POLITECNICO MILANO 1863

SCUOLA DI INGEGNERIA INDUSTRIALE E DELL'INFORMAZIONE

Executive Summary of the Thesis

Efficient Satellite Manoeuvre Detection in Batch Orbit Determination Processes

Laurea Magistrale in Space Engineering - Ingegneria Spaziale

Author: Carlos Alonso Vélez

Advisor: Prof. Pierluigi Di Lizia

Co-advisor: Dr. Diego Escobar Antón, Pau Gago Padreny

Academic year: 2025-2026

1. Introduction

The rapid expansion of the space sector has resulted in a near-Earth space environment that is increasingly crowded with resident objects [1, 2]. This congestion poses significant challenges for satellite operations and the reliability of key services, highlighting the importance of space sustainability. To properly operate, a Space Traffic Management (STM) system requires accurate, comprehensive and up-to-date databases, which are created within the framework of Space Situational Awareness (SSA). A space object catalogue can be considered the foundation of space environment sustainability and is one of the primary outputs of Space Surveillance and Tracking (SST) activities.

However, when a satellite performs a station-keeping, orbital transfer or collision avoidance manoeuvre, standard propagation models fail. In these scenarios, stored orbits cannot be properly updated if a manoeuvre plan is not available. This failure is due to the fact that standard procedures fail to represent the physical reality, typically leading to an unsuccessful track-to-orbit (T2O) correlation or a degradation in the Orbit Determination (OD) process. To address this issue, the scope of the present thesis is the development of a novel and operationally feasible methodology to detect satellite manoeuvres from observational data.

2. Scope of the Thesis

A comprehensive review of the state of the art reveals a significant gap between the operational simplicity required for catalogue maintenance and the complexity of the available detection algorithms. While propagator-based methods [3, 4] offer low computational cost, their reliance on derived data limits their application to forensic analysis rather than precise, actionable operations. Conversely, filter-based approaches [5, 6] excel at custody maintenance but suffer from high computational burdens and the risk of divergence when facing unmodelled dynamics.

The methodology proposed in this thesis addresses these limitations by leveraging the infrastructure of a standard batch OD process. Specifically, the proposed algorithm functions as an extension of a batch Weighted Least-Squares (WLSQ) scheme. The primary objective of the method is to minimise the computational effort required to detect the occurrence of a manoeuvre, while providing an initial estimation of its magnitude and epoch. The method reuses computations from the nominal OD, namely the predicted Weighted Root Mean Square (WRMS) and partial derivative matrices. This procedure avoids the complexity of running parallel filter banks while maintaining the precision associated with high-fidelity perturbations models.

Furthermore, a distinct advantage of the developed approach over many correlation-driven techniques found in the literature [7, 8] is its flexibility regarding the manoeuvre epoch and the quality of the pre-manoeuvre orbit fit. The algorithm is designed to identify manoeuvres at a generic instant without relying on prior hypotheses regarding the manoeuvre epoch. While traditional optimisation-based algorithms often constrain the search space to the data gap between a catalogue state and a new track, the proposed strategy performs a grid search over the entire fitting interval. This is due to the fact that it does not assume the pre-manoeuvre orbit has been accurately estimated previously; in the present proposal, this segment of the trajectory is optimally adjusted as a part of the correction application, therefore allowing to utilize the degraded product of the standard OD process.

Regarding operational limitations, the method addresses the tracking problem rather than the surveillance problem. Consequently, it assumes the existence of a prior correlation between the incoming tracks and a catalogued object (orbit). Therefore, the algorithm intervenes in the OD and manoeuvre characterisation phase, rather than the initial association phase. The method developed herein aims to generate manoeuvre trend patterns that facilitate and enable future association processes, but it does not perform the association of Uncorrelated Tracks (UCTs) itself. Ultimately, the method offers a balanced compromise, providing the accuracy of high-fidelity dynamics with a computational cost compatible with routine operational constraints.

3. Methodology

3.1. Problem Formulation

Let $\boldsymbol{r}(t) \in \mathbb{R}^3$ and $\boldsymbol{v}(t) \in \mathbb{R}^3$ denote, respectively, the position and velocity vectors of a generic satellite orbiting the Earth at time t, in the GCRF frame. Both variables can be aggregated to define the state of the spacecraft at epoch t as $\boldsymbol{x}(t) = [\boldsymbol{r}(t), \boldsymbol{v}(t)]^\intercal \in \mathbb{R}^6$. A set of N observations $\{\boldsymbol{z}_i\}$ is obtained by a sensor network, modelled as a function of the state: $\hat{\boldsymbol{z}}_i = \boldsymbol{h}(\boldsymbol{x}(t_i)) \in \mathbb{R}^m$, where $\boldsymbol{h}$ is the measurement function and $\hat{\boldsymbol{z}}_i$ is the predicted observation. Due to (1) the inherent uncertainties associated with the measurement device (sensor noise) and (2) the knowledge of the real state of the object through the force model (process noise), a discrepancy exists between the real and the expected value of the observation taken at a certain instant, called observation residual and defined as: $\boldsymbol{\rho}_i = \boldsymbol{z}_i - \boldsymbol{h}(\boldsymbol{x}(t_i))$.

To account for the estimation of free dynamical parameters, such as the drag and solar radiation pressure coefficients, an extended state vector is defined as $\boldsymbol{y}(t) = [\boldsymbol{x}(t), \boldsymbol{p}]^\intercal \in \mathbb{R}^{6+n_p}$, where $\boldsymbol{p} \in \mathbb{R}^{n_p}$ contains the values of the n_p free parameters.

In a nominal scenario in which the satellite follows a ballistic trajectory, the state at a generic epoch t can be propagated, through the model $\boldsymbol{f}(\boldsymbol{x}, t)$, from a reference time t_0 at which the state is known, considering only natural perturbations. However, in the presence of a manoeuvre, these dynamics change. In the present work, propulsive events are assumed to be impulsive, consisting of an instantaneous change of velocity $\boldsymbol{u}$ occurring at an epoch t_M:

$$\boldsymbol{u} = \boldsymbol{v}(t_M^+) - \boldsymbol{v}(t_M^-) \in \mathbb{R}^3 \tag{1}$$

In this context, the trajectory over the whole observation interval is divided into two different ballistic orbits, Orbit A (pre-manoeuvre) and Orbit B (post-manoeuvre), which intersect at $\boldsymbol{r}(t_M)$ (see Figure 1). The objective of the OD process is to find the optimal set of parameters

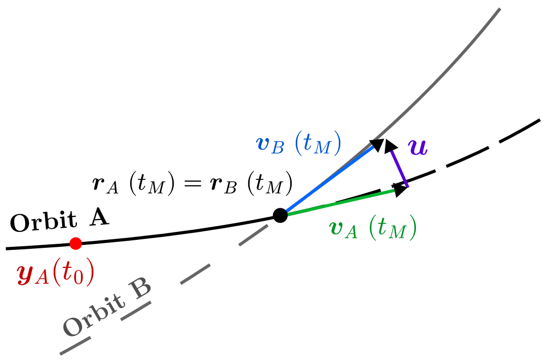

Figure 1: Impulsive manoeuvre model used for the developed procedure.

to fully characterise the trajectory: the extended state vector of Orbit A at a reference epoch t_0, $\boldsymbol{y}_A(t_0)$, and the velocity increment $\boldsymbol{u}$.

3.2. Model Linearisation

The core objective of the proposed approach is the minimisation of computational cost. For this reason, a linearised perturbation model based on a reference orbit defined by $\boldsymbol{y}_{ref}(t_0)$ is employed. The extended State Transition Matrix (STM), defined as $\boldsymbol{\Psi}(t,t_0) = \partial \boldsymbol{y}(t)/\partial \boldsymbol{y}(t_0)$, maps how small deviations from the reference affect the state at a generic time t.

Following the previous definition and using a first-order Taylor expansion with respect to the reference, the full state of the satellite in Orbit A can be approximated as:

$$\boldsymbol{y}_A(t) \approx \boldsymbol{y}_{ref}(t) + \boldsymbol{\Psi}(t,t_0)\,\Delta\boldsymbol{y} \tag{2}$$

where $\Delta\boldsymbol{y} = \boldsymbol{y}_A(t_0) - \boldsymbol{y}_{ref}(t_0)$ is the discrepancy between Orbit A and the reference trajectory at t_0.

On the other hand, substituting the relations of the partial derivatives and the instantaneous velocity change ($\boldsymbol{y}_B(t_M) = \boldsymbol{y}_A(t_M) + \left[\boldsymbol{0}_{3\times1}, \boldsymbol{u}, \boldsymbol{0}_{n_p\times1}\right]^{\mathsf{T}}$), the post-manoeuvre orbit is given by:

$$\begin{aligned} \boldsymbol{y}_B(t) &\approx \boldsymbol{y}_A(t) + \frac{\partial \boldsymbol{y}_A(t)}{\partial \boldsymbol{y}_A(t_M)}\left(\boldsymbol{y}_B(t_M) - \boldsymbol{y}_A(t_M)\right) \\ &= \boldsymbol{y}_A(t) + \boldsymbol{\Psi}_{yv}(t,t_M)\,\boldsymbol{u} \\ &= \boldsymbol{y}_{ref}(t) + \boldsymbol{\Psi}(t,t_0)\,\Delta\boldsymbol{y} + \boldsymbol{\Psi}_{yv}(t,t_M)\,\boldsymbol{u} \end{aligned} \tag{3}$$

with $\boldsymbol{\Psi}_{yv}(t,t_M)$ being the sub-matrix corresponding to the effect of a variation of velocity on the full state.

Consequently, the following linearisation allows to recover the observation residuals at any point of the trajectory as a function of the reference orbit and the (small) perturbations $\Delta\boldsymbol{y}$ and $\boldsymbol{u}$:

$$\begin{aligned} \boldsymbol{\rho}_i &\approx \boldsymbol{\rho}_{i,ref} - \frac{\partial \boldsymbol{h}(\boldsymbol{x}(t_i))}{\partial \boldsymbol{y}_A(t_0)}\Delta\boldsymbol{y} - \frac{\partial \boldsymbol{h}(\boldsymbol{x}(t_i))}{\partial \boldsymbol{u}}\boldsymbol{u} = \\ &= \boldsymbol{\rho}_{i,ref} - \boldsymbol{G}_{y,i}\,\Delta\boldsymbol{y} - \boldsymbol{G}_{u,i}\,\boldsymbol{u} \end{aligned} \tag{4}$$

Matrices $\boldsymbol{G}_{y,i} \in \mathbb{R}^{m\times(6+n_p)}$ and $\boldsymbol{G}_{u,i} \in \mathbb{R}^{m\times3}$ represent the effect on the expected observation at time t_i of the deviation of Orbit A (from the reference trajectory) and the manoeuvre, respectively. The former can be computed as $\boldsymbol{G}_{y,i} = \boldsymbol{H}_i\,\boldsymbol{\Psi}_{xy}(t_i,t_0)$, where the evaluation of the partials $\boldsymbol{H}_i = \partial \boldsymbol{h}(\boldsymbol{x}(t_i))/\partial \boldsymbol{x}(t_i) \in \mathbb{R}^{m\times6}$ is straightforward from the measurement model (often analytically). In the second case, the matrix is non-zero only for observations taken after the manoeuvre ($t_i > t_M$), and is then obtained using $\boldsymbol{G}_{u,i}(t_i) = \boldsymbol{H}_i\,\boldsymbol{\Phi}_{xv}(t_i,t_M)$.

3.3. Manoeuvre Detection Scheme and Differential Correction

The algorithm requires the prior execution of a batch Least-Squares (LS) estimation procedure, yielding the ballistic trajectory, characterised by $\boldsymbol{y}_{LS}(t_0)$, which best fits the provided observations. The quality of the fit is assessed via the Weighted Root Mean Square (WRMS), a metric defined as $\mathrm{WRMS} = \sqrt{1/N\sum_{i=1}^{N}\boldsymbol{\rho}_i^{\mathsf{T}}\boldsymbol{W}\boldsymbol{\rho}_i} \in \mathbb{R}$, being $\boldsymbol{W}$ the weighting matrix that optimally balances the contributions of the measurement types depending on their noise level. In the presence of a manoeuvre, the unmodelled dynamics cause the WRMS to grow substantially, triggering the detection scheme when it exceeds a certain threshold WRMS_{th}.

The baseline strategy of the detection approach is to find, for a fixed and hypothesised epoch t_M, the pre- and post-manoeuvre orbits A and B that best fit the observations in the LS sense. To this aim, the cost function $J = \mathrm{WRMS}^2$ (mean square of the residuals) must be minimised. Using the linearised residuals model given by Equation 4 results in an estimation problem where variables $\Delta\boldsymbol{y}$ and $\boldsymbol{u}$ allow to retrieve the optimal trajectory. Setting the first-order conditions with respect to both generates the so-called Gauss Normal Equations (GNE):

$$\sum_{i=1}^{N}(\boldsymbol{G}_i^{\mathsf{T}}\,\boldsymbol{W}\,\boldsymbol{G}_i)\begin{bmatrix}\Delta\boldsymbol{y}\\ \boldsymbol{u}\end{bmatrix} = \sum_{i=1}^{N}(\boldsymbol{G}_i^{\mathsf{T}}\,\boldsymbol{W}\,\boldsymbol{\rho}_{i,ref}) \tag{5}$$

where $\boldsymbol{G}_i = \left[\boldsymbol{G}_{y,i}\;\;\boldsymbol{G}_{u,i}\right]^{\mathsf{T}}$ is an aggregation of the partial derivative matrices of the i-th observation.

The solution of the previous system provides the optimal set $\{\Delta\boldsymbol{y}^*, \boldsymbol{u}^*\}$ which minimises the WRMS of the linearised residuals. Expressing the GNE as $\boldsymbol{N}\,\boldsymbol{X} = \boldsymbol{b}$, the reduced value of WRMS due to the inclusion of the manoeuvre in the estimation problem, under linear hypothesis, can be evaluated as follows:

$$\mathrm{WRMS}_{\mathrm{man}} = \sqrt{\mathrm{WRMS}_{LS}^2 - \frac{1}{N}\,\boldsymbol{b}^{\mathsf{T}}\,\boldsymbol{N}^{-1}\,\boldsymbol{b}} \tag{6}$$

with WRMS_{LS} being the weighted root mean square of the residuals $\boldsymbol{\rho}_{i,ref}$ obtained with the prior WLSQ procedure.

The main contribution to the efficiency of this method is the reuse of information that has already been obtained in the preceding batch process by setting the WLSQ product, $\boldsymbol{y}_{LS}(t_0)$, as the reference orbit for the differential correction. The predicted baseline WRMS_{LS}, as well as the partial derivative, transition, and sensitivity matrices used to construct $\boldsymbol{G}_i$, are already a product of the last iteration of the no-manoeuvre WLSQ. Consequently, the algorithm achieves manoeuvre characterisation without additional heavy orbit propagations; the only non-negligible cost is that associated with solving the linear system.

Furthermore, this framework assesses the effect that the introduction of a manoeuvre $\boldsymbol{u}^*$ and a correction on Orbit A $\Delta\boldsymbol{y}^*$ have on the residuals. This constitutes a distinct advantage over other approaches in the literature, which assume the pre-manoeuvre orbit estimate has high accuracy and maintain Orbit A fixed for the impulse determination. By optimally adjusting this segment of the trajectory alongside the manoeuvre, the method successfully utilises the degraded product of the standard OD process.

3.4. Selection of the Optimal Manoeuvre Epoch

To estimate the actual ignition epoch, a discretisation strategy is proposed, treating the manoeuvre time as a grid-search parameter instead of a continuous estimation variable in a non-linear optimisation problem, which often converges to local minima. The estimation is carried out by solving a sequence of linear LS problems for different values of t_M, and selecting the epoch t_M^* that results in the lowest augmented metric, $\mathrm{WRMS}_{\mathrm{man}}$.

To assess whether a real manoeuvre has taken place or not, the following residual relative improvement ratio is used:

$$\gamma\ [\%] = \frac{\mathrm{WRMS}^*_{\mathrm{man}}}{\mathrm{WRMS}_{LS}} \cdot 100 \tag{7}$$

If the obtained γ falls below a rigorously calibrated threshold γ_{th}, a manoeuvre at t_M^* is flagged. This is due to the fact that, in most common cases, manoeuvres have a larger influence on the residuals than other unmodelled effects (e.g., measurement noise, dynamic perturbations), meaning that, if the reduction of WRMS achieved by modelling a manoeuvre $\boldsymbol{u}^*$ at epoch t_M^* is substantial, an actual manoeuvre is likely to have been performed. Scenarios in which the manoeuvre magnitude is low are more difficult to tackle, as the distinction of its effect from the other sources is not evident. The selected value of γ_{th} needs to be high enough to capture low-magnitude operational manoeuvres which do not pose a significant degradation of residuals, but low enough to avoid detection when noise introduced by other sources, such as outliers, is substantial.

3.5. Algorithm Implementation

The practical implementation of the proposed scheme consists of a sequential execution, as observed in Figure 2.

1. A set of associated tracks is processed in a batch WLSQ scheme using an appropriate solver, including dynamic parameter estimation, covariance inflation, and relaxation factors for optimal correction calculation (theoretical foundations detailed in Montenbruck and Gill [9]).
2. The WRMS of the output estimate residuals is compared to a threshold WRMS_{th}; if exceeded, the detection scheme is triggered. Otherwise, the accuracy of the fit is deemed acceptable, and existence of a manoeuvre in the observation time interval is discarded.
3. The observation span $\mathcal{T}$ is discretised into n_t points. For each point, the optimal manoeuvre $\boldsymbol{u}^*$ and pre-manoeuvre orbit correction $\Delta\boldsymbol{y}^*$ are computed, together with the predicted WRMS. Because the transition and sensitivity matrices are available from the no-manoeuvre WLSQ, no additional orbit propagations are required to solve the GNE.
4. The epoch t_M^* providing the lowest WRMS is retrieved. If $\gamma < \gamma_{th}$, the manoeuvre is officially flagged.

This complete routine has been robustly implemented extending local GMV libraries and ESA's GODOT astrodynamics library [10], seamlessly processing both newly associated tracks and historical data with zero prior assumptions.

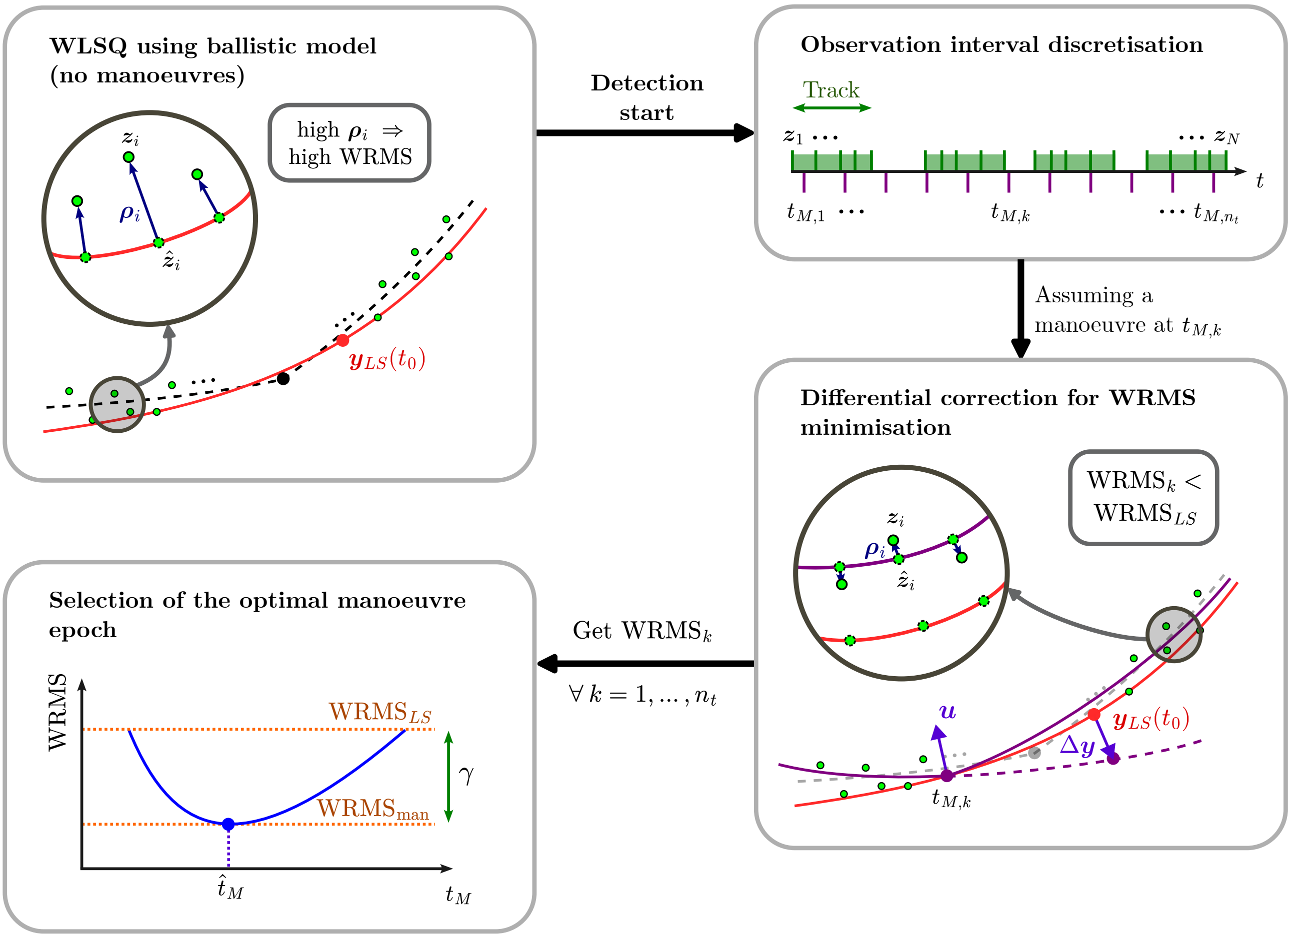

Figure 2: Full scheme of the developed procedure for a case containing a manoeuvre.

4. Results

4.1. Simulation Environment

To provide a rigorous evaluation, a comprehensive simulation campaign was designed and implemented. Unlike validation strategies based solely on limited real data scenarios, a Monte Carlo approach was adopted, which allows for the statistical assessment of the algorithm across a broad spectrum of orbital geometries and manoeuvre characteristics. The validation procedure emulates the complete operational loop of an SST operations centre, and is structured in sequential steps:

1. **Scenario generation**: a population of synthetic LEO orbits is generated by randomising Keplerian elements, covering the orbital regimes of interest, as shown in Table 1. Random impulsive manoeuvres, with a probability of occurrence of 70%, are injected into these trajectories. Angles α and β are the so-called in-plane and out-of-plane angles of manoeuvre $\boldsymbol{u}$ in the RTN (Radial-Tangential-Normal) reference frame. On the other hand, it is required to select a reference epoch t_0, which acts as the last catalogue update time, and a manoeuvre ignition epoch t_M. The start of the propagation interval is fixed to $t_{\mathrm{prop},0}$ = 2020-01-01T00:00:00 UTC, while its duration d is chosen so that every propagation interval contains roughly the same number of orbital revolutions, regardless of satellite altitude, which equates to 7 days for 700 km. Reference and manoeuvre epoch are randomly selected. After converting the Keplerian elements to a Cartesian state vector, all the synthetic satellite data is written to an appropriate OPM (Orbit Parameter Message) format file.
2. **High-fidelity propagation**: the synthetic states are propagated to generate "ground truth" ephemerides using local GMV software and stored in an OEM (Or-

Parameter	Distribution	Lower / Mean (μ)	Upper / Std Dev (σ)
Altitude (h)	Uniform	400 km	800 km
Eccentricity (e)	Uniform	0	10^{-3}
Inclination (i)	Uniform	30 deg	90 deg
Ω / ω / $\theta(t_0)$	Uniform	0 deg	360 deg
Δv	Gaussian	3 cm/s	2 cm/s
$\alpha = \angle(\hat{\boldsymbol{r}}, \hat{\boldsymbol{t}})$	Uniform	-45 deg	45 deg
$\beta = \angle(\hat{\boldsymbol{n}}, \hat{\boldsymbol{t}})$	Uniform	-30 deg	30 deg
Reference epoch (t_0)	Gaussian	$t_{\text{prop},0} + d/2$	$d/4$
Manoeuvre epoch (t_M)	Uniform	$t_{\text{prop},0}$	$t_{\text{prop},0} + d$

Table 1: Random generation parameters for orbital elements and epoch. The values correspond to Min/Max for Uniform distributions and Mean/Std Dev for Gaussian distributions.

bit Ephemerides Message) file.

3. **Measurement simulation**: realistic tracking radar data are synthesised, including visibility constraints and sensor noise, for four different devices, with a minimum elevation of 20 deg. The observations are then composed of azimuth, elevation, range and range-rate measurements. The resulting tracks are organized in TDM (Tracking Data Message) files.
4. **Nominal OD attempt (batch WLSQ)**: the simulated observations are processed using a model that excludes the manoeuvre and depends on the orbital regime of the particular object. The relative tolerance of the solver is set to 10^{-8} (in terms of WRMS), with a maximum of 4 diverging iterations. If the routine stops prematurely or does not meet the convergence criterion after 10 iterations, the algorithm is switched from Gauss-Newton (GN) to Levenberg-Marquardt (LM), in order to perform another 10 (maximum) iterations improving the probability of convergence.
5. **Manoeuvre detection**: in case the output WRMS of the WLSQ scheme exceeds threshold WRMS_{th}, the proposed algorithm is applied to classify the case (manoeuvre vs. no-manoeuvre) and, if detected, estimate the propulsive event.

This generation procedure was employed to create two distinct datasets: a primary training dataset consisting of 1,000 simulations and a secondary independent validation dataset of 700 simulations.

4.2. Detection Logic Calibration

Depending on the real scenario and the manoeuvre detection algorithm output flag, four different status categories are possible for a given simulation: True Positives (TP), True Negatives (TN), False Positives (FP) and False Negatives (FN). Based on these results, two accuracy metrics are mainly considered in the validation campaign: (1) precision, which quantifies the fraction of positive manoeuvre predictions that corresponds to a real event, and (2) recall, which measures the fraction of real manoeuvres that trigger the detection flag. To evaluate the overall performance of the classifier, the F_1 score is used, which combines precision and recall into a single scalar that weights their importance equally.

As explained, the proposed detection logic relies on two control parameters to distinguish between ballistic propagation and actual propulsive events: WRMS_{th} and γ_{th}. The calibration of these bounds is critical; an overly conservative selection may lead to missed events (FN), while an overly sensitive tuning could corrupt the catalogue with non-existent manoeuvres (FP). The decision space is visualised in Figure 3, which presents a scatter plot of the 1,000 Monte Carlo simulations. Qualitatively, a clear separability is observed: ballistic cases tend to cluster in the region of low initial residuals or high γ values (indicating no significant improvement in the fit when a manoeuvre is estimated), whereas manoeuvring targets typically exhibit high initial residuals ($\text{WRMS}_{LS} \gg 1$) that are drastically reduced by the algorithm ($\gamma \ll 100\%$). The value of WRMS_{th} sets the limit below which the case instances are automatically considered free of manoeuvres, while those that surpass it and experience a reduction $\gamma < \gamma_{th}$ are labelled as positives.

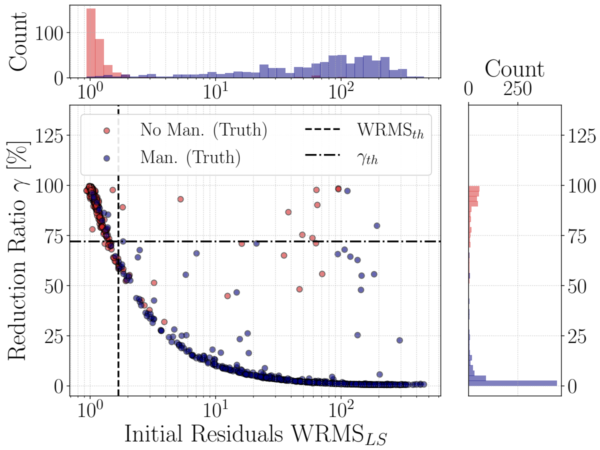

Figure 3: Decision map of the simulation campaign. The scatter plot shows the distribution of manoeuvring (red) and non-manoeuvring (blue) cases in the $\mathrm{WRMS} - \gamma$ plane.

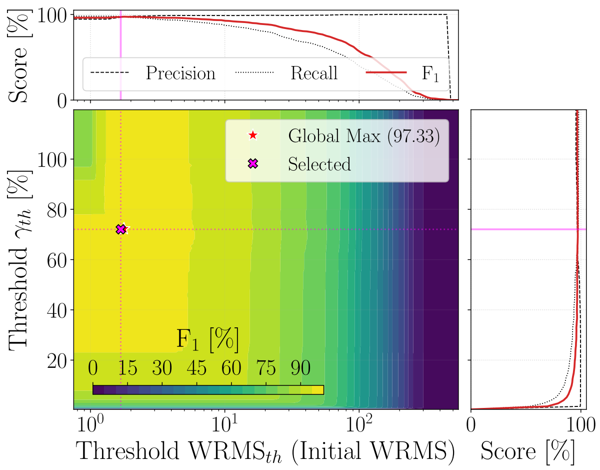

Figure 4: Optimisation landscape showing the F_1 score as a function of the detection thresholds. The global maximum is highlighted, with sensitivity cuts projected on the axes.

As depicted by Figure 4, a grid search optimisation was conducted over the dataset to identify the combination that maximises classification performance (F_1 score). To mitigate the risk of overfitting and verify the generalisation capability of the selected thresholds, a 5-fold cross-validation procedure was implemented. The final calibrated operational parameters were set to $\mathrm{WRMS}_{th} = 1.68$ and $\gamma_{th} = 72.04\%$, representing the best compromise for the simulated environment.

4.3. Statistical Assessment

The statistical validation of the proposed methodology is conducted using the aforementioned global operational thresholds applied to the primary dataset. The raw classification outcomes are visually summarised in the confusion matrix depicted in Figure 5, which contrasts the predicted labels against the ground truth of the synthetic scenarios. As observed in the matrix, the vast majority of cases accumulate along the main diagonal, indicating a high degree of success in both identifying propulsive events and rejecting ballistic trajectories, with misclassifications representing a marginal fraction of the sample space.

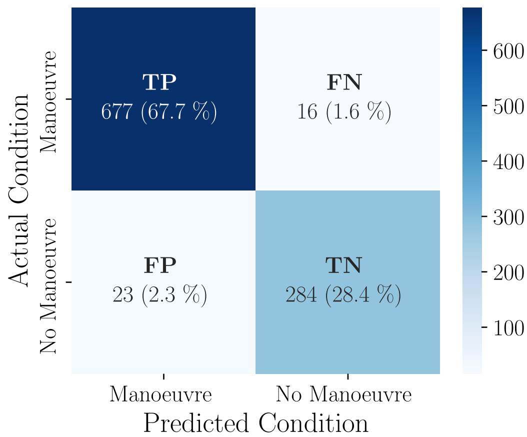

Figure 5: Classification confusion matrix. The grid presents both the absolute counts and the relative percentages of TPs, FPs, TNs and FNs obtained during the validation campaign.

To provide a rigorous quantitative assessment, these raw counts are translated into the standard performance metrics. The system achieves a precision of 96.71%, which ensures that the catalogue integrity is preserved by minimising the inclusion of non-existent manoeuvres. Simultaneously, the recall stands at 97.69%, demonstrating the algorithm's sensitivity to recover the vast majority of performed manoeuvres. Consequently, the overall performance, measured by the F_1 score, reaches a value of 97.20%, validating the algorithmic strategy as a balanced solution for the detection problem.

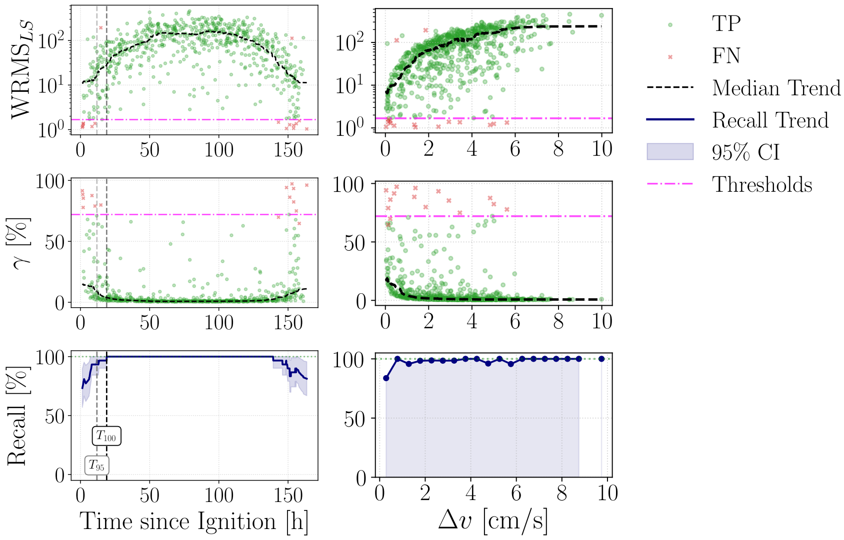

Figure 6: Sensitivity analysis of the detection performance with respect to detection latency $(t - t_M)$ and manoeuvre magnitude (Δv).

Sensitivity analyses indicate that the primary limiting factor for detection is observation geometry, i.e., the temporal distribution of measurements relative to the ignition epoch. As shown in the left column of Figure 6, when a manoeuvre is recent $(t - t_M \to 0)$, the orbital deviation is not yet fully developed, resulting in low residuals that are difficult to distinguish from noise; conversely, at the end of the interval, the lack of pre-manoeuvre history allows the batch estimator to absorb the propulsive event into a distorted ballistic arc, lowering the post-manoeuvre residuals and making the fraction of observations with high WRMS substantially small. Consequently, in some of these cases, the introduction of a manoeuvre into the dynamic model does not significantly improve the fit (higher γ). From an operational perspective, the lack of pre-manoeuvre data is a non-critical state; in a rolling catalogue update cycle, a recent manoeuvre (few post-manoeuvre tracks) will progressively accumulate observations, naturally moving into the region of high observability. The bottom row characterizes the detection probability (recall) using a rolling mean, which demonstrates a rapid saturation, achieving 90% in less than 8 hours after the ignition, and 100% (full reliability) before 19 hours. This latency may vary for different sensor network configurations, as the presence of four stations might not always be the case. In any case, the recall remains consistently above 75% throughout the entire observable domain, confirming the robustness of the proposed method even in early detection scenarios.

The final stage of the sensitivity assessment investigates the influence of the intrinsic properties of the propulsive event on the detection rate. The analysis of the second column of Figure 6 confirms the physical intuition governing the estimation process: the nominal residuals exhibit a clear upward trend proportional to the magnitude of the impulse, as larger Δv values introduce greater deviations from the ballistic trajectory. Furthermore, the reduction ratio γ improves (decreases) for larger events. In terms of reliability, the recall remains consistently high across the entire spectrum, experiencing a performance drop to approximately 80% only for small manoeuvres ($\Delta v < 1$ cm/s), where the induced orbital perturbation becomes comparable to the measurement noise floor and propagation errors. In contrast, the geometric orientation of the impulse does not appear to be a critical driver of the system's performance.

While the fundamental purpose of the proposed methodology is the binary classification of trajectories (detection), the algorithm produces a preliminary estimate of the manoeuvre parameters as a byproduct of the residual minimization. However, the raw output of the solver can occasionally produce physically implausible values (e.g., extremely large velocity increments) when numerical convergence is poor, even if the detection flag is correct. In an operational environment, these estimation instances (not their detection labels) would be automatically discarded by an appropriate filter based on typical Δv magnitudes. To simulate this process and derive representative accuracy metrics, an iterative σ-clipping filter [11] was applied to the estimated Δv of the TP cases, obtaining a cut-off threshold of 14 cm/s and rejecting 21 outliers (3% of TPs). The statistical profile of the estimation residuals is summarised in Table 2. The 90^{th} percentile (P_{90}) establishes a reliable operational upper bound, certifying that 90% of the accepted estimations exhibit an epoch error below 15% of the orbital period and a magnitude error of roughly 50% of the mean value (3 cm/s). These results support the suitability of operationally implementing the threshold for accepting estimations since, once the clear outliers are filtered out, the vast majority of detections possess sufficient precision to successfully seed a subsequent, more refined manoeuvre estimation process.

Metric	Mean	Median	σ	P_{90}
Abs. epoch error [rev]	0.321	0.057	2.79	0.147
Norm. abs. Δv error [%]	23.9	13.0	35.8	51.1

Table 2: Statistical summary of the estimation errors for TP detections.

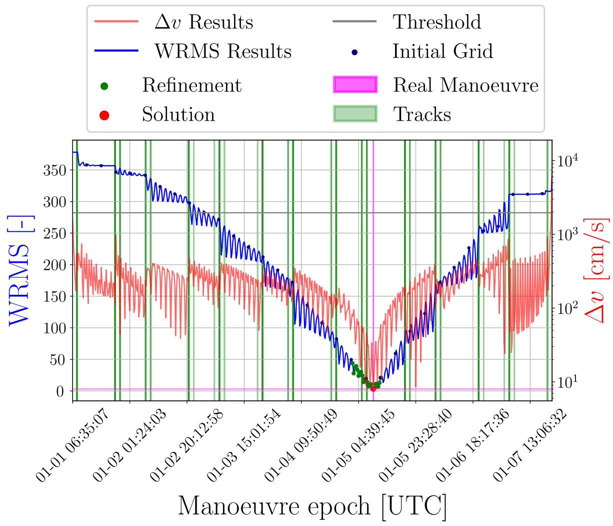

Figure 7: Global landscape of a manoeuvre estimation objective function. The plot illustrates the evolution of the post-fit residuals (WRMS) and the estimated Δv across the entire propagation interval. The markers indicate the sampling points of the M4 strategy.

4.4. Cost Optimization

Computational profiling identified the manoeuvre detection stage as a significant contributor to the total runtime. However, an analysis of the residual evolution function (WRMS vs. t_M, see Figure 7) reveals that the manoeuvre fitting landscape is generally convex over the propagation arc, superimposed with local oscillations proportional to the orbital period. This topological characteristic suggests that the initial dense grid single-sweep approach is unnecessarily conservative. Consequently, a hierarchical "coarse-to-fine" search strategy is designed, as shown in Table 3, and validated to drastically reduce execution time without compromising classification accuracy. The most aggressive one, M4, achieved a mean runtime of 11.1 s (less than 2% of the total OD time: WLSQ + detection), representing an approximate 10-fold reduction compared to the baseline.

This analysis is executed on the secondary dataset of 700 simulations, whose blind validation yields a stable F_1 score of 96.18%, confirming the algorithm's generalisation capability.

Strategy	Initial Grid	Refined Grid	Mean Cost	Mean Contribution [%]
Baseline (M1)	6	–	1 min 53 s	15.45
Conservative (M2)	1	6	21.1 s	2.88
Intermediate (M3)	1/3	6	13.3 s	1.82
Aggressive (M4)	1/3	4	11.1 s	1.52

Table 3: Search grid configurations evaluated for operational optimization, with grid densities in points per orbit.

4.5. Impact of Miscorrelated Tracks

In an operational SST catalogue maintenance environment, the association of tracks to catalogue objects (T2O) is not error-free. Occasionally, a track belonging to a different object may be incorrectly correlated to the satellite under analysis (cross-tagging). It is critical that the detection algorithm does not mistake these severe outliers for a propulsive event, generating FP instances. To assess this vulnerability, a specific test scenario is created, in which a single TDM file belonging to a different satellite from the database is injected into the nominal ballistic trajectory observation set, acting as an "intruder" track.

The first analysis focuses on the correlation between the occurrence of false alarms and the temporal location of the miscorrelated track relative to the end of the propagation interval (representing the current epoch). The results show that when the defective track is located near the end of the estimation arc (<18 hours), the solver is often able to find a manoeuvre solution that significantly reduces the residuals (low γ), thereby triggering a FP. Physically, this occurs because a recent deviation is indistinguishable from a post-manoeuvre trajectory branch. Consequently, filtering detections solely based on γ implies accepting a number of FPs. However, the solver, in its attempt to fit the uncorrelated outlier tracks (FP cases), converges to impulsive solutions with no physical meaning, typically yielding Δv values in the order of km/s.

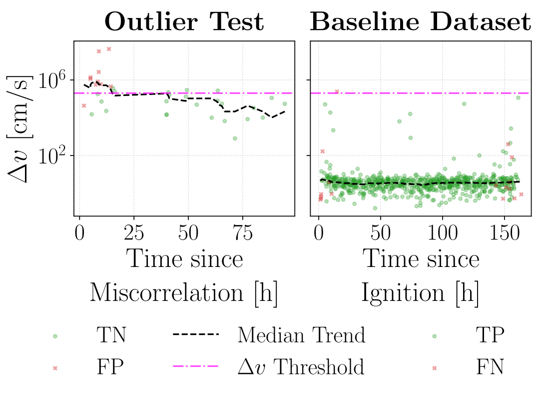

Figure 8: Separability analysis of estimated manoeuvre magnitudes. The plot contrasts the Δv values obtained from miscorrelated tracks (outliers) analysis against the population of real manoeuvres (TP) in the baseline dataset.

Figure 8 highlights that it is operationally possible to discriminate the artifacts (FPs in the left column) using a threshold based on this estimated Δv, e.g., 2 km/s, without affecting the valid detections (TPs in the right column), and thus maintaining the same recall rate.

5. Conclusions

This thesis has presented a novel methodology for the detection of satellite manoeuvres within the context of operational SST. Unlike other approaches that frame manoeuvre detection as a surveillance or uncorrelated track association problem, the proposed algorithm addresses the issue from the tracking domain. It leverages the infrastructure of a standard batch OD process to identify, characterise, and correct for non-cooperative propulsive events without disrupting the operational flow.

The numerical validation, performed via Monte Carlo simulations involving synthetic LEO scenarios, demonstrated the robustness of the approach, leading to the following conclusions:

- **Operational feasibility and efficiency**: the method meets the strict requirement of low computational cost ($<2\%$ of OD time). By embedding the detection logic within the final steps of the batch process and recycling propagation data, transition and sensitivity matrices, and the predicted WRMS of residuals, the algorithm enables the systematic processing of large catalogues without necessitating complex routines.
- **Detection performance and robustness**: the algorithm exhibited a high recall rate and precision, reflected in an F_1 score exceeding 96%. Furthermore, false positives caused by misassociated tracks tend to result in non-physical, high-magnitude impulsive solutions ($\Delta v > 1$ km/s), allowing for the implementation of simple magnitude-based filters to effectively reject artifacts without compromising the recall of legitimate manoeuvres.
- **Estimation capability and catalogue enrichment**: while the primary goal is detection, the method provides a preliminary estimation of the manoeuvre's epoch and magnitude that can serve as an initial guess for subsequent refinement algorithms, provided physically implausible es-

timates are rejected using an appropriate filter, e.g., based on the estimated Δv impulse. Ultimately, generating and maintaining a chronological archive of propulsive events facilitates future T2O association tasks.

Despite the promising results, the methodology operates under specific assumptions that define its current operational boundary. The most significant limitation is its reliance on a successful prior correlation. For instance, if a manoeuvre is large enough to completely decorrelate the track from its catalogue entry, the problem reverts to the uncorrelated track (UCT) domain. Consequently, this method serves as an "orbit maintenance" tool rather than a "lost object discovery" tool.

The development and validation of this algorithm open several avenues for future research. Algorithmic enhancements could include further studying grid search refinement techniques and improving the WLSQ solver efficiency to further reduce execution time, analysing convergence criteria and fallback mechanisms. Furthermore, to ensure broader applicability, future validation campaigns should investigate the impact of uncalibrated sensor biases, test performance within the Geostationary (GEO) regime, and evaluate the system against extended manoeuvre profiles to determine the absolute bounds of detectability.

References

[1] ESA Space Debris Office. ESA Space Environment Report 2025. LOG GEN-DB-LOG-00288-OPS-SD Rev. 9.1, European Space Agency, October 2025.

[2] Phillip Anz-Meador, John Opiela, and Jer-Chyi Liou. History of On-orbit Satellite Fragmentations. Technical Report NASA/TP-20220019160, 16th Edition, National Aeronautics and Space Administration, January 2023.

[3] T. Kelecy, D. Hall, K. Hamada, and D. Stocker. Satellite Maneuver Detection Using Two-line Elements Data. In *Advanced Maui Optical and Space Surveillance Technologies Conference*, page E19, January 2007.

[4] M. Pittelkau. Space Object Maneuver Detection Algorithms Using TLE Data. In *Advanced Maui Optical and Space Surveillance Technologies Conference*, page 105, September 2016.

[5] Gary M. Goff, Jonathan T. Black, and Joseph A. Beck. Tracking maneuvering spacecraft with filter-through approaches using interacting multiple models. *Acta Astronautica*, 114:152–163, September 2015. ISSN 0094-5765. doi: 10.1016/j.actaastro.2015.05.009.

[6] Christoph M. Bergmann, Andrea Zollo, Johannes Herzog, and Thomas Schildknecht. Integrated manoeuvre detection and estimation using nonlinear Kalman filters during orbit determination of satellites. In *8th European Conference on Space Debris*, 2021.

[7] Alejandro Pastor, Guillermo Escribano, Manuel Sanjurjo-Rivo, and Diego Escobar. Satellite maneuver detection and estimation with optical survey observations. *The Journal of the Astronautical Sciences*, 69 (3):879–917, June 2022. ISSN 2195-0571. doi: 10.1007/s40295-022-00311-5.

[8] Lorenzo Porcelli, Alejandro Pastor, Alejandro Cano, Guillermo Escribano, Manuel Sanjurjo-Rivo, Diego Escobar, and Pierluigi Di Lizia. Satellite maneuver detection and estimation with radar survey observations. *Acta Astronautica*, 201:274–287, December 2022. ISSN 0094-5765. doi: 10.1016/j.actaastro.2022.08.021.

[9] Oliver Montenbruck and Eberhard Gill. *Satellite Orbits*. Springer, Berlin, Heidelberg, 2000. ISBN 978-3-540-67280-7 978-3-642-58351-3. doi: 10.1007/978-3-642-58351-3.

[10] ESA Flight Dynamics. GODOT 1.14.1 documentation. https://godot.io.esa.int/godotpy/, 2022.

[11] Mohammad Akhlaghi. GNU Astronomy Utilities. *Free Software Foundation*, July 2024. doi: 10.5281/zenodo.12738457.