that, internally to the mapping/transformation, the mean-to-osculating step is still applied. This is essential since mean elements are a mathematical abstractions, while the osculating orbital elements represent the actual, instantaneous orbit. However, the mappings do not necessarily require the conversion to cartesian elements, as done in Eq. (13). This appears clear in the derivation of GDA-MAP-r, discussed in Gaias et al. [8] and not reported here for brevity.

In this work several mappings found in the literature and listed in Table 2 undergo a comparative analysis. The simplest (but less accurate) of them is the aforementioned direct linear mapping GDA-MAP used in [8]. Two variants of this model are analyzed here: the rectilinear one (“-r”) and the curvilinear one (“-c”). The latter uses a correction of to account for orbit curvature [26,39]. Secondly, the model developed by Gim and Alfriend [39] GA-MAP is considered, again in both the rectilinear and curvilinear variants. Notice that this model is the only one not based on quasi-non-singular ROE but uses differences of the OE of the two spacecraft, making it possible to develop simple interfaces that allow it to seamlessly fit in a framework based on ROE. All the above methods are first order. On the other hand, the GAC-MAP, introduced by Gaias et al. [30] and the WD-MAP developed by Willis [31] are second order approximations of the

Table 3
Ground-truth high fidelity simulator.

Simulator	Configuration
Epoch (UTC)	2020-10-15 14:30:00
Integrator	Dormand-Prince 5th order
Step	Fixed
Geo-potential	GGM01S (up to 120×120)
Atmospheric Drag	Cannonball model
Solar Radiation Pressure	Flat plate, conical Earth shadow
Atmospheric Density	NRLMSISE-00
Third-body gravity	Luni-solar (point masses)
Luni-solar Ephemerides	Analytic
Relativistic Effects	Approximate general relativity
Maneuvers	Treated as continuous

transformation in Eq. (13). Lastly, this work considers also a non-linear transformation, developed by Gaias et al. [30] that simplifies the exact one of Eq. (13) by using an approximate solution of the Kepler's equation [42] and by computing only ${}^R\delta\mathbf{r}$ (avoiding the relative velocity). The approximate evaluation of the true anomaly using the simplified approach of [30,42] is based on an expansion in the orbital eccentricity, which restricts its validity to small values of e . In addition to GAC-MAP-e, the implementations of GA-MAP, GAC-MAP and WD-MAP also adopt this approximation, hence their limited applicability with respect to eccentricity.

2.3. Batch LSQ filtering framework

The relative orbit determination is cast as a batch least-squares problem, building on the ROE-based formulation of Gaias et al. [8] and Koenig and D'Amico [5]. The latter proved that augmenting the estimated state with the chief semi-major axis allowed the improvement of the initial guess on a_C without incurring any noticeable loss in the observability of $a\delta\alpha$. This improves the estimation of $a\delta a$ and $a\delta\lambda$, thus justifying an augmented (estimated) state formulation. This framework extends our previous work [41], where only the ROE were estimated.

Consider the full state introduced in Eq. (3), which collects the information of the *mean* relative orbit, the *mean* orbit of the chief at a generic time instant and is augmented by the constant DD empirical parameters. Without loss of generality we assume to estimate the state at the initial time t_0 since we often have some *a priori* knowledge of the full state at this instant. This could result from a process of initial relative orbit determination or from other sources, such as NORAD two-line-elements. We can then partition the full state as:

$$\mathbf{x}_{\text{FULL}}(t_0) = \begin{pmatrix} \mathbf{x}_{\text{EST}}(t_0) \\ \mathbf{x}_{\text{PRIOR}}(t_0) \end{pmatrix}. \quad (15)$$

The first part represents the portion of the state that will be estimated by the LSQ algorithm. The second consists of parameters that are known *a priori* but will not be updated by the filter. The partition is done to highlight that only some components of the full state are updated by the filtering process. In practice, we will either estimate only the relative orbital elements at t_0 , meaning $\mathbf{x}_{\text{EST}}(t_0) = a\delta\alpha(t_0)$ or the ROE augmented by the chief semi-major axis, meaning $\mathbf{x}_{\text{EST}}(t_0) = [a\delta\alpha(t_0)^T, a_C(t_0)]^T$. The estimation of additional parameters, such as the empirical DD parameters [36,44] used by the aforementioned models, is certainly possible. However, it requires first a proper observability analysis which is beyond the scope of this work. Inspection of the bottom-right block in (6) shows that these parameters are treated as constants.

The modeled bearing angles $\mathbf{y}(t_i)$ at each measurement epoch t_i are obtained by propagating the mean ROE and chief OE from t_0 via one of the models in Table 1, accounting for any maneuvers as described in Section 2.1.2 and subsequently computing the measurement using one the mappings in Table 2. When $\alpha_C(t_0)$ is not part of the estimated state, the chief absolute OE at each epoch may be supplied

directly from a exogenous solution rather than propagated internally. Such information might be provided by the absolute navigation (A-NAV) filter of the chief. This strategy must not be applied to any component of $\alpha_C(t_0)$ that is being estimated, as it would break the correlation between modeled measurements and that state component, rendering it unobservable. We can formally summarize the chain of operations needed to obtain the bearing angles by defining the following propagated-*modeled*-measurement function:

$$\mathbf{y}(t) = \mathbf{h}_{\text{MAP}}(\mathbf{f}_{\text{STM}}(t_0, t, \mathbf{x}_{\text{FULL}}(t_0), \{{}^R\mathbf{u}(t)\}), {}^V\mathbf{R}_I(t)), \quad (16)$$

where ${}^V\mathbf{R}_I(t)$ is the VBS attitude with respect to the inertial frame. Stacking Eq. (16) over the full batch of $|B|$ measurements yields:

$$\mathbf{Y}(\mathbf{t}^{\text{MEAS}}) = \begin{pmatrix} \mathbf{h}_{\text{MAP}}(\mathbf{f}_{\text{STM}}(t_0, t_1, \mathbf{x}_{\text{FULL}}(t_0), \{{}^R\mathbf{u}(t_1)\}), {}^V\mathbf{R}_I(t_1)) \\ \vdots \\ \mathbf{h}_{\text{MAP}}(\mathbf{f}_{\text{STM}}(t_0, t_i, \mathbf{x}_{\text{FULL}}(t_0), \{{}^R\mathbf{u}(t_i)\}), {}^V\mathbf{R}_I(t_i)) \\ \vdots \\ \mathbf{h}_{\text{MAP}}(\mathbf{f}_{\text{STM}}(t_0, t_{|B|}, \mathbf{x}_{\text{FULL}}(t_0), \{{}^R\mathbf{u}(t_{|B|})\}), {}^V\mathbf{R}_I(t_{|B|})) \end{pmatrix}, \quad (17)$$

where $\mathbf{t}^{\text{MEAS}}$ collects the $|B|$ measurement epochs. Defining $\mathbf{Z}(\mathbf{t}^{\text{MEAS}})$ as the analogously stacked vector of *observed* measurements, the batch residual vector reads:

$$\rho(\mathbf{t}^{\text{MEAS}}) = \mathbf{Y}(\mathbf{t}^{\text{MEAS}}) - \mathbf{Z}(\mathbf{t}^{\text{MEAS}}). \quad (18)$$

The batch-LSQ filter seeks the estimate $\hat{\mathbf{x}}_{\text{EST}}^{\text{LSQ}}(t_0)$ minimizing the weighted residual norm plus an *a priori* penalty [8,43]. The solution is obtained with the following iterative scheme:

$$\begin{aligned} \mathbf{P}_0^{[j+1]} &= \left(\mathbf{A}^{[0]} + \mathbf{H}^{[j]T} \mathbf{W}^{-1} \mathbf{H}^{[j]} \right)^{-1} \\ \hat{\mathbf{x}}_{\text{EST}}^{[j+1]}(t_0) &= \hat{\mathbf{x}}_{\text{EST}}^{[j]}(t_0) + \mathbf{P}_0^{[j+1]} \left[\mathbf{A}^{[0]} \left(\hat{\mathbf{x}}_{\text{EST}}^{[j]}(t_0) - \hat{\mathbf{x}}_{\text{EST}}^{[0]}(t_0) \right) \right. \\ &\quad \left. + \mathbf{H}^{[j]T} \mathbf{W}_z^{-1} \rho^{[j]} \right], \end{aligned} \quad (19)$$

where $\mathbf{W}_z = \text{blkdiag}(\mathbf{W}_1, \dots, \mathbf{W}_{|B|})$ with $\mathbf{W}_i = \text{diag}(\sigma_a^{-2}, \sigma_e^{-2})$ and $\mathbf{A}^{[0]} = (\mathbf{P}_{\text{EST}}^{[0]})^{-1}$ is the *a priori* information matrix. The matrix $\mathbf{H}(\mathbf{x}_{\text{FULL}}(t_0))$ is the Jacobian of the propagated measurements [5] and is defined as:

$$\mathbf{H}(\mathbf{x}_{\text{FULL}}(t_0)) = \frac{\partial \mathbf{Y}(t_0, \mathbf{t}^{\text{MEAS}}, \mathbf{x}_{\text{FULL}}(t_0))}{\partial \mathbf{x}_{\text{EST}}(t_0)} \bigg|_{\hat{\mathbf{x}}_{\text{EST}}(t_0)}. \quad (20)$$

The matrix is computed numerically via central differences [5], since the non-linear measurement models of Table 2 preclude closed-form analytical expressions. Importantly, only the estimated portion of the state is perturbed in the computation. The recursion of Eq. (19) starts with $j = 0$ and continues until convergence, yielding $\hat{\mathbf{x}}_{\text{EST}}^{\text{LSQ}}(t_0)$ and $\mathbf{P}_0^{\text{LSQ}}$.

3. Sensitivity analyses

All analyses are conducted against a simulated ground-truth generated by a high-fidelity numerical simulator based on Montenbruck and Gill [43], whose settings are reported in Table 3. The simulation epoch fixes both the solar and geomagnetic indices, internally retrieved by the NRLMSISE-00 model, and the analytical luni-solar ephemerides. The selected epoch corresponds to a geomagnetically quiet period (atmospheric density $\approx 1 \times 10^{-13} \text{ kg/m}^3$ at the selected altitude). This deliberate choice balances the drag prominence without masking the effects of competing perturbations in the sensitivity analyses. The simulator integrates the Cartesian equations of motion of chief and deputy under a user-selectable perturbation set. Spacecraft physical parameters, summarized in Table 4, are inspired from the AVANTI experiment [45] and are chosen to accentuate differential drag effects in drag-active scenarios.

Simulator outputs are partitioned into *true* quantities, obtained directly from numerical integration without corruption and *measured* quantities, which are a noise-corrupted subset of the true outputs generated by a VBS sensor block. The VBS block produces time-tagged

Table 4
Spacecraft physical parameters.

Parameter	Chief	Deputy
Mass (kg)	130	1
Area (m ²)	0.5984	0.0100
Drag Coefficient (–)	2.5	2.2
Ballistic Coefficient (kg/m ²)	0.0115	0.0220
Reflectivity Coefficient (–)	1.3	1.2

Table 5
Main outputs of simulator — True (time-series format).

Output	Description
$\mathbf{t}^{\text{TRUE}}$	Simulation time-instants vector
$\{\alpha_C^{\text{TRUE}}(t)\}$	Time series of chief OE
$\{a\delta\alpha^{\text{TRUE}}(t)\}$	Time series of ROE
$\{^V\mathbf{R}^{\text{TRUE}}(t)\}$	Time series of VBS inertial attitude
$\{^R\delta\mathbf{x}^{\text{TRUE}}(t)\}$	Time series of relative Cartesian state in R
$\{\mathbf{z}^{\text{TRUE}}(t)\}$	Time series of ideal (noise-free) measurements

Table 6
Main outputs of simulator — Measured (time-series format).

Output	Description
$\mathbf{t}^{\text{MEAS}}$	Measurement time-instants vector
$\{^V\mathbf{R}^{\text{MEAS}}(t)\}$	Time series of noisy VBS inertial attitude
$\{\mathbf{z}^{\text{MEAS}}(t)\}$	Time series of noisy measurements

azimuth-elevation measurements at a user-defined rate by corrupting true values with Gaussian white noise of standard deviations σ_α and σ_ϵ . The block also outputs the noisy VBS inertial attitude, corrupted by off-axis (σ_{off}) and bore-sight (σ_{bore}) noise components. Illumination effects (deputy eclipse, camera blinding) are disabled in all analyses. This is relevant only for the batch configuration study, where it comes necessary to select arbitrary batch configurations. The simulator outputs are listed in [Tables 5 and 6](#).

In order to determine the optimal configuration of the estimation pipeline components illustrated in [Fig. 1](#), this work adopts a sequential strategy, in which each sub-block is assessed independently. The underlying hypothesis is that a filter composed of the most accurate dynamics, the most effective maneuver modeling and the most precise mapping yields the highest estimation accuracy and filter robustness. The remainder of the section explains the performance assessment strategy and the metrics employed to quantify performance of each sub-component. The terminology and scheme used consistently throughout this and subsequent sections is defined in [Table 7](#).

3.1. Accuracy of maneuver-free propagation

This analysis assesses the accuracy of relevant analytic models reported in [Table 1](#) with the aim of studying their sensitivity to the inter-satellite range and the relative semi-major axis. The aimed outcome is the identification of the most accurate.

Scenario: Ground-truth simulations are performed under the fully perturbed scenario (see [Table 3](#)). The geo-potential is set to 30th-degree-and-order non-homogeneous Earth gravity field. No maneuvers are executed. Each simulation runs for 24 h at a fixed time step of 10 s.

Families: Two families can be identified in this analysis. Both share a single single chief orbit, LEO1, which is reported in [Table 8](#). Its altitude of 500km is deliberately chosen to produce a pronounced DD effect on the relative orbit.

The two families are differentiated by their relative orbits: RO1 and RO2, respectively. These are defined in [Table 9](#).

Family RO1 produces bounded and centered relative motion, where the only non zero components are the $a\delta e_x$ and $a\delta i_x$. This configuration maximizes the influence of J_2 perturbation, making it well suited to

Table 7
Terminology and scheme used throughout the sensitivity analyses.

Analysis	Study of one pipeline component: unforced dynamics, maneuver modeling, measurement mapping and batch configuration.
Scenario	Defined by the perturbative forces active in the ground-truth simulator and the presence or absence of maneuvers.
Family	Identified by a family of initial conditions (ICs) for the chief mean orbit (LEO#) and the mean relative orbit (RO#).
Case	Identified by specific ICs expressed via the parameters of Eqs. (1) and (2). Initialization from mean ROE and mean OEs follows Sullivan et al. [26]. Maneuver timing, duration and intensities further identify a case when maneuvers are present.
Delta-case	Arises when sensitivity variables do not parametrize ICs, or when the noise-corrupted outputs of Table 6 are used; one case may then generate multiple delta-cases.
Sensitivity variable	Quantities with respect to which an analysis is conducted; defines cases when it parametrizes ICs, delta-cases otherwise.
Evaluation metric	Scalar performance indicator computed per case or delta-case; defined individually for each analysis in the sections that follow.
Metric computation	Procedure for computing evaluation metrics for a generic case or delta-case.

Table 8
Initial mean absolute orbital elements of the chief (LEO1).

a (km)	e (–)	i (deg)	Ω (deg)	ω (deg)	M (deg)
6878	0.0001	98.2	30	0	0

discriminating models that capture this effect from those that do not. Family RO2 introduces instead a non-zero semi-major axis difference ($a\delta a \neq 0$), causing the deputy and chief orbits to drift apart, a case typical of rendezvous missions. This scenario specifically stresses propagation models that cannot account for the coupled effect of J_2 and the energy-induced drift [8].

Cases: Each combination of LEO1 with RO1(c_1) and RO2(c_2) give rise to a specific case which is simulated under the selected scenario. To obtain the different initial conditions, parameter c_1 and c_2 are each sampled at 10 times, yielding 10 cases per family and 20 in total. The sampling intervals for these parameters are: $c_1 \in [1, 10^5]$ m, $c_2 \in [1, 10^3]$ m.

Sensitivity Variables: The first sensitivity variable, adopted for the cases involving RO1 family is the (dimensionless) maximum inter-satellite distance achieved during propagation:

$$\delta r_{\max}/a_C = \frac{\max_{t_i} \|\delta \mathbf{r}^{\text{TRUE}}(t_i)\|_2}{a_C^{\text{TRUE}}(t_0)} \quad (21)$$

The sensitivity variable adopted for the RO2-cases is the relative semi-major axis $a\delta a^{\text{TRUE}}(t_0)$ which specifies c_2 .

Evaluation metric: Propagation accuracy is quantified by a single scalar metric ν , defined as:

$$\nu = \max_{t_i} \|\mathbf{W} (\delta \mathbf{x}^{\text{STM}}(t_i) - \delta \mathbf{x}^{\text{TRUE}}(t_i))\|_2 \quad \text{for } t_i \in [t_0, t_F], \quad (22)$$

$$\mathbf{W} = \text{diag}(1, 1, 1, n_C^{-1}, n_C^{-1}, n_C^{-1}),$$

where $\delta \mathbf{x}^{\text{TRUE}}(t_i)$ is one of the mentioned “true” outputs (see [Table 5](#)) and the weighting matrix $\mathbf{W}$ is used to normalize the velocity components by the chief mean motion n_C , rendering all terms dimensionally homogeneous. Taking the maximum over all simulation time instants yields a single scalar per case. The “STM” Cartesian relative state results from the analytical propagation of $a\delta \alpha(t_0)$ and $\alpha_C(t_0)$ via one of the

Table 9

Initial mean ROE families for maneuver-free propagation tests. All values in (m).

	$a\delta a$	$a\delta\lambda$	$a\delta e_x$	$a\delta e_y$	$a\delta i_x$	$a\delta i_y$
RO1	0	0	c_1	0	c_1	0
RO2	c_2	0	500	0	500	0

Table 10

Initial mean ROE for maneuver propagation test. All values in (m).

	$a\delta a$	$a\delta\lambda$	$a\delta e_x$	$a\delta e_y$	$a\delta i_x$	$a\delta i_y$
RO3	0	−30000	500	0	−500	0

models in Table 1. This last point is different from Sullivan et al. [26], where the authors used the ground-truth $\alpha_C^{\text{TRUE}}(t_i)$.

Metric computation: The initial mean ROE ($a\delta\alpha(t_0)$) and the initial mean chief OE ($\alpha_C(t_0)$), are propagated analytically to each time instant t_i using the model under evaluation. The initial *mean* quantities are obtained from the “true” *osculating* ones in Table 5 using the mean-to-osculating transformation associated to each model. The models accounting for DD use a $a\delta\alpha$ obtained by pre-fitting the ground-truth. This is considered fair since providing a corrupted value would hide the benefit of using these models properly. Moreover, utilizing a fitted value is far from knowing the actual *instantaneous* value of $a\delta\alpha$ which oscillates considerably during simulation. The resulting propagated quantities ($a\delta\alpha^{\text{STM}}(t_i)$, $\alpha_C^{\text{STM}}(t_i)$) are mapped to the Cartesian relative state $\delta\mathbf{x}^{\text{STM}}(t_i)$ via the exact non-linear transformation of Eq. (13). The normalized error vector $v_i = \|\mathbf{W}(\delta\mathbf{x}^{\text{STM}} - \delta\mathbf{x}^{\text{TRUE}})\|_2$ is computed at each t_i . Then v is recorded as the maximum across all time instants. This procedure is repeated for each model in Table 1 and every case, enabling direct cross-model comparison across the full range of initial conditions.

3.2. Maneuver propagation accuracy

This analysis assesses how the inclusion of a maneuver affects propagation accuracy when the impulsive and quasi-continuous propagation models described in Section 2.1.2 are used. In particular, the sensitivity of the propagation accuracy to burn intensity and duration is studied with the specific objective of identifying the burn-duration threshold beyond which the impulsive assumption becomes inadequate.

Scenario: Ground-truth simulations are performed under Keplerian, unperturbed motion, obtained by disabling all perturbations. Under purely Keplerian dynamics, the KEP-STM analytical model reproduces the ground-truth with negligible error, thereby eliminating initialization errors from the osculating-to-mean transformation and accumulation of errors along maneuver-free arcs. A single maneuver is present in every simulation, implemented as a continuous thrust. Each simulation runs for 24 h with a maximum (fixed) time-step of 1% the duration of the maneuver for maximum fidelity.

Family: The chief orbit is again defined by LEO1, reported in Table 8. A single relative orbit family, RO3, is considered and reported in Table 10. This RO introduces a large along-track separation consistent with a generic starting point for a mid-range rendezvous scenario [8].

Sensitivity variables: The first sensitivity variable is the maneuver magnitude $\|\Delta\mathbf{v}\|_2^{\text{MAN}}$, which measures the intensity of the maneuver. The second one is the burn duration Δt^{MAN} , which determines how $\|\Delta\mathbf{v}\|_2^{\text{MAN}}$ is distributed.

Cases: All the cases are characterized by the unique relative orbit RO3. However, the selected sensitivity variables define several cases. The two

Table 11

Initial mean ROE for measurement accuracy test. All values in (m).

	$a\delta a$	$a\delta\lambda$	$a\delta e_x$	$a\delta e_y$	$a\delta i_x$	$a\delta i_y$
RO4	0	c_4	$f(c_4)$	0	$-f(c_4)$	0

are sampled as follows:

$$\Delta t^{\text{MAN}} \in \left[\begin{array}{ccccc} 1, & 3, & 7, & 20, & 50, \\ 120, & 320, & 830, & 2170, & 5680 \end{array} \right] \text{ s}$$

$$\|\Delta\mathbf{v}\|_2^{\text{MAN}} \in [0.0001, 0.001, 0.01, 1] \text{ m/s}$$

The burn duration goes from 1 s to 5680 s, which corresponds to approximately one orbital period of the chief at t_0 for LEO1. The combinations of the two variables defines a 10×4 grid of cases. In each of them, the maneuver starts at the very beginning of the simulation ($t_0 = 0$ s) and its duration and intensity are one of the mentioned combinations. In all of them, the maneuver is oriented in a generic direction in the RTN:

$${}^R d_{\Delta\mathbf{v}} = \left[1/\sqrt{3}, -1/\sqrt{3}, 1/\sqrt{3} \right]^\top.$$

The maneuver is kept constant throughout its duration. Thus the chief spacecraft is subject to the following force (applied at its COM) during the burn time:

$${}^R \mathbf{F} = m_C \frac{\|\Delta\mathbf{v}\|_2^{\text{MAN}}}{\Delta t^{\text{MAN}}} {}^R d_{\Delta\mathbf{v}},$$

where m_C is the mass of the chief spacecraft.

Evaluation metric: Two scalar metric are used to fully characterize the accuracy of maneuver modeling. The first is the propagation accuracy v defined in Eq. (22) (see Section 3.1). The second is a simple variation of it that takes value of weighted error norm at t_F instead of the maximum. We refer to this quantity as v_{end} . This allows us to study the long-term effects of maneuver models.

Metric computation: The same computational procedure described in Section 3.1 is adopted. For each delta-case, the initial mean ROE of RO3 is propagated with the KEP-STM model under the two maneuver representations introduced in Section 2.1.2 (single impulsive and quasi-continuous). The resulting $\delta\mathbf{x}^{\text{STM}}(t_i)$ is compared against the simulator output and v is extracted as before. The second metric is obtained by considering only the final time instant. The process is repeated for all 40 cases with both impulsive and quasi-continuous models, yielding two (v, v_{end}) pairs per case, enabling comparison as a function of $\|\Delta\mathbf{v}\|_2^{\text{MAN}}$ and Δt^{MAN} . The equivalent impulses computed by for the quasi-continuous model uses a discretization step of $\Delta t_{\text{max}}^{\text{MAN}} = 0.01$ s to closely reproduce the continuous thrust.

3.3. Measurement accuracy

This analysis assesses the accuracy of the models listed in Table 2 in converting the ROE, combined with the OE of the chief, into the relative position vector. The objective is to quantify the error relative to the exact non-linear transformation of Eq. (13) and to evaluate whether any of the models confers an observability advantage over the others.

Scenario: Ground-truth simulations are performed under purely Keplerian dynamics and no maneuver is executed. In this way we isolate the mapping accuracy by excluding the influence of the mean-to-osculating transformation and of propagation errors. Each simulation runs for 24 h at a fixed integration time step of 10 s.

Family: The chief orbit is defined by LEO1, reported in Table 8. A single relative orbit family RO4 is considered, whose initial mean ROE are reported in Table 11.

The function $f(c_4)$, which shapes the periodic in-plane and out-of-plane motion, takes the following linear form:

$$f(c_4) = 10 + \frac{(1000 - 10)(c_4 - 10)}{(10^5 - 10)} \text{ m.}$$

The shape of RO4 is selected to obtain relative trajectories representative of successive stages of a passively-safe rendezvous approach/drift apart operation.

Cases: Ten cases are obtained by sampling the free parameter c_4 over the range $c_4 \in [10, 10^5]$ m. The eccentricity and inclination vector magnitudes consequently range from 10 m in the first case to 1 km in the last.

Sensitivity variables: The mapping accuracy is expected to degrade with increasing inter-satellite separation [31]. The sensitivity variable is therefore the mean relative longitude $a\delta\lambda(t_0)$, which serves as an analogue for inter-satellite range.

Evaluation metrics: Two complementary scalar metrics are defined to fully characterize mapping accuracy. The first quantifies the angular error between the mapped and true directions:

$$\theta = \max_{t_i} \left(\arccos \left(\frac{R_{\delta\mathbf{r}}^{\text{MAP}}(t_i) \cdot R_{\delta\mathbf{r}}^{\text{TRUE}}(t_i)}{\|R_{\delta\mathbf{r}}^{\text{MAP}}(t_i)\| \cdot \|R_{\delta\mathbf{r}}^{\text{TRUE}}(t_i)\|_2} \right) \right). \quad (23)$$

The above defines the maximum angular deviation across all time instants between the ground-truth relative position vector and the mapped one (superscript “MAP”). Depending on the mapping, $R_{\delta\mathbf{r}}^{\text{MAP}}$ is obtained either directly [8] or extracted from the first three components of $R_{\delta\mathbf{x}}^{\text{MAP}}$ [30]. Since a mapping may closely reproduce the LOS direction while introducing a significant magnitude error, the metric θ alone is insufficient. A second scalar metric is therefore introduced:

$$\gamma = \max_{t_i} (\|\delta\mathbf{r}_{\parallel}^{\text{MAP}}(t_i) - \delta\mathbf{r}^{\text{TRUE}}(t_i)\|_2), \quad (24)$$

$$\delta\mathbf{r}_{\parallel}^{\text{MAP}}(t_i) = \left(\frac{R_{\delta\mathbf{r}}^{\text{MAP}}(t_i) \cdot R_{\delta\mathbf{r}}^{\text{TRUE}}(t_i)}{R_{\delta\mathbf{r}}^{\text{TRUE}}(t_i) \cdot R_{\delta\mathbf{r}}^{\text{TRUE}}(t_i)} \right) \delta\mathbf{r}^{\text{TRUE}}(t_i).$$

This measures the maximum magnitude error in the direction of the true LOS, obtained by projecting the mapped relative position onto the true one and subtracting the true value. Together, θ and γ provide a complete characterization of the accuracy.

In addition to accuracy, the analysis evaluates whether the choice of mapping model affects the observability of the relative orbit estimation problem. This is quantified through the observability index (OI) and the condition number (CN) of the observability Gramian (OG), defined as in [32]:

$$OG = \mathbf{H}^T \mathbf{H}, \quad (25)$$

where $\mathbf{H}$ is the Jacobian of the stacked measurements model given in Eq. (20). The condition number quantifies the sensitivity of the observability to the inter-satellite range across models, while the observability index measures the sensitivity to measurement noise.

Metrics computation: To compute θ and γ , at each time instant t_i , the true mean ROE $a\delta\alpha^{\text{TRUE}}(t_i)$ and the true mean chief orbital elements $\alpha_C^{\text{TRUE}}(t_i)$ are retrieved from the ground-truth dataset of Table 5. The mean-to-osculating transformation is not applied since under Keplerian motion, mean and osculating elements coincide. The mentioned quantities are then passed to each mapping model under evaluation to obtain the mapped relative position vector $R_{\delta\mathbf{r}}^{\text{MAP}}(t_i)$. The ground-truth relative position vector $R_{\delta\mathbf{r}}^{\text{TRUE}}(t_i)$ is extracted from the first three components of the state. Then, the instantaneous angular error is computed as the $\arccos(\cdot)$ of the normalized dot product of $R_{\delta\mathbf{r}}^{\text{MAP}}(t_i)$ and $R_{\delta\mathbf{r}}^{\text{TRUE}}(t_i)$ and the projection $\delta\mathbf{r}_{\parallel}^{\text{MAP}}(t_i)$ is obtained according to Eq. (24). Then, θ and γ are recorded as the maximum of these values across all time instants in $[t_0, t_F]$.

To compute the CN and OI we first build the observability Gramian. A subset of 1000 time instants is uniformly extracted from the true dataset Table 5, retaining the corresponding time series $\{\alpha_C^{\text{TRUE}}(t)\}$




Fig. 2. Visualization of four batches with varying lengths Δt_B and sizes $|B|$. Each data point is shown as a blue dot along the relative orbit. The full relative trajectory is depicted by a black solid line.

and $\{^V \mathbf{R}_I^{\text{TRUE}}(t)\}$. The partial derivatives of the stacked propagated-measurement model of Eq. (17) with respect to the initial ROE state are then computed, assembling the matrix $\mathbf{H}$, from which OG is computed via its definition (25). Finally, OI and CN are obtained directly from OG .

These procedures are repeated for each mapping model in Table 2 and each case to compare the models across the full range of initial conditions.

3.4. Optimal batch length and size

The focus of this analysis is to study the sensitivity of the estimation error to two fundamental parameters of a batch filter:

- **Batch Length Δt_B :** the elapsed time between the first and last observations in the batch. This interval determines the temporal span over which the measurements are acquired.
- **Batch Size $|B|$:** the number of data-points grouped in the batch. Each data-point includes the time instant, a pair of azimuth-elevation angles, the absolute orbit and attitude of the chief spacecraft and the time series of the maneuvers. We assume strict synchrony across all data-points (excluding the maneuvers), so every quantity in a data-point refers to the same time stamp.

A visualization of these parameters is presented in Fig. 2 where the same relative orbit is sampled using four different combinations of Δt_B and $|B|$ to generate batches of data-points. The batch length relates to propagation accuracy [24]. In fact, longer propagation times are expected to result in greater propagation errors due to approximations and unmodeled dynamics, which progressively degrade the accuracy of the estimation process. On the other hand, an excessively small Δt_B might not contain enough information about the relative motion to allow the estimation process. This problems might verify when the illumination is such that measurements can be taken only during a small portion of the chief orbit. The batch size directly relates to the computational feasibility of the procedure. Recalling Eq.s (20)–(19), it is evident that the computational time scales with size of the batch size (each element adds a block row to matrix $\mathbf{H}$). Moreover, a very large number of tightly grouped measurements might lead to ill-conditioning of the OG , as its rows would become close to being linearly dependent.

As secondary objective, we aim at understanding the influence of augmenting the estimated state (normally composed by the ROE in ROD) with the chief semi-major axis a_C to complement the analysis of Koenig and D’Amico [5].

Scenario: Ground truth simulations are performed enabling all perturbations in Table 3. No maneuver is executed in this analysis to avoid

Table 12

Initial mean ROE for batch configuration test. All values in (m).

	$a\delta a$	$a\delta\lambda$	$a\delta e_x$	$a\delta e_y$	$a\delta i_x$	$a\delta i_y$
RO5	−9	−30180	−79	−395	46	558

Table 13

Standard deviation of bearing angles, attitude noise and errors.

Parameter	Standard deviation	Value
Azimuth	σ_a	40 arcsec
Elevation	σ_e	40 arcsec
Attitude (boresight)	$\sigma_{\text{boresight}}$	40 arcsec
Attitude (off-axis)	σ_{off}	6 arcsec

artificially improving observability. Each simulation runs for 24 h with fixed time step of 10 s. The simulation duration implicitly imposes a limit on the batch length. However, this does not introduce limitations on the analysis as will be clear from Section 4.4.

Family: The same relative orbit used in Gaias et al. [8] to test the first version of the batch LSQ filter developed for ARGON experiment is employed in this work as well. The relative orbit is reported in Table 12 where we can notice that RO5 is mid-range and characterized by relatively small values of the x components of the relative eccentricity and inclination vector, which are oriented in an almost anti-parallel configuration.

This is typically done to obtain a stable passively-safe orbit. In fact, it combines limited J_2 effect (small $a\delta e_x$ and $a\delta i_x$), with the minimum distance in the RN-plane imposed by the anti-parallel configuration, as explained in D'Amico and Montenbruck [46]. The limited effect for J_2 is positive from the perspective of mission design (prolongs passive safety and reduces fuel consumption for formation keeping), but limiting for navigation performance. In fact, modeled perturbations have the same observability-improving effects of (small) calibrated maneuvers [8]. Additionally, the selected RO, has a moderate approaching drift of almost 90 m per chief orbit. Differently from Gaias et al. [8], we substitute the high-altitude chief orbit of PRISMA with the 500 km orbit reported in Table 8 to increase the effect of the differential aerodynamic drag. In conclusion, the family identified by LEO1 and RO5 is realistic as it balances observability with mission-design constraints.

Sensitivity variables: The primary sensitivity variables are the batch duration, Δt_B and the batch size, $|B|$. In addition, the analysis considers whether $a_C(t_0)$ is included in the parameter estimation vector, as well as the impact of measurement noise.

Delta-cases: All delta cases share common IC and the ground-truth is uniquely determined by LEO1 and RO5. Eclipses and camera blinding are turned off and the VBS block runs at same rate of the simulation. In this way all the time-series in Tables 5 and 6 have the same number of elements and we can batch the data arbitrarily as function of Δt_B and $|B|$. For convenience, we split the delta-cases in two categories: *ideal* (ID) and *noisy* (NY). In the first, a subset of the true simulation output of Table 5 is directly used to generate the batches of data for the LSQ filter. In the latter, the batch of data-points used for relative orbit determination is formed by synthetic bearing angles and attitude of Table 6. These are corrupted with Gaussian white noise whose statistics are summarized in Table 13. We can split the delta-cases also by the number of parameters estimated by the LSQ filter $\mathbf{x}_{\text{EST}}(t_0)$; thus by the optional augmentation of the ROE state by the chief semi-major axis $a_C(t_0)$.

To finally define the delta-cases we construct a 10×10 grid of values of the remaining two sensitivity variables:

$$\Delta t_B = \begin{bmatrix} 0.1, & 0.5, & 1, & 2, & 3, \\ 4, & 5, & 6, & 7.62, & 15.22 \end{bmatrix} T_C \quad \text{s} \quad (26)$$

$$|B| = \begin{bmatrix} 10 & 65 & 120 & 175 & 230 \\ 280 & 335 & 390 & 445 & 500 \end{bmatrix} \quad (-)$$

Table 14

Standard deviation of ROE initialization error.

$\sigma_{a\delta a}$	$\sigma_{a\delta\lambda}$	$\sigma_{a\delta e_x}$	$\sigma_{a\delta e_y}$	$\sigma_{a\delta i_x}$	$\sigma_{a\delta i_y}$
20 m	3000 m	100 m	100 m	100 m	100 m

The batch lengths are conveniently expressed in terms of the chief orbital period T_C . Each one of the grid entries defines the length and size of the batch of data-points used by the LSQ. The batch itself is then constructed by sampling uniformly the time series in Table 5 (for the ideal conditions) and Table 6 (for the noisy conditions). The number of samples is dictated by $|B|$, while Δt_B imposes the limits in the time-span along which the samples can be extracted. Again, Fig. 2 helps visualizing this concept by showing four extreme examples.

Once built the data batch, prior to the actual estimation task, one has to regard the filter initialization step. This is done by providing to the filter an initial guess for $\mathbf{x}_{\text{EST}}^{[0]}(t_0)$ and the associated covariance $\mathbf{P}_{\text{EST}}^{[0]}(t_0)$ and the remaining components of the full state $\mathbf{x}_{\text{PRIOR}}(t_0)$. These always include the constant empirical differential-drag parameters $a\delta\alpha$ [35,36] and the non-estimated subset of $\alpha_C(t_0)$ (when a_C is not estimated, the subset includes all the chief absolute OE). Notice that when the ideal delta-cases are considered, the absolute orbit of the chief at t_0 is assumed perfectly known and the filter is initialized without error on $\alpha_C(t_0)$. When the state vector is augmented to include the chief semi-major axis (a_C), the chief absolute OE initial guess is obtained from $\alpha_C^{\text{TRUE}}(t_0)$ by adding zero-mean Gaussian noise. Specifically, the semi-major axis is corrupted with a standard deviation $\sigma_a = 100$ m. The remaining (non-estimated) chief OE are corrupted with standard deviations scaled according to $\sigma_{e,i,\Omega,\omega,M} = \sigma_a/a_C$, following the normalization proposed by Koenig and D'Amico [5]. Moreover, in all delta-cases, the filter is provided with fitted $a\delta\alpha$ to avoid polluting the analysis with additional parameters. Finally, the ROE state is initialized with a conservative error whose standard deviation is described in Table 14.

In conclusion, in the analysis of 3.4 there are four possible combinations of the delta-cases categories (ideal vs noisy and augmented vs ROE-only state). For each of those, a total of 100 cases are defined by the grid of batch lengths and sizes in Eq. (26), leading to 400 delta-cases in total, for which the batch-LSQ filter described in Section 2.3 provides an estimate.

Additional hyper-parameters: As suggested by Koenig and D'Amico [5], convergence of the filter is governed by the maximum number of iterations (19), $N_{\text{max}}^{\text{ITER}} = 5$; the minimum chief semi-major axis variation between iterations, $\Delta_{\text{min}}^{\text{ITER}}(a_C) = 1$ m (applicable only when a_C is estimated) and the minimum ROE-state component variation between iterations, $\Delta_{\text{min}}^{\text{ITER}}(a\delta\alpha) = 0.1$ m. Finally, the chief absolute OE are propagated by the analytical model regardless of whether a_C is estimated. The batch LSQ filter employs the analytical GCL-STM-dd alongside the exact numerical measurement model based on the chain of non-linear transformations of Eq. (13), so as to isolate as much as possible the influence of the sensitivity variables on the estimation error. This combination follows the sequential strategy described previously and in [41].

Evaluation metrics: The first performance metric considers the residuals between the predicted and true azimuth and elevation measurements after the last iteration of the batch-LSQ. By predicted, we mean the bearing angles obtained applying the propagated-modeled-measurement function (16) to the LSQ-estimate. Following the approach introduced by Gaias et al. [8], the quality of the estimate can be evaluated by computing the (sample) mean of azimuth and elevation residuals across the batch and their respective standard deviations:

$$\mu_{\rho\beta} = \frac{1}{|B|} \sum_{i=1}^{|B|} (\beta_i^{\text{PRED}} - \beta_i^{\text{TRUE}}) \quad (27a)$$

$$\sigma_{\rho\beta} = \sqrt{\frac{1}{|B|} \sum_{i=1}^{|B|} \left((\beta_i^{\text{PRED}} - \beta_i^{\text{TRUE}}) - \mu_{\rho\beta} \right)^2}. \quad (27b)$$

Where $\beta = \{\alpha, \epsilon\}$ and i indicates the element of the batch corresponding to the time instant t_i . The superscript “PRED” indicates that the quantity is obtained by re-propagating the final estimate. These two quantities provide insight into the consistency and spread of the measurement fit; however, they may not capture convergence issues [24]. Particularly, if the filter settles around an incorrect local minimum, it leads to very low values of these metrics but an erroneous estimate. To address this limitation, an additional metric is introduced based on state estimation accuracy. Specifically, the error in the estimated dimensional mean relative longitude ($a\delta\lambda(t_0)$). As highlighted in the analysis conducted by Sullivan [3], $a\delta\lambda$ is the weakly observable component of the ROE-state and when convergence is achieved in this component, all the other estimated elements of $a\delta\hat{\alpha}(t_0)$ converge as well. Thus, it is decided to use the estimation error on this quantity to evaluate the estimation accuracy:

$$e_{\delta\lambda} = \left| a\delta\lambda^{\text{LSQ}}(t_0) - a\delta\lambda^{\text{TRUE}}(t_0) \right|. \quad (28)$$

The superscript “LSQ” is used to specify that the quantity is the outcome of the estimation process after convergence conditions are met or the maximum number of iterations is reached.

An additional metric is needed when the filter is asked to estimate the chief semi-major axis $a_C(t_0)$. This can be defined in an equivalent manner to the previous:

$$e_{a_C} = \left| a_C^{\text{LSQ}}(t_0) - a_C^{\text{TRUE}}(t_0) \right|. \quad (29)$$

Together, these complementary metrics enable a comprehensive evaluation of the batch LSQ filter performance.

Additional secondary metrics can be defined to study the sensitivity of the observability to the aforementioned variables. We can leverage the fact that matrix $\mathbf{H}$ (20) is available from the LSQ-iterations (19). Then, we can immediately compute the OG (25). From the observability Gramian, it is straightforward to compute the OI and CN .

Metrics computation: The computation of the metrics starts with execution of the batch LSQ filter which iteratively process a batch of data-points of length Δt_B and size $|B|$ until a final estimate $\hat{\mathbf{x}}^{\text{LSQ}}(t_0)$ is reached. From this, the computation of the metrics of Eq. (28) just requires the extraction of the second component and to subtract from it the known ground-truth value (see Table 12). When $a_C(t_0)$ is estimated, the related metric of Eq. (29) is computed similarly.

The residual-based metrics defined in Eq. (27) require instead some additional steps. First, the ROE and chief OE contained in the final estimate $\hat{\mathbf{x}}^{\text{LSQ}}(t_0)$ are re-propagated to each time instant of the batch. Then, the measurement model is applied to each propagated state to obtain the predicted (“PRED”) azimuth and elevation at each time instant t_i . Considering then the “true” values of azimuth and elevation (these not corrupted by noise/error) at each time instant we can employ the definition in Eq. (27) to compute the mean residuals and the corresponding standard deviations.

4. Results

In this section the results of the analysis previously introduced are discussed. Then, a set of simple actionable guidelines will be formulated.

4.1. Propagation accuracy

Three performance *strata* are identifiable across the bounded RO1 family in Fig. 3. The least accurate one contains the KEP-STM exclusively: its Keplerian-only dynamics and the absence of any osculating-to-mean transformation at initialization produce a ν growing from roughly 10^1 m at close range to 10^5 m at $\delta r_{\text{max}}/a_C \approx 10^{-2}$. The intermediate *stratum* groups GDA-STM, GA-STM, KGD-STM and GCL-STM




Fig. 3. Propagation error metric. Unforced dynamics subject to all perturbations. Initial conditions: bounded cases of LEO1-RO1.

(all without differential drag). All four account for secular J_2 effects and apply a consistent osculating-to-mean transformation, placing them uniformly below KEP-STM. Within this *stratum*, GCL-STM achieves the lowest ν , with a gap relative to GA-STM and KGD-STM that widens with inter-satellite separation. This separation is attributable to the combined effect of the J_4, J_6 geo-potential terms, the second-order STT and the more accurate osculating-to-mean transformation used by GCL-STM, which is the only one adopting a second-order perturbation method. Although each contribution is individually modest, their cumulative effect becomes measurable at separation ratios approaching 10^{-2} . The non-monotonic behavior of GDA-STM, where ν decreases slightly before resuming its growth, reflects a partial cancellation between the initialization error and the dynamic modeling deficiency, which temporarily reduces the net propagation error before the latter dominates. The best *stratum* comprises KGD-STM-dd and GCL-STM-dd. The inclusion of the (fitted) empirical differential drag coefficients reduces ν by approximately one order of magnitude relative to the intermediate *stratum*, confirming that differential drag is a dominant perturbation for bounded relative motion at 500 km altitude with the AVANTI-like physical parameters of Table 4. Within this, GCL-STM-dd consistently achieves lower ν than KGD-STM-dd, with the gap widening at larger separations due to the compounding advantages of higher-order geo-potential modeling, the second-order STT and the more accurate initialization.

The same three-*strata* structure persists in the unbounded RO2 family of Fig. 4, but the dominant error mechanism shifts. KEP-STM and GDA-STM grow rapidly with $a\delta a$: KEP-STM because it lacks any J_2 modeling and GDA-STM because its simplified J_2 formulation neglects the secular coupling between $a\delta a$ and the J_2 , which becomes the primary error source beyond $a\delta a \approx 50$ m. GA-STM and KGD-STM remain approximately flat across the full $a\delta a$ range, confirming that their formulations correctly capture this coupling. GCL-STM maintains the lowest ν within the intermediate *stratum* across all cases, again reflecting the advantage of higher-order geo-potential modeling and better initialization. In the best *stratum*, KGD-STM-dd and GCL-STM-dd are very close at small $a\delta a$, where the DD fit absorbs the dominant secular error regardless of model fidelity. At larger $a\delta a$, the two models begin to diverge: KGD-STM-dd accumulates additional error from linearization and lower-fidelity




Fig. 4. Propagation error metric. Unforced dynamics subject to all perturbations. Initial conditions: unbounded cases of LEO1-RO2.

initialization, while GCL-STM-dd maintains a lower ν for the same reasons discussed in Fig. 3.

At 500 km altitude, empirical differential drag modeling proves to be a critical design choice for analytical relative motion propagation, outweighing the combined contribution of higher-order geo-potential terms, second-order dynamics and initialization fidelity. The one-order-of-magnitude reduction in ν achieved by DD-capable models over their non-DD counterparts holds across both bounded and drifting configurations at all tested separations. For onboard implementations targeting LEO altitudes at comparable altitudes, DD modeling is therefore a necessary condition for adequate propagation accuracy and its omission cannot be compensated by increasing the geo-potential degree.

4.2. Maneuvers accuracy

The propagation error metric ν (Fig. 5) compares impulsive and quasi-continuous maneuver modeling across four Δv magnitudes spanning three orders of magnitude (0.001–1 m/s) and burn durations ranging from 1 s to approximately one orbital period (5680 s). Under purely Keplerian dynamics, any discrepancy between the analytical propagation and the numerical ground-truth is attributable almost entirely to the maneuver model, as the KEP-STM reproduces unforced Keplerian motion with negligible error.

The impulsive model (dashed lines) exhibits ν that grows with burn duration, although the logarithmic scale partially masks this dependence. This growth reflects the fundamental limitation of the impulsive assumption: the entire Δv is applied instantaneously at a single orbital location regardless of actual burn arc. The quasi-continuous model (solid lines) shows the opposite trend: ν decreases monotonically with burn duration because distributing the same total Δv over a longer arc implies lower instantaneous thrust, and the finer discretization into equivalent impulses better approximates the true continuous profile. At short burn durations ($\Delta t^{\text{MAN}} < 7$ –10 s), both models yield comparable accuracy. Beyond this threshold, the models diverge progressively: at orbit-duration burns, the quasi-continuous model reduces ν by up to four orders of magnitude.

However, the maximum error metric ν advantages the quasi-continuous model in a potentially misleading way. When an impulsive model is used, the error accumulates during the first half of the physical




Fig. 5. Propagation error metric. Forced dynamics subject to Keplerian acceleration and thrust command. Initial conditions: bounded case of LEO1-RO3.




Fig. 6. Final time propagation error metric. Forced dynamics subject to Keplerian acceleration and thrust command. Initial conditions: bounded case of LEO1-RO3.

burn (since the modeled impulse is applied at midpoint), peaks at the instant of the “jump”, then decreases as the true and modeled trajectories reconverge. The metric ν captures this worst-case instant rather than the sustained error.

The final-time error metric ν_{end} (Fig. 6) provides a complementary view. The stratification by maneuver intensity persists, but the impulsive model now exhibits the expected behavior: for short burn durations, the error stabilizes rather than remaining at peak value. The quasi-continuous model remains nearly unaffected by maneuver duration. A clear threshold emerges at $\Delta t^{\text{MAN}} \approx 50$ s, beyond which the impulsive assumption breaks down definitively.

These results indicate that as maneuver magnitude increases, the impulsive assumption must be abandoned in favor of quasi-continuous modeling. Formation flying maneuvers are typically small (mm/s to cm/s), and for such magnitudes the impulsive model may provide adequate accuracy up to moderately long burn durations. A conservative operational guideline is to employ the impulsive model only




Fig. 7. Mapping error angular metric. Unforced dynamics subject to Keplerian acceleration only. Initial conditions: bounded cases of LEO1-RO4.

when $\Delta t^{\text{MAN}} < 10\text{--}20$ s. This threshold is set conservatively because measurements acquired during the burn interval may coincide with the transient error peak, potentially causing large residuals and filter divergence.

These conclusions are expected to remain valid under perturbed dynamics and may even be amplified by perturbation-induced errors. In fact, during the free-propagation arcs, the error introduced by the maneuver will accumulate. Secondly, the maneuvers enter through the input matrix, which depends on the mean latitude and semi-major axis (via n_C). Errors in u_C and a_C therefore lead to an inaccurate input matrix, and the resulting maneuver modeling error is expected to be adversely affected.

4.3. Measurement accuracy

The angular error metric θ , reported in Fig. 7, stratifies the mapping models into three performance *strata*. The exact non-linear transformation NUM-e (Eq. (13)) settles at negligible errors, confirming that it achieves the absolute accuracy floor set by analytical KEP-STM propagation. The simplest mappings based on the GDA-MAP occupy the highest portion of the plot. Even when the curvilinear correction is applied, the error remain unacceptably large, especially at intermediate/far ranges. The curvilinear-to-rectilinear mapping proves however to be quite useful when applied to the GA-MAP. As the inter-satellite range grows, the importance of modeling orbit curvature increases and the benefits of applying this transformation appear clear. The curvilinear-to-rectilinear correction accounts for orbit curvature and reduces θ by approximately one order of magnitude at far range. The second order mapping GAC-MAP, which also applies the aforementioned curvilinear-to-rectilinear transform, achieves results virtually indistinguishable from the GA-MAP-c.

The second order mapping GAC-MAP, which also applies the aforementioned curvilinear-to-rectilinear transform, achieves results virtually indistinguishable from the GA-MAP-c. On the other hand, the WD-MAP stays at the absolute floor error up to intermediate ranges, after which the errors grows rapidly, but still outperforming all other linearized model.

GAC-MAP-e, which is a simplified version of NUM-e, exhibits a non-monotonic trend: θ initially decreases before increasing with $a\delta\lambda$. At




Fig. 8. Mapping error magnitude metric. Unforced dynamics subject to Keplerian acceleration only. Initial conditions: bounded cases of LEO1-RO4.

very close range, the line-of-sight subtends a large solid angle, amplifying the geometric sensitivity to small position errors. At intermediate separations, the deputy-chief geometry stabilizes, reducing this sensitivity. Beyond $a\delta\lambda \approx 10^3$ m, even minor errors in true anomaly accumulate over the 24-hour propagation arc, re-emerging as measurable angular deviations. In any case, the error metric $\theta < 10^{-1}$ arcsec across the full range.

Typical vision-based sensor noise lies in the range 25–50 arcsec [3], two to three orders of magnitude above the accuracy floor. Adopting the criterion $\theta < \sigma_{\text{noise}}$ yields conservative validity bounds: GDA-MAP and GA-MAP-r remain acceptable up to $a\delta\lambda \approx 10^2$ m. The validity is extended by one order of magnitude for GA-MAP-c and GAC-MAP. On the other hand, WD-MAP and the simplified exact mapping satisfy the criterion across the entire tested range.

The magnitude error metric γ (Fig. 8) provides the starkest differentiation among models and complements the previous angular metric. In fact, the error spans six orders of magnitude between the exact transformation and first-order rectilinear mappings at far range. NUM-e remains at the numerical-error floor ($\gamma \approx 10^{-8}$ m) regardless of inter-satellite separation. WD-MAP performs exceptionally well up to intermediate range ($a\delta\lambda \approx 10^4$ m), maintaining centimeter-level accuracy and slowly degrades to meter-level error at $a\delta\lambda \approx 10^5$ m. GAC-MAP-e, while producing larger error relative to WD-MAP at short ranges, produces very consistent and always acceptable performances over the full range. The other second order model GAC-MAP instead produces worse results with respect to the GA-MAP family. This might be due to the along-track refinement of the mapping which not applied in the radial and out-of-plane directions. The worst overall results are again obtained by the GDA-MAP whose simple but simplistic model quickly becomes ineffective as the range increases.

The observability index (Fig. 9) and condition number (Fig. 10) reveal that mapping fidelity has a limited but non-negligible influence on observability. All curvilinear and higher-order models collapse onto nearly identical curves: OI decreases from $\approx 10^0$ to a plateau near 10^{-12} , while CN rises correspondingly from $\approx 10^5$ to $\approx 10^{10}$. Rectilinear mappings, however, continue to degrade beyond this plateau: OI drops toward 10^{-20} and CN climbs to 10^{15} at far range. This divergence reflects the geometric inconsistency of applying a rectilinear approximation to inherently curvilinear motion; the mismatch reduces the information content extractable from bearing measurements and




Fig. 9. Observability index. Unforced dynamics subject to Keplerian acceleration only. Initial conditions: bounded cases of LEO1-RO4.




Fig. 10. Condition number. Unforced dynamics subject to Keplerian acceleration only. Initial conditions: bounded cases of LEO1-RO4.

increases sensitivity to measurement noise. Observability is governed primarily by the dynamics, but neglecting orbit curvature produces a measurable penalty.

For operational filter design, GAC-MAP-e and WD-MAP offer accuracy close to the exact transformation while avoiding iterative solution of Kepler's equation and the several transformations of Eq. (13), making them the recommended default. At close range ($a\delta\lambda < 10^2$ m), first-order rectilinear mappings introduce errors well below sensor noise and remain acceptable where computational simplicity is prioritized.

4.4. Optimal batch length and size

The results of this analyses is presented for each of the four delta-cases, starting from the ideal ones and then presenting how the results are affected by noise.




Fig. 11. Azimuth residual statistics. Unforced dynamics subject to all perturbations. Initial conditions: unbounded case of LEO1-RO5. Delta case: ID, only $a\delta\alpha(t_0)$ estimated.

4.4.1. ID: Delta-cases without semi-major axis estimated

The azimuth residual statistics are reported in Fig. 11. They exhibit a pattern of increasing standard deviations with longer batch lengths. The spread starts from being almost null at $\Delta t_B = 0.1T_C$ and exceeds ± 100 arcsec at $\Delta t_B = 15T_C$. As explained by [8,24], the near-zero mean across a broad range of batch lengths indicates adequate modeling fidelity. In fact, a biased mean would signal unmodeled dynamics or systemic modeling errors. However, the steady growth in standard deviation does suggest a worsening of model quality over longer propagation lengths, which is to be expected given the assumptions needed to derive the model [38].

The elevation residual statistics (Fig. 12) follow a similar pattern. However, a slight negative bias emerges at batch lengths exceeding $7.6T_C$. This might be additional proof that the analytical propagation model is becoming progressively less accurate over extended arcs. Sullivan [3] explored the geometric connection between the elevation angle mean (over one period) and $(a\delta\lambda, a\delta a)$. When the VBS is oriented in the anti-flight direction, the elevation is the angular equivalent of the $\hat{\mathbf{R}}$ unit vector of the RTN frame. Thus, an increase in mean relative longitude causes a negative trend while an increase in relative semi-major axis a positive one. Given this geometric connection, the appearance of bias in elevation residuals might also correlate with increased estimation error in $a\delta\lambda$ (and of $a\delta a$).

The mean relative longitude estimation error confirms this expectation, as displayed in Fig. 13. At short batch lengths, the estimation error exceeds the initialization error (dashed line), with minor but existing dependence on batch size: $|B| = 10$ yields approximately 6 km error while $|B| = 500$ reduces this to approximately 3 km. The green and red markers at the grid corners indicate the best and worst results, respectively. A distinctive double-dip pattern emerges, with minima near $0.5T_C$ and $3 - 4T_C$. The first minimum arises because differential drag has insufficient time to accumulate significant effect over such a short arc, allowing a good estimate from limited but dynamically “clean” data. The second minimum represents the optimal trade-off where sufficient geometric diversity has been collected while propagation errors remain bounded. Beyond $\Delta t_B \approx 7T_C$, the filter diverges for most configurations, consistent with the elevation residual bias observed in Fig. 12. This double-dip behavior confirms the nonlinearity and non-convexity of the AO-ROD problem, further confirming that the filter can converge to suboptimal local minima which might depend not only on the initial guess, but also on the batch length.

The observability index (Fig. 14) increases monotonically with batch length, rising from approximately 10^{-20} at $\Delta t_B = 0.1T_C$ to




Fig. 12. Elevation residual statistics. Unforced dynamics subject to all perturbations. Initial conditions: unbounded case of LEO1-RO5. Delta case: ID, only $a\delta\alpha(t_0)$ estimated.




Fig. 13. Mean relative longitude estimation error. Unforced dynamics subject to all perturbations. Initial conditions: unbounded case of LEO1-RO5. Delta case: ID, only $a\delta\alpha(t_0)$ estimated.

approximately 10^{-15} beyond $5T_C$. This trend reflects the intuitive expectation that observing a larger portion of the relative orbit improves its observability. Notably, increasing batch size $|B|$ also increases OI monotonically at fixed batch length, indicating improved robustness to measurement noise [32]. This is again supported by the intuition that more data points enable better filtering of noise-induced fluctuations.

The condition number (Fig. 15) exhibits the complementary trend, decreasing from approximately 10^{12} at short batch lengths to approximately 10^6 as Δt_B increases. The near-complete insensitivity to $|B|$ suggests that batch length governs the fundamental observability structure, while batch size primarily influences noise rejection.

4.4.2. ID: Delta-cases with semi-major axis estimated

The azimuth (Fig. 16) and elevation (Fig. 17) residual statistics exhibit patterns coherent with the ROE-only estimation case, but with some important differences. In fact, the mean elevation residuals present stronger oscillation over Δt_B , regardless of the batch size. The reason of the difference is that now the filter is solving a more demanding task. In the previous delta-case, a_C was not estimated and the true value of $a_C(t_0)$ was provided to the filter. Now, also the guess on $a_C(t_0)$ is corrupted by an error of ≈ 55 m. However, an error




Fig. 14. Observability index. Unforced dynamics subject to all perturbations. Initial conditions: unbounded case of LEO1-RO5. Delta case: ID, only $a\delta\alpha(t_0)$ estimated.




Fig. 15. Condition number. Unforced dynamics subject to all perturbations. Initial conditions: unbounded case of LEO1-RO5. Delta case: ID, only $a\delta\alpha(t_0)$ estimated.

on $a_C(t_0)$ reflects on the elevation trend as an additional semi-major axis difference. Thus, now the mean elevation residuals are affected by both estimation errors in $a\delta\lambda$ and in a_C . This complicates the diagnostic process, since the two effects might (partially) cancel each other. This explains the almost absent bias at small batch lengths, when the estimation error on both mean relative longitude and chief semi-major axis are the greatest, as clear from Fig. 18 and Fig. 19. Basically, the filter fits very well the bearing angles trend, producing very low residuals (almost unbiased) but does so converging to estimates worse than the initial guess. This also happens for $\Delta t_B \geq 6$, where both the estimation errors are increasing and the elevation residuals display only a small positive bias. However, elevation residuals develop a stronger bias when small batch sizes are employed, particularly for longer batch lengths. Thus we might think to use a down-sampled batch to check the results obtained with all the available data-points. Both the mean and standard deviation of the residuals seem to remain strong indicators of filter health.

The mean relative longitude estimation error (Fig. 18) shows an improvement over the ROE-only case. The double-dip pattern persists, but error magnitudes are reduced: the filter achieves sub-2 km accuracy by $\Delta t_B \approx 2T_C$ and sub-500 m at longer batches. These values correspond to $\approx 7\%$ and $\approx 2\%$ of the true range, respectively.




Fig. 16. Azimuth residual statistics. Unforced dynamics subject to all perturbations. Initial conditions: unbounded case of LEO1-RO5. Delta case: ID, $a\delta\alpha(t_0)$ and $a_c(t_0)$ estimated.




Fig. 17. Elevation residual statistics. Unforced dynamics subject to all perturbations. Initial conditions: unbounded case of LEO1-RO5. Delta case: ID, $a\delta\alpha(t_0)$ and $a_c(t_0)$ estimated.




Fig. 18. Mean relative longitude estimation error. Unforced dynamics subject to all perturbations. Initial conditions: unbounded case of LEO1-RO5. Delta case: ID, $a\delta\alpha(t_0)$ and $a_c(t_0)$ estimated.




Fig. 19. Chief semi-major axis estimation error. Unforced dynamics subject to all perturbations. Initial conditions: unbounded case of LEO1-RO5. Delta case: ID, $a\delta\alpha(t_0)$ and $a_c(t_0)$ estimated.




Fig. 20. Azimuth residual statistics. Unforced dynamics subject to all perturbations. Initial conditions: unbounded case of LEO1-RO5. Delta case: NY, only $a\delta\alpha(t_0)$ estimated.

The chief semi-major axis estimation error (Fig. 19) decreases from approximately 0.4 km at short batch lengths toward 0.05 km at $\Delta t_B > 7T_C$. The double-dip structure is also present but shifted toward longer batch lengths compared to $e_{\delta\lambda}$, indicating that a_c requires more extended observation arcs to achieve comparable convergence.

The observability index and condition number do not vary appreciably when the estimated state is augmented with $a_c(t_0)$; these metrics are therefore omitted to avoid redundancy.

4.4.3. NY: Delta-cases without semi-major axis estimated

The azimuth (Fig. 20) and elevation (Fig. 21) residual statistics reflect the presence of 40 arcsec measurement noise. Residual spread increases accordingly, but the overall convergence behavior is preserved. Reaching the noise floor in most configurations provides a positive indication of filter health. The elevation residual bias at extended batch lengths persists and remains correlated to a jump in estimation error of $a\delta\lambda$, as clear from Fig. 22. Again, the smallest batch size ($|B| = 10$) highlights this property.

With respect to the ID counterpart, the mean relative longitude estimation error exhibits greater sensitivity to batch size and the double-dip becomes harder to spot (it remains valid for $|B| = 445$). Indeed, we can see that larger $|B|$ yields lower errors and reduced oscillations




Fig. 21. Elevation residual statistics. Unforced dynamics subject to all perturbations. Initial conditions: unbounded case of LEO1-RO5. Delta case: NY, only $a\delta\alpha(t_0)$ estimated.




Fig. 22. Mean relative longitude estimation error. Unforced dynamics subject to all perturbations. Initial conditions: unbounded case of LEO1-RO5. Delta case: NY, only $a\delta\alpha(t_0)$ estimated.

in $e_{\delta\lambda}$. As batch length increases, this sensitivity diminishes progressively, indicating that increasingly important propagation errors, united with noise, make the filtering process divergent. Despite the increased variability, some features of the underlying pattern from the ideal case remains discernible: the best estimates are obtained consistently at approximately $1 T_C$ and $4 - 6 T_C$, for most batch sizes. A somewhat unexpected behavior is instead observed for small batch length, where the LSQ converges to quite good estimates, despite the intuition would suggest that noise would degrade the results even more with respect to Fig. 13. Instead, the presence of noise is able to “unstuck” the estimate from the local minima obtained in the ID case, leading to an improved estimate. Clearly, it is not guaranteed that different noise levels would always improve the estimate.

The observability metrics do not change significantly and are therefore omitted.

4.4.4. NY: Delta-cases with semi-major axis estimated

This delta-case acts collects and reproduce most of the features identified in previous ones. The residual statistics, reported in Figs. 23 and 24, reach the noise floor for most batch configurations. However, several sudden increases of the standard deviations are registered and




Fig. 23. Azimuth residual statistics. Unforced dynamics subject to all perturbations. Initial conditions: unbounded case of LEO1-RO5. Delta case: NY, $a\delta\alpha(t_0)$ and $a_C(t_0)$ estimated.




Fig. 24. Elevation residual statistics. Unforced dynamics subject to all perturbations. Initial conditions: unbounded case of LEO1-RO5. Delta case: NY, $a\delta\alpha(t_0)$ and $a_C(t_0)$ estimated.

are especially evident in the azimuth plot. Notice that in correspondence of this peaks, the LSQ converges to bad estimates of $a\delta\lambda$ and/or a_C . This confirms the diagnostic power of the standard deviation of the residuals (not only the elevation but also the one of the azimuth). However, the plot of elevation residuals (Fig. 24) is once more of particular interest. A clear bias is present even at short batch lengths (independently of $|B|$) and at long batch lengths (when $|B| = 10$). Consider now the estimation errors $e_{\delta\lambda}$ and e_{a_C} plots, reported in (Fig. 25) and Fig. 26, respectively. At short batch lengths, $a\delta\lambda$ is well estimated, while a_C worsens with respect to the initial guess. On one hand, it confirms the fact that a_C requires longer batches to be correctly identified. On the other hand, the large estimation error e_{a_C} cannot be compensated by $e_{\delta\lambda}$, thus generating the large elevation bias seen in Fig. 24. The opposite is true for long batches of data, where the estimation errors are relatively large in both parameters. In fact, they seem to compensate each other, making the elevation bias almost null. Again, the small batch size $|B| = 10$ still displays additional bias, but the effect might be due to a good estimate of $a\delta\lambda$.

The mean relative longitude estimation error (Fig. 25) shows improved convergence relative to the ROE-only noisy case, with errors reduced to the $0.5 - 3$ km range and more consistent behavior across batch configurations.




Fig. 25. Mean relative longitude estimation error. Unforced dynamics subject to all perturbations. Initial conditions: unbounded case of LEO1-RO5. Delta case: NY, $a\delta a(t_0)$ and $a_C(t_0)$ estimated.




Fig. 26. Chief semi-major axis estimation error. Unforced dynamics subject to all perturbations. Initial conditions: unbounded case of LEO1-RO5. Delta case: NY, $a\delta a(t_0)$ and $a_C(t_0)$ estimated.

The chief semi-major axis estimation error (Fig. 26) remains in the 20–60 m range across batch configurations, with a clear minimum near $4 T_C$.

The observability metrics are again omitted as they do not vary appreciably.

4.5. Guidelines

The following recommendations synthesize the sensitivity analyses into actionable design guidance for angles-only relative navigation filters. While developed for batch LSQ architectures, the propagation and measurement model recommendations apply directly to sequential filters (EKF, UKF), as these components are shared across filtering paradigms.

Propagation model selection

- At LEO altitudes where atmospheric drag and different ballistic coefficients are significant, empirical differential drag modeling is essential. DD-capable models (KGD-STM-dd, GCL-STM-dd) can reduce propagation error up to approximately one order of magnitude compared to their non-DD counterparts. This conclusion is consistent with flight experience: the ARGON experiment at

750 km altitude operated successfully without DD modeling [2], whereas the AVANTI experiment at 550 km required explicit DD treatment [1].

- Among DD-capable models, GCL-STM-dd achieves the lowest propagation error due to higher-order geo-potential terms, second-order STT and consistent high-fidelity initialization. KGD-STM-dd provides a computationally lighter alternative with modest accuracy penalty.
- For bounded relative orbits, all J_2 -capable models perform adequately. For drifting configurations ($a\delta a \neq 0$), models must capture the secular coupling between $a\delta a$ and J_2 ; GDA-STM does not and should be avoided in rendezvous scenarios with fast drifting orbits.
- Each propagation model should be paired with an osculating-to-mean transformation of matching fidelity.

Maneuver modeling

- Under Keplerian dynamics, the impulsive assumption remains valid for burn durations $\Delta t^{\text{MAN}} < 10\text{--}20$ s regardless of maneuver magnitude. This conservative threshold accounts for the transient error peak that occurs when the modeled impulse is applied: measurements acquired during this interval may produce large residuals leading to filter divergence.
- Beyond this threshold, quasi-continuous modeling should be adopted. At orbit-duration burns, it reduces maximum propagation error by up to four orders of magnitude relative to the impulsive model.
- Formation flying maneuvers are typically small (mm/s to cm/s). For such magnitudes, final-time errors remain bounded up to $\Delta t^{\text{MAN}} \approx 50$ s, but the conservative threshold of 10–20 s is recommended to avoid transient peak errors during measurement acquisition.
- Under perturbed dynamics, the conclusions are expected to extend beyond the Keplerian test case, potentially with reduced achievable accuracy. This further supports adopting a conservative threshold.

Measurement model selection

- GAC-MAP-e and WD-MAP are recommended as default choices. Both offer accuracy close to the exact transformation while avoiding iterative solution of Kepler's equation. GAC-MAP-e provides consistent performance across the full range. WD-MAP maintains centimeter-level magnitude accuracy up to $a\delta\lambda \approx 10^4$ m and degrades slowly to meter-level at far range.
- First-order mappings with curvilinear correction (GA-MAP-c) and second-order GAC-MAP satisfy $\theta < \sigma_{\text{noise}}$ up to $a\delta\lambda \approx 10^3$ m and are suitable for mid-range operations.
- First-order rectilinear mappings (GDA-MAP-r, GA-MAP-r) introduce errors below sensor noise only up to $a\delta\lambda \approx 10^2$ m. Their use should be restricted to close-range proximity operations where computational simplicity is prioritized.
- The GDA-MAP family produces the worst overall results; its simple formulation becomes ineffective as range increases and should be avoided beyond close range.
- Mapping fidelity has secondary but non-negligible influence on observability. Rectilinear approximations produce measurable degradation in OI and CN at far range, reflecting the geometric inconsistency of neglecting orbit curvature.

Batch configuration

The results of the batch configuration analysis are significant primarily for the trends they reveal rather than for specific numerical values. Different initialization guesses, noise realizations, convergence tolerances or orbital geometries will influence the precise error magnitudes; however, the qualitative behavior, observability-accuracy trade-off, divergence at excessive batch lengths, generalizes more robustly

than specific numerical values. From an operational perspective, although illumination constraints are disabled in this analysis, the sensitivity study directly informs practical design choices: batch length can be controlled by selecting the observation arc, while batch size can be influenced through the VBS sampling rate.

- Batch length governs fundamental observability; batch size influence noise rejection. Increasing Δt_B improves OI and reduces CN until propagation errors accumulate; increasing $|B|$ improves robustness to measurement noise at fixed batch length.
- The estimation error in $a\delta\lambda$ exhibits a characteristic double-dip pattern as a function of batch length. The first minimum arises from dynamically clean short arcs before DD accumulates; the second minimum represents the optimal trade-off between geometric diversity and propagation error.
- Augmenting the estimated state with the chief semi-major axis a_C improves $a\delta\lambda$ estimation accuracy and enables joint refinement of the reference orbit. However, a_C requires longer observation arcs than $a\delta\lambda$ to achieve comparable convergence. Moreover, estimation errors in both parameters affect elevation residuals through similar geometric mechanisms, and their effects may partially cancel, complicating residual-based diagnostics.
- Residual statistics, both mean and standard deviation, are primary diagnostics for filter health. Large residual standard deviations indicate poor filter performance; the underlying cause may be excessively long propagation arcs (where model approximations accumulate) or convergence to an incorrect local minimum. Biased residual means signal unmodeled dynamics or systematic modeling errors. These diagnostics are observable in both azimuth and elevation channels.
- Elevation residual bias provides additional diagnostic information linked to specific state components. When the VBS is oriented in the anti-flight direction, systematic elevation bias correlates with estimation error in $a\delta\lambda$ and a_C due to the geometric coupling identified by Sullivan [3]. This diagnostic is accentuated when small batch sizes are combined with long batch lengths; a down-sampled batch can therefore serve as a cross-check on results obtained with all available data. However, a small value cannot be considered conclusive, due to the competition of the effects due to estimation errors in $a\delta\lambda$ and a_C .
- Residuals approaching zero do not guarantee convergence to the correct solution Ardaens [24]. The filter may fit bearing angles well while converging to estimates worse than the initial guess. The estimation errors $e_{\delta\lambda}$ and e_{a_C} should be monitored as complementary convergence indicators when independent estimates are available and examining their covariance.

5. Conclusions

This work presented a systematic sensitivity analysis of the four key modeling components of an angles-only batch least-squares relative orbit determination filter: unforced dynamics, maneuver modeling, measurement mapping and batch configuration. The analysis was conducted across representative near-circular LEO scenarios, including bounded and drifting relative orbits at 500 km altitude where differential drag is significant.

Four principal contributions emerge from this study. First, empirical differential drag modeling reduces propagation error by approximately one order of magnitude compared to non-DD models, establishing DD treatment as a necessary condition for accurate analytical propagation at low LEO altitudes; this improvement cannot be compensated by increasing geo-potential degree. Second, the impulsive maneuver assumption remains valid for burn durations below 10 s, beyond which quasi-continuous modeling yields up to four orders of magnitude improvement in propagation accuracy. Third, the simplified exact mapping GAC-MAP-e and the second-order WD-MAP

maintain sub-arcsecond angular accuracy across the full tested range; first-order curvilinear mappings and GAC-MAP remain acceptable up to approximately 10^3 m separation, while first-order rectilinear mappings are acceptable only below approximately 10^2 m. Fourth, batch length governs observability whereas batch size influences noise rejection capability; the estimation error exhibits a characteristic structure whose qualitative features, rather than specific numerical optima, generalize across scenarios. Residual statistics are confirmed as primary filter health diagnostics, confirming once more that residuals approaching zero do not guarantee correct convergence.

The propagation and measurement model recommendations apply directly to sequential filters (EKF, UKF), as these components are shared across filtering architectures. The batch configuration insights provide useful starting points for sequential implementations, though the degree to which these guidelines transfer warrants dedicated investigation. Future work should extend this analysis to eccentric orbits, where the perturbative environment changes substantially: solar radiation pressure and third-body effects become dominant at higher altitudes and eccentricities and the assumption of quasi-circular dynamics underlying several of the tested models no longer holds. Additionally, a systematic comparison of how batch-derived configuration guidelines translate to sequential filter tuning would complete the bridge between the two filtering paradigms.

CRedit authorship contribution statement

Alessandro Scalvini: Writing – review & editing, Writing – original draft, Visualization, Validation, Software, Methodology, Investigation, Formal analysis, Data curation, Conceptualization. **Giacomo Borelli:** Writing – review & editing, Supervision, Software. **Gabriella Gaia:** Writing – review & editing, Supervision, Software, Resources, Project administration, Funding acquisition.

Declaration of generative AI and AI-assisted technologies in the manuscript preparation process

During the preparation of this work the authors used Anthropic Claude (Opus 4.6 version) in order to assist with manuscript compliance review, proofreading, rephrasing, language refinement, code debugging and figure formatting. After using this tool, the authors reviewed and edited the content as needed and take full responsibility for the content of the published article.

Declaration of competing interest

The authors declare that they have no known competing financial interests or personal relationships that could have appeared to influence the work reported in this paper.

Acknowledgments

This research is funded by Infinite Orbits SAS, under the Agreement for Activation and Financing of n. 1 Ph.D. with Doctoral Scholarship in “Aerospace Engineering” 39° Cycle. The authors acknowledge GNC Engineer Daniel Opoka and GNC Engineer Raffaele Tarone for their support.

References

- [1] G. Gaia, J.-S. Ardaens, Flight demonstration of autonomous noncooperative rendezvous in low Earth orbit, *J. Guid. Control Dyn.* 41 (2018) 1337–1354, <http://dx.doi.org/10.2514/1.g003239>.
- [2] S. D’Amico, J.-S. Ardaens, G. Gaia, H. Benninghoff, B. Schlepp, J.L. Jørgensen, Noncooperative rendezvous using angles-only optical navigation: System design and flight results, *J. Guid. Control Dyn.* 36 (2013) 1576–1595, <http://dx.doi.org/10.2514/1.59236>.

- [3] J.A. Sullivan, Nonlinear Angles-Only Orbit Estimation for Autonomous Distributed Space Systems, Ph.D. thesis, Stanford University, 2020, URL: <https://purl.stanford.edu/yx801sv8482>.
- [4] J. Kruger, Flight Algorithms for Autonomous Tracking And Navigation of Distributed Space Systems Using Inter-Satellite Bearing Angles, Ph.D. thesis, Stanford University, 2024, URL: https://justin.zone/thesis_kruger_final_online.pdf.
- [5] A.W. Koenig, S. D'Amico, Observability-aware numerical algorithm for angles-only initial relative orbit determination, in: Proceedings of the AAS/AIAA Astrodynamics Specialist Conference, 2020, URL: https://slab.sites.stanford.edu/sites/g/files/sbiybj25201/files/media/file/bod_aas_2020_final.pdf.
- [6] D.C. Woffinden, D.-K. Geller, Observability criteria for angles-only navigation, IEEE Trans. Aerosp. Electron. Syst. 45 (2009) 1194–1208, <http://dx.doi.org/10.1109/taes.2009.5259193>.
- [7] P. Bodin, R. Noteborn, R. Larsson, T. Karlsson, S. D'Amico, J.-S. Ardaens, M. Delpech, J.-C. Berges, Risma formation flying demonstrator: Overview and conclusions from the nominal mission, in: Proceedings of the 35th Annual AAS Guidance and Control Conference, 2012, URL: https://www.researchgate.net/profile/Jean-Claude-Berges/publication/259897193_, paper AAS 12-072.
- [8] G. Gaïas, S. D'Amico, J.-S. Ardaens, Angles-only navigation to a noncooperative satellite using relative orbital elements, J. Guid. Control Dyn. 37 (2014) 439–451, <http://dx.doi.org/10.2514/1.61494>.
- [9] H. Reile, E. Lorenz, T. Terzibaschian, The FireBird mission – A scientific mission for Earth observation and hot spot detection, in: Small Satellites for Earth Observation: Digest of the 9th International Symposium of the International Academy of Astronautics, Wissenschaft und Technik Verlag, Berlin, Germany, 2013, pp. 184–196, URL: https://elib.dlr.de/83866/1/IAA-B9-0201_DLR_FireBird_Mission_extended_abstract.pdf.
- [10] J. Kruger, A.W. Koenig, S. D'Amico, Starling formation-flying optical experiment (StarFOX): System design and preflight verification, J. Spacecr. Rockets 60 (2023) 1755–1777, <http://dx.doi.org/10.2514/1.a35598>.
- [11] J. Kruger, S.S. Hwang, S. D'Amico, Starling formation-flying optical experiment: Initial operations and flight results, 2024, URL: <https://arxiv.org/abs/2406.06748>.
- [12] NASA, Starling mission final report: Swarm technology validation, 2024.
- [13] J. Sullivan, A.W. Koenig, J. Kruger, S. D'Amico, Generalized angles-only navigation architecture for autonomous distributed space systems, J. Guid. Control Dyn. 44 (2021) 1087–1105, <http://dx.doi.org/10.2514/1.g005439>.
- [14] J.J. Kruger, T. Guffanti, T.H. Park, M. Murray-Cooper, S.Y.W. Low, T. Bell, S. D'Amico, C.W. Roscoe, J. Westphal, Adaptive end-to-end architecture for autonomous spacecraft navigation and control during rendezvous and proximity operations, in: Proceedings of the AIAA SciTech Forum, American Institute of Aeronautics and Astronautics, 2024, p. 0430, <http://dx.doi.org/10.2514/6.2024-0430>.
- [15] J. Sullivan, S. D'Amico, Nonlinear kalman filtering for improved angles-only navigation using relative orbital elements, J. Guid. Control Dyn. 40 (2017) 2183–2200, <http://dx.doi.org/10.2514/1.g002719>.
- [16] J. Kruger, S. D'Amico, Autonomous angles-only multitarget tracking for spacecraft swarms, Acta Astronaut. 189 (2021) 514–529, <http://dx.doi.org/10.1016/j.actaastro.2021.08.049>.
- [17] J. Kruger, S. D'Amico, Observability analysis and optimization for angles-only navigation of distributed space systems, Adv. Space Res. 73 (2024) 5464–5483, <http://dx.doi.org/10.1016/j.asr.2023.08.055>.
- [18] C. Weeden, C. Blackerby, N. Okada, E. Yamamoto, J. Forshaw, J. Auburn, Industry implementation of the long-term sustainability guidelines: An astroscale perspective, in: 70th International Astronautical Congress, IAC, 2019.
- [19] A. Scalvini, A. Fois, A. Suarez, S.R. Nekoo, A. Ollero, Adaptive state-dependent Riccati equation attitude control for rotation-flying dual-arm manipulators, in: Proceedings of the 2024 International Conference on Space Robotics, ISPaRo, IEEE, Luxembourg, 2024, pp. 372–379, <http://dx.doi.org/10.1109/isparo60631.2024.10687494>, URL: <https://ieeexplore.ieee.org/abstract/document/10687494>.
- [20] R. Biesbroek, S. Aziz, A. Wolahan, S. Cipolla, M. Richard-Noca, L. Piguet, The clearspace-1 mission: Esa and clearspace team up to remove debris, in: Proceedings of the 8th European Conference on Space Debris, ESA Space Debris Office, 2021, pp. 1–3, URL: <https://conference.sdo.esoc.esa.int/proceedings/sdc8/paper/320>.
- [21] S. Woicke, J. Jipp, C. Steimle, N. Pokrupa, L. Metrailler, Adrios clearspace-1: In orbit demonstration of the removal of a non-cooperative spacecraft, in: Proceedings of the 9th European Conference on Space Debris, ESA Space Debris Office, 2025, URL: <https://conference.sdo.esoc.esa.int/proceedings/sdc9/paper/317>.
- [22] Y. Chihabi, S. Ulrich, Nonlinear adaptive angle-only relative navigation on perturbed eccentric orbits, Acta Astronaut. 228 (2025) 494–514, <http://dx.doi.org/10.1016/j.actaastro.2024.12.012>, URL: <https://www.sciencedirect.com/science/article/pii/S0094576524007598>.
- [23] J.-S. Ardaens, G. Gaïas, Flight demonstration of spaceborne real-time angles-only navigation to a noncooperative target in low Earth orbit, Acta Astronaut. 153 (2018) 367–382, <http://dx.doi.org/10.1016/j.actaastro.2018.01.044>.
- [24] J.-S. Ardaens, Angles-Only Relative Navigation in Low Earth Orbit, Ph.D. thesis, Delft University of Technology, 2020, URL: <https://pure.tudelft.nl/ws/portafiles/portal/82719048/dissertation.pdf>.
- [25] D.C. Woffinden, D.-K. Geller, Optimal orbital rendezvous maneuvering for angles-only navigation, J. Guid. Control Dyn. 32 (2009) 1382–1387, <http://dx.doi.org/10.2514/1.45006>.
- [26] J. Sullivan, S. Grimberg, S. D'Amico, Comprehensive survey and assessment of spacecraft relative motion dynamics models, J. Guid. Control Dyn. 40 (2017) 1837–1859, <http://dx.doi.org/10.2514/1.g002309>.
- [27] G. Gaïas, M. Lara, C. Colombo, Accurate osculating/mean orbital elements conversions for spaceborne formation flying, in: Proceedings of the 27th International Symposium on Space Flight Dynamics, ISSFD, 2019, pp. 1–15.
- [28] J. Grzymisch, W. Fichter, Observability criteria and unobservable maneuvers for in-orbit bearings-only navigation, J. Guid. Control Dyn. 37 (2014) 1250–1259, <http://dx.doi.org/10.2514/1.62476>.
- [29] M. Chernick, S. D'Amico, Closed-form optimal impulsive control of spacecraft formations using reachable set theory, J. Guid. Control Dyn. 44 (2021) 25–44, <http://dx.doi.org/10.2514/1.g005218>.
- [30] G. Gaïas, J.-S. Ardaens, C. Colombo, Precise line-of-sight modelling for angles-only relative navigation, Adv. Space Res. 67 (2021) 3515–3526, <http://dx.doi.org/10.1016/j.asr.2020.05.048>.
- [31] M.B. Willis, Analytical Theory of Satellite Relative Motion with Applications to Autonomous Navigation and Control, Ph.D. thesis, Stanford University, 2023.
- [32] J. Wang, E.A. Butcher, T.A. Lovell, Use of nonlinearities for increased observability in relative orbit estimation, in: Astrodynamics 2015: Proceedings of the AAS/AIAA Astrodynamics Specialist Conference, in: Advances in the Astronautical Sciences, vol. 156, American Astronautical Society / American Institute of Aeronautics and Astronautics, Univelt, Inc. Vail, Colorado, United States, 2016, pp. 3525–3544, URL: https://www.researchgate.net/profile/Jingwei-Wang/publication/318447396_, aAS/AIAA Astrodynamics Specialist Conference, August 9–13 2015.
- [33] D.A. Vallado, Fundamentals of astrodynamics and applications, second ed. in: Space Technology Library, vol. 12, Microcosm Press, 2001, <http://dx.doi.org/10.2514/2.4291>.
- [34] S. D'Amico, Autonomous Formation Flying in Low Earth Orbit, Ph.D. thesis, Delft University of Technology, 2010, URL: https://elib.dlr.de/63481/1/Damico_PhD_01022010.pdf.
- [35] G. Gaïas, J.-S. Ardaens, O. Montenbruck, Model of J_2 j_2 perturbed satellite relative motion with time-varying differential drag, Celest. Mech. Dyn. Astron. 123 (2015) 411–433, <http://dx.doi.org/10.1007/s10569-015-9643-2>.
- [36] A.W. Koenig, T. Guffanti, S. D'Amico, New state transition matrices for spacecraft relative motion in perturbed orbits, J. Guid. Control Dyn. 40 (2017) 1749–1768, <http://dx.doi.org/10.2514/1.g002409>.
- [37] G.V.M. Gaïas, C. Colombo, Analytical framework for precise relative motion in low Earth orbits, in: Proceedings of the 7th International Conference on Astrodynamics Tools and Techniques, ICATT, American Institute of Aeronautics and Astronautics (AIAA), 2018, pp. 1–10, <http://dx.doi.org/10.2514/1.g004716>.
- [38] G. Gaïas, C. Colombo, M. Lara, Analytical framework for precise relative motion in low Earth orbits, J. Guid. Control Dyn. 43 (2020) 915–927, <http://dx.doi.org/10.2514/1.g004716>.
- [39] D.-W. Gim, K.T. Alfriend, State transition matrix of relative motion for the perturbed noncircular reference orbit, J. Guid. Control Dyn. 26 (2003) 956–971, <http://dx.doi.org/10.2514/2.6924>.
- [40] S. D'Amico, Relative orbital elements as integration constants of hill's equations, Technical Report, German Aerospace Center (DLR), German Space Operations Center (GSOC), Oberpfaffenhofen, Germany, 2005, URL: https://www.researchgate.net/profile/Simone-Damico/publication/265495397_.pdf.
- [41] A. Scalvini, G. Borelli, G. Gaïas, Towards robust angles-only relative orbit navigation: Assessing key modelling factors using batch least squares filtering, in: Proceedings of the 11th European Conference for Aeronautics and Aerospace Sciences, EUCASS, 2025, <http://dx.doi.org/10.13009/EUCASS2025-555>.
- [42] R.H. Battin, An introduction to the mathematics and methods of astrodynamics, revised edition ed. in: AIAA Education Series, American Institute of Aeronautics and Astronautics, Reston, Virginia, USA, 1999, <http://dx.doi.org/10.2514/4.861543>.
- [43] O. Montenbruck, E. Gill, Satellite Orbits: Models, Methods, and Applications, first ed. vol. 55, Springer-Verlag, 2000, <http://dx.doi.org/10.1115/1.1451162>.
- [44] G. Gaïas, S. D'Amico, J.-S. Ardaens, Generalised multi-impulsive manoeuvres for optimum spacecraft rendezvous in near-circular orbit, Int. J. Space Sci. Eng. 3 (2015) 68–88, <http://dx.doi.org/10.1504/ijspacese.2015.069361>.
- [45] G. Gaïas, J.-S. Ardaens, C. Schultz, The avanti experiment: Flight results, in: Proceedings of the 10th International ESA Conference on Guidance, Navigation & Control Systems, European Space Agency (ESA), Salzburg, Austria, 2017, p. n/a, URL: <https://elib.dlr.de/112652/>, conference dates: 29 May–2 June 2017.
- [46] S. D'Amico, O. Montenbruck, Proximity operations of formation-flying spacecraft using an eccentricity/inclination vector separation, J. Guid. Control Dyn. 29 (2006) 554–563, <http://dx.doi.org/10.2514/1.15114>.