ALGORITHMS FOR ROBUST TRACKING OF MANOEUVRING SPACE OBJECTS IN CATALOGUE MAINTENANCE

Álvaro Flórez González(1), José Ramón García Espinosa(1), Ivan Llamas de la Sierra(1), Emanuele Nasini(2), Pierluigi Di Lizia(3), Diego Aranda Pradana(2), Pau Gago Padreny(2), Antonio Malaver(4), Rafael Vázquez(4), Felix Stechowsky(5), Juan Velasco(5), and Alejandro Pastor(5)

(1)GMV, Rue Hermès 17, 31520 Ramonville St. Agne, France, Email: aflorez@gmv.com, jrgarcia@gmv.com
(2)GMV, Calle Isaac Newton 11, 28670 Tres Cantos, Spain, Email: enasini@gmv.com, daranda@gmv.com
(3)Politecnico di Milano, Via Privata Giuseppe La Masa 34, 20156 Milan, Italy, Email: pierluigi.dilizia@polimi.it
(4)Universidad de Sevilla, Escuela Técnica Superior de Ingeniería, Camino de los Descubrimientos s.n, 41092 Sevilla, Spain, Email: antmaljam@alum.us.es, rvazquez1@us.es
(5)GMV, Europaplatz 2, 64293 Darmstadt, Germany, Email: felix.stechowsky@gmv.com, juan.velasco.mariscal@gmv.com, apastor@gmv.com

ABSTRACT

More than one third of the trackable population of space objects is classified as operational payload. Most of them are satellites and spacecraft with manoeuvring capabilities. The proliferation of satellite constellations has increased and will continue to escalate the complexity and volume of space traffic, urging the need for solutions for cataloguing. This paper presents several innovative algorithms designed for an automated, efficient and robust tracking of manoeuvring space objects during catalogue maintenance. For cases where correlated observations are available (i.e., already associated with a specific space object), we introduce a manoeuvre detection strategy that supports optical, radar, laser, and passive ranging sensors. Additionally, we propose a dynamics-agnostic alternative aimed at maintaining the trackability of the space object by characterising the uncertainty introduced by the manoeuvre using covariance inflation and smoothing within sequential estimation. Finally, we address the challenging issue of uncorrelated observations by proposing a multiple hypothesis tracking-like methodology. The different methodologies are presented comprehensively, along with relevant test cases that assess the performance and suitability of the different approaches. The results demonstrate the potential of these algorithms to significantly improve the accuracy and reliability of the catalogue of space objects.

Keywords: Manoeuvre Detection; Manoeuvre Estimation; Catalogue Maintenance; Tracking.

1. INTRODUCTION

At the current time of space exploration, Space Traffic Management (STM) and Space Situational Awareness (SSA) have become a crucial issue to ensure the safety of all stakeholders. Congestion in the most densely populated Earth orbits poses a challenge to their exploitation. The latest reports estimate a population of 40500 objects larger than 10 cm, usually considered trackable population [1] and this situation is expected to aggravate in the coming years with more launches and the generation of new space debris.

In this context, one of the main strategies in the Space Surveillance and Tracking field (SST) is the creation and maintenance of space object catalogues. It allows to know and predict the status of the population of objects orbiting Earth. A space objects catalogue is defined as a robust, automated and reliable database containing the information of the detected space objects. It has to be built and maintained through a series of data processing techniques, known as the cataloguing chain. The main inputs to the catalogue are the observations from a sensor network, which are grouped into tracks or tracklets, i.e. batch of observations belonging to a unique object generated when it passes over a sensor. From these, it is possible to establish an orbit for each space object, which will be updated and maintained over time as more information is received. The cataloguing process includes techniques to correlate the generated tracks to existing objects (track-to-orbit correlation), include new space object to the catalogue and estimate an initial orbit (track-to-track association), and avoid duplication of objects (orbit-to-orbit correlation), among other applications.

However, this process is hindered by the presence of manoeuvrable objects. In these cases, the trajectory cannot be properly maintained without a priori knowledge of the manoeuvre plan. Unfortunately, this information is rarely published by most operators, which hinders orbit update. In current catalogues, for example the one maintained by Kelso [2], the fraction of objects listed as 'active' is on the order of 33%. In fact, it is estimated that this situation will become more and more concerning in the com-

Proc. 9th European Conference on Space Debris, Bonn, Germany, 1–4 April 2025, published by the ESA Space Debris Office
Editors: S. Lemmens, T. Flohrer & F. Schmitz, (http://conference.sdo.esoc.esa.int, April 2025)

ing years, since a great deal of scheduled launches correspond to large constellations such as Starlink, Oneweb and Quianfan.

The main issues regarding manoeuvrable objects can be sorted into two kind: survey and tracking problems. Survey problems arises when the tracks cannot be correlated to any of the objects in the catalogue. The usual strategy to correlate tracks is to compute the residuals between the actual observations and the predicetd ones, i.e. the difference between the computed measurements coming from the orbit prediction and the actual ones. However, if a manoeuvre happens between the last available measurement and the target track date, the correlation process may fail. The predicted trajectory will not take into consideration the manoeuvre information and it will not match the observations. Even more, this situation could lead to generation of duplicated objects during catalogue maintenance. Provided that the track will not be assigned to any of the existing objects, catalogue maintainer could try to conduct track-to-track association strategies to estimate the state of a possible new object. Then, if they succeed, the catalogue could be extended with an RSO which is in fact a duplication of an existing one.

On the other hand, tracking problem might occur when the track has been pre-correlated, so the involved object is known and thus there is no need to perform a correlation process. However, catalogue maintainers might not be able to estimate an accurate orbit. Most of the current Orbit Determination (OD) processes for SST rely on Batch Least Squares solutions [3]. These algorithms use as initial information a set of measurements associated to a specific object and an initial estimation of the state vector at a certain epoch. The aim is to fit the parameters of the considered dynamical model to the observations, by minimizing the root mean square of the residuals. These are also usually weighted with the expected sensor measurement noise. In this process, the tracks are treated in groups, so if the actual manoeuvres in the observation time span have not been taken into account, the dynamical model may not be able to reproduce the trajectory and the orbit determination may fail. One possible way to mitigate this is to reduce the orbit determination interval to one without manoeuvres. Nevertheless, this will result in less accurate orbits, since the number of observations used to generate them will be much smaller. In certain cases, it could even lead to custody loss of those objects.

This paper presents three innovative methodologies to address both the tracking and survey problems. Firstly, we consider the case in which correlated observations are available but the manoeuvre plan of the target object is unknown. A manoeuvre detection and estimation algorithm has been developed to cover the cases of radar, laser and passive ranging sensors. The starting point for this research can be found in [4] and [5]. Manoeuvre detection is done by trend analysis of the residuals between the actual observations and the predicted orbit, weighted by the reported sensors noises. The algorithm is capable of inferring a time interval to search for the manoeuvre based on thresholds. Then a linearized propagation model based on the use of State Transition Matrix allows to estimate the manoeuvre date and magnitude by a weighted non-linear least squares process that tries to fit the post-manoeuvre tracks. The methodology in this paper is an extension of the latter, in which a multiple burns manoeuvre scenario is considered. The previous propagator is extended, taking into account the linearized effect of successive impulsive manoeuvres and a cost function function is defined to minimize both the residuals with respect to the observations and the expected manoeuvre magnitudes.

Furthermore, a dynamics-agnostic solution for the tracking problem is proposed, aimed at maintaining the trackability of the space object by characterising the uncertainty introduced by the manoeuvre. In this case we rely on sequential estimators (e.g., Extended Kalman Filter) [3] to process the observations continuously. Manoeuvre detection in this case explores, among others, the Mahalanobis distance [6] of track attributables, using the predicted orbits and covariances [8]. An attributable is defined as an observation compression consisting in the transformation of a track into a single observation [7]. One of the main issues with sequential estimators is covariance shrinkage with each new estimation. This strategy intends to mitigate this problem by covariance inflation. Once a manoeuvre is detected, the covariance matrix is artificially increased in an attempt to capture the uncertainty increase due to this discontinuity. Since the orbit uncertainty would have risen, the filter is able to assimilate the new observations coming from post manoeuvre tracks. In addition, a smoothing step with several forward and backwards estimations is applied to achieve the most refined solution for the orbit estimation, until the algorithm converges. This may prevent us from inferring the manoeuvre characteristics, which is not required for catalogue maintenance purposes, but allow us to keep on tracking the object.

Finally, we address the survey problem by proposing a multiple hypothesis tracking-like methodology. This approach uses track-to-track association techniques [9] [10] to group uncorrelated tracks, retrieved by the sensor network. These clustered observations are considered as post-manoeuvre tracks. Thus, it is possible to try to estimate a manoeuvre that allows to fit these measurements to the orbits present in the catalogue prior to the reception of the observations. The algorithm used for the manoeuvre estimation is based again on the same linearized method presented for the tracking problem previously mentioned. This results in multiple solutions, which allow to reach the desired orbit from different candidate orbits. A pruning and ranking process based on the weighted root mean square of the residuals it and the magnitude of the estimated manoeuvre is used to promote the final solution.

This paper is structured in 5 different sections: Section 1 presents the challenges related to catalogue maintenance in the presence of manoeuvring objects, summarizes the motivation behind this research and introduces the proposed solutions. Section 2 and Section 3 describe the developed methods targeting the tracking problem and outlines use cases for each of them. Section 4 proposes a solution for the survey problem, with the correspondent description and use case. Finally, Section 5 compiles the conclusions and future work derived from these approaches.

2. MULTIPLE BURNS MANOEUVRE ESTIMATION

This work focuses on the development of a new methodology to perform manoeuvre determination for satellites in LEO, starting from a set of ground station measurements acquired at different epochs, and based on trajectory optimization using a cost function.

2.1. Methodology

In a multi-manoeuvre scenario, the orbit is divided into several arcs, with manoeuvres occurring at specific times. A Linear Multi-Manoeuvre Orbit Propagator method (LMMOP), based on the use of State Transition Matrix (STM) has been developed to properly simulate this behaviour in an operational environment. Through the STM the effects of single manoeuvres are added consecutively to the reference orbit. Generalizing the *enhanced STM correction method* of Porcelli [5], tailoring it for a multi-manoeuvre scenario, Eq. 1 can be defined for the correction of the state of a space object after a number of manoeuvres, at time $\hat{t}$:

$$\begin{aligned}\mathbf{x}(\hat{t}) = \mathbf{x}_A(\hat{t}) + \sum_i^{n_{\text{man}}} \Big\{ &\Phi_A(t_i, \hat{t})\delta\mathbf{x}_i + \\ &+ R_i \Big[\mathbf{x}_{K,post}(\hat{t}) - \mathbf{x}_{K,pre}(\hat{t}) + \\ &- \Phi_{K,pre}(\mathbf{x}_{K,pre}(\hat{t}), t_i, \hat{t})\delta\mathbf{x}_i\Big]\Big\}\end{aligned} \tag{1}$$

being $\mathbf{x}_A$ the reference pre-manoeuvre state, n_{man} the total number of manoeuvres, t_i the manoeuvring times, Φ_A the STM of the reference orbit, and $\delta\mathbf{x}_i$ the manoeuvre vector at time t_i, considered as a perturbation. R_i is the keplerian correction rotation matrix, while $\mathbf{x}_{K,post}$, $\mathbf{x}_{K,pre}$ and $\Phi_{K,pre}$ are the keplerian correction terms (see [5]), with the subscripts *post* and *pre* referring to the post-manoeuvre and pre-manoeuvre keplerian orbit arcs, corresponding to the considered $\mathbf{u}_i$, array holding the components for manoeuvre i.

An important assumption is made when two or more manoeuvres are considered: that the post manoeuvre arc is close to the reference orbit, since the used STMs are always the ones of the initial reference orbit. The effect of manoeuvre, starting from the second one onwards, are computed as they were applied to the reference orbit, while in reality they are acting on different orbit arcs.

As it can be seen from Fig. 1 the real correction would be $\Delta\mathbf{x}_1 + \Delta\mathbf{x}_2$, which are the increment in the state vector

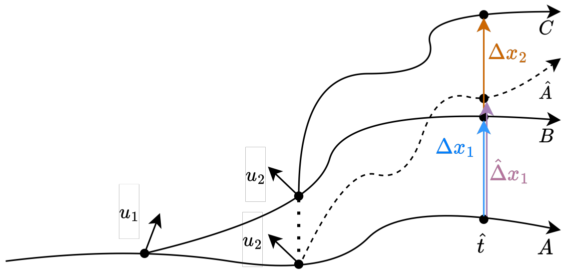

Figure 1. STM approximation.

due to the effect of u_1 and u_2. Instead of using always the STM of the reference orbit, it is approximated as $\Delta\mathbf{x}_1 + \hat{\Delta}\mathbf{x}_1$, with both manoeuvre acting on orbit A.

The effect of this approximation can be mitigated performing a rotation of the STM before computing the correction. Indeed after each manoeuvre the orbit changes its orientation in space, by an angle α, and the STM can be rotated by the same angle, re-aligning it with the correct orbital dynamics. To realign the STM with each new orbit arc, a rotation matrix $\mathbf{\Gamma}$ is used:

$$\Gamma(t) = \left(R^{\text{corrected}}_{\text{TNW}}(t)\right)^T R^{\text{reference}}_{\text{TNW}}(t) \tag{2}$$

being $R^{\text{corrected}}_{\text{TNW}}(t)$ and $R^{\text{reference}}_{\text{TNW}}(t)$ the rotation matrices from the Cartesian frame to the local TNW frames associated with the current corrected orbit arc (after the manoeuvres) and the reference orbit.

The STM is therefore rotated as follow:

$$\hat{\Phi}_A(t_i, t) = \Gamma(t)\Phi_A(t_i, t)\Gamma(t)^T \tag{3}$$

and Eq. 1 can be reformulated accordingly:

$$\begin{aligned}\mathbf{x}(\hat{t}) = \mathbf{x}_A(\hat{t}) + \sum_i^{n_{\text{man}}} \Big\{ &\hat{\Phi}_A(t_i, \hat{t})\delta\mathbf{x}_i + \\ &+ R_i \Big[\mathbf{x}_{K,post}(\hat{t}) - \mathbf{x}_{K,pre}(\hat{t}) + \\ &- \Phi_{K,pre}(\mathbf{x}_{K,pre}(\hat{t}), t_i, \hat{t})\delta\mathbf{x}_i\Big]\Big\}\end{aligned} \tag{4}$$

A genetic algorithm, the *Differential Evolution* (DE) algorithm, has been chosen for the optimization routine. The chosen logic for the optimizer is a Global Optimization strategy. This approach considers all the measurements at the same time, processing them in parallel. The founding assumption for this is that at maximum one manoeuvre per day is performed, as this is typical for Starlink satellites [11]. Measurement passes are therefore grouped accordingly, starting from the first pass where a manoeuvre is detected, with intervals of one day, and a manoeuvre is assumed to be present in each of them.

At each iteration of the optimization a guess on the manoeuvres magnitudes and epochs will be produced. This guess is used to propagate the orbit with the LMMOP, to then compute the simulated measurements for the current guess (without noise). Those simulated measurements will be compared to the real ones and residuals will be computed.

The selected optimization variables are, for each manoeuvre, its magnitude and its epoch. Manoeuvres are considered always in the along track direction, neglecting the radial and out of plane components. By doing so the total number of optimization variables is twice the number of manoeuvres ($2 \cdot n_{\text{man}}$).

Parameters entering the cost function are the Root Mean Square Errors (RMS) of the measurement residuals ρ (one for each group of passes, for each day), together with the RMS of the manoeuvre magnitudes. All the RMS are scaled through the adimensionalization of the measurement residuals and the manoeuvre magnitude to standardize their influence (Weighted Root Mean Square, WRMS), allowing the optimizer to explore the landscape more evenly and converge faster. The scaling parameters are the noise values for the ground stations together with the expected manoeuvre magnitude (Δv) value (Tab. 1), such that at the optimum point the cost function will be unitary. Manoeuvre magnitude is extracted from [11].

Table 1. Scaling factors for the adimensionalization. Values for range (R), range rate ($\dot{R}$), azimuth (Az), elevation (El) and manoeuvre magnitudes

σ_R	$\sigma_{\dot{R}}$	σ_{Az}	σ_{El}	$\sigma_{\Delta v}$
10 m	600 mm/s	0.3 deg	0.3 deg	10 mm/s

After adimensionalizing the residuals, the WRMS for each group of measurement can be computed:

$$WRMS_{\text{meas},j} = \sqrt{\frac{\sum_i^{n_{\text{meas},j}} \left[\left(\frac{\rho_{\text{meas}}}{\sigma_{\text{meas}}} \right)^2 \right]}{4 n_{\text{meas},j}}} \tag{5}$$

With $n_{\text{meas},j}$ being the number of observations in the considered group of passes. The scaled WRMS for the manoeuvre magnitudes can be computed similarly:

$$WRMS_{\Delta v} = \sqrt{\frac{\sum_i^{n_{\text{man}}} \left(\frac{\Delta v_i}{\sigma_{\Delta v}} \right)^2}{n_{\text{man}}}} \tag{6}$$

Since at the optimum point the measurement residuals will be around the noise values, and the manoeuvre magnitudes will be the expected ones, the scaled WRMSs will all be unitary. They finally enter the cost function, J, which is divided by the total number of entries, such that it remains unitary at optimum:

$$J = \frac{\sum_j^{n_{\text{days}}} (WRMS_{\text{meas},j}) + WRMS_{\Delta v}}{n_{\text{days}} + 1} \tag{7}$$

With n_{days} being the number of days over which passes are observed, that corresponds to the number of grouped passes and optimized manoeuvres.

Regarding the manoeuvre magnitudes, an upper limit is set both in the positive and negative along track directions (Tab. 2):

Table 2. Bounds for manoeuvre magnitude [11]

$UB_{\Delta v}$	$LB_{\Delta v}$
30 mm/s	-30 mm/s

Bounds for the manoeuvring epochs are instead designed ad-hoc for each manoeuvre, depending on which day they are happening, and considering that each batch of grouped passes is positioned after the current manoeuvre and before the next one. This is summarized in Tab. 3 with $t^{\text{last}}_{\text{meas},i}$ being the last epoch of the batch of grouped passes on the i-th day.

Table 3. Bounds for manoeuvring epochs

n_{man}	UB_t	LB_t
$i = 0$	$t^{\text{last}}_{\text{meas},0}$	t_i
$i > 0$	$t^{\text{last}}_{\text{meas},i}$	$t^{\text{last}}_{\text{meas},i-1}$

2.2. Use Cases

The performance of the optimizer is analysed, in different designed test cases, from the nominal scenario to a more complex one, identifying the strengths and the criticalities of the method. The observation data has been simulated using own software, assuming a LEO object similar to STARLINK-5885 (NORAD ID: 56352, Keplerian elements and considered manoeuvres are described in Tab. 4). The simulated stations are located in Southern Spain (37°10'N, 5°36'W) and in Kiruna (67° 51'N, 20° 26'E)

Table 4. Keplerian elements of the observed satellite and manoeuvring epochs

a	e	i	Ω	ω
6935.69 km	$1.384 \cdot 10^{-3}$	97.5°	274.9°	46°
First manoeuvre	October 2^{nd} 2024, 00:00:00 UTC			10 mm/s
Second manoeuvre	October 3^{rd} 2024, 00:00:00 UTC			10 mm/s
Third manoeuvre	October 4^{th} 2024, 00:00:00 UTC			10 mm/s

2.2.1. Nominal case

In the nominal case scenario the tracked space object performs one manoeuvre per day, and the ground stations correctly observe the passes over their field of view. A total of 67 observations per day are available for the analysis.

The chosen figure of merit to evaluate the goodness of the solution at convergence is the cumulative error in the epochs:

$$e_{\text{time}} = \sum_{i}^{n_{\text{man}}} |\hat{t}_i^{\text{man}} - t_i^{\text{man}}| \tag{8}$$

with $\hat{t}$ being the true manoeuvring epoch, and t being the result of the optimization. The obtained cumulative probability distribution, on a total of **75** tests, out of **80**, filtered with $e_{\text{time}} < 1000$ s, is shown in Fig. 2.

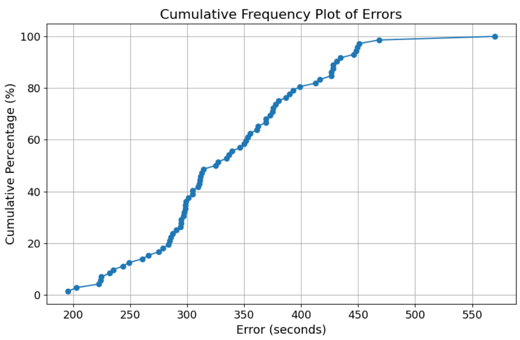

Figure 2. Filtered error cumulative probability distribution of the population evolved in genetic algorithm. e_{time} in abscissa.

As it can be seen with the selected configuration the optimizer almost always converges to an acceptable solution. This means that the cost function has been correctly designed, with all the contributions weighted equally.

The distribution of the solutions can also be shown in a graph linking the errors in the epochs and the magnitudes (Fig. 3), with each solution overlayed with the associated WRMS. Two clusters of solutions can be identified for each manoeuvre, with the WRMS values always close to unity, as in Eq. 5.

Considering the result of a single optimization run, an orbit comparison between the result at convergence, and the real manoeuvres, can be performed, with the propagation performed by a high-fidelity propagator. Differences in radial, along-track and out-of-plane components of the state vectors of the true and optimized orbits are shown in Fig. 4. The error is always in the order of few meters also in the end of the propagation, and it is mainly along track.

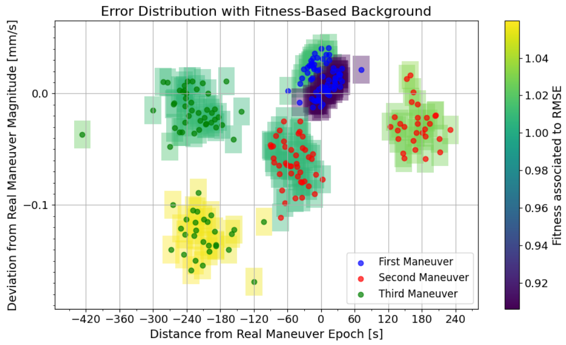

Figure 3. Heat map of solutions for each of the starting individuals of the genetic algorithm.

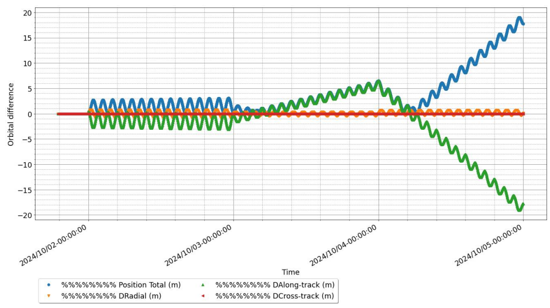

Figure 4. Orbital comparison between simulated and estimated orbit for nominal case optimal solution.

2.2.2. Data Scarcity subcase

In this subcase observations come only from one of the two ground stations, which is the one located in Spain. This situation ought to simulate a malfunction on one of the radars. A total of 28 observations per day are available.

With measurements from only one station the amount of available information for the optimizer diminishes, making it more challenging for it to accurately converge towards the true solution, as can be seen from Fig. 5.

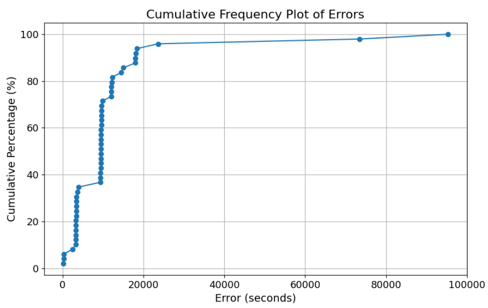

Figure 5. Error cumulative probability distribution for tests in data scarcity subcase.

The results are much worse than when two stations are available: out of a total of **49** performed tests, only **3** ($\approx 6\%$) were below the 1000 s threshold.

Even though results are bad, it doesn't mean that the optimizer does not converge: indeed in all the tests the obtained WRMSs for the measurements, entering the cost function, were around unity.

This means that the optimizer does not have enough information on the orbital status of the tracked space object to correctly determine the manoeuvring epochs and magnitudes.

2.2.3. No-measurement case

In this scenario, the tracked space object performs one manoeuvre per day as in Tab. 4. However observations from the radar stations are not produced during one day (the first one), to simulate a malfunction on ground.

This translates into the fact the optimizer has to guess two manoeuvres with only one group of measurements. The obtained result, on a total of **30** tests, is a cloud of points spread around the central optimal solution for the first and second manoeuvres, while the third one almost always converged close to the optimum (Fig. 6). However this does not imply that the optimizer didn't actually work as intended, since looking at the values of the cost function, they are always close to unity.

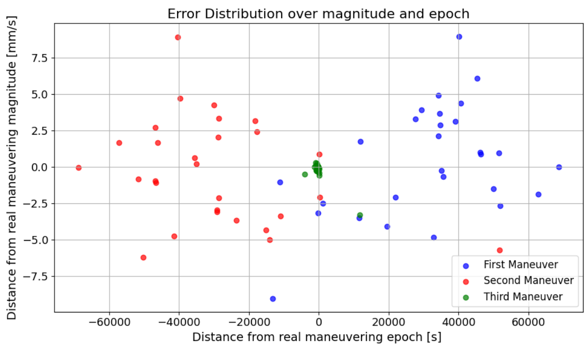

Figure 6. Set of solutions for no-measurement case.

An explanation for the fact that the optimizer converged, but to a solution that is not the real one, is that since in the first day measurements are missing, the problem is less constrained. The algorithm converged to a solution that guarantees a good fit with the available measurements of the second and third days. This is similar to what happens in the case with only one station: there is not enough information to correctly converge to the real solution, since the optimization space is vast and multiple minima are possible.

Analysing the orbit comparison between one of the solutions and the real orbit, as it has been done in the nominal case, the same conclusion can be inferred, considering that its solution at convergence for the considered optimization run is the one presented in Tab. 5:

Table 5. Optimization result test 26 in no-measurements case

Manoeuvring epochs	Manoeuvring magnitudes
October 2^{nd} 2024, 09:39:09 UTC	9.166 mm/s
October 2^{nd} 2024, 14:20:57 UTC	10.96 mm/s
October 4^{th} 2024, 00:00:38	9.742 mm/s

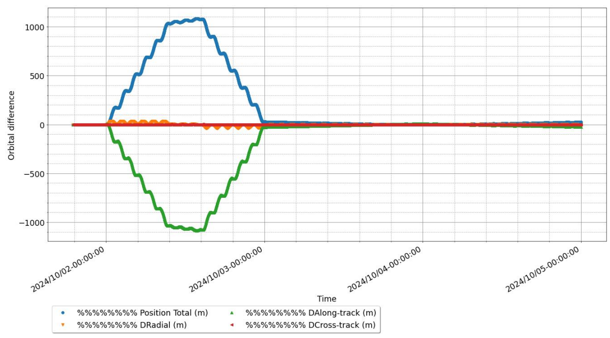

Figure 7. Orbital comparison between simulated reference and estimated for no-measurements cases.

From Fig. 7 it is seen that in the first day the orbit diverges from the true one, since the problem is unconstrained and the optimizer cannot rely on measurement residuals to determine the true orbital fit. In particular the first manoeuvre is missed by approximately 9 hours. However, together with the second one, they are placed such that at the epoch of the second real manoeuvre, the 3^{rd} of October 00:00:00 UTC, the orbital difference drops drastically again. What happens is that the optimizer finds one of the multiple solutions to the unconstrained problem (the optimal one), which is different from the real one but in the end leads to the same orbital position.

3. MANOEUVRE HANDLING WITH SEQUENTIAL ESTIMATION

An alternative is presented aiming at maintaining the trackability of the space object by characterising the uncertainty introduced by the manoeuvre, rather than the manoeuvre itself, using covariance inflation, attributables and smoothing within sequential estimation. This may prevent us from inferring the manoeuvre characteristics, which is not required for catalogue maintenance purposes, while ensuring continuous tracking of the object. The detection and handling of the manoeuvre is performed as an additional step of the so-called Manoeuvre Detection Filter (MDF) built on an Extended Kalman Filter (EKF).

3.1. Methodology

For the tracking scenario, a high density of observations is available at frequent intervals. The EKF is particularly well-suited for this application due to its computational efficiency and sequential, forward-moving prediction capabilities [12]. The MDF is implemented for this task and is described in Algorithm 1. It consists of an additional step in the sequential filtering after the prediction step and before the correction step.

The MDF consists of two parts: the detection and the handling. A manoeuvre detection metric has been computed by modeling the squared mahalanobis distance (MD) as a χ^2 distribution function. To declare a manoeuvre event, a residual χ-square test is used. For this, MD between the measurements and the predicted state and covariance is used in the observation domain. Clearly, the MD is a scalar function of the observational residual which is continuously updated during the estimation process. If a manoeuvre occurs, MD increases proportionally to the growth of the residual. Furthermore, by accounting for the inherent noise in the residual, MD can be used to identify the onset of a manoeuvre. The probability of manoeuvre PR_{md} defined in [8] is used as the final metric for manoeuvre detection to reduce the number of false positives. When PR_{md}, which is a function of MD, exceeds a predefined threshold, a manoeuvre is assumed to have happened. Among the four radar measurements, angle measurements are usually the least accurate and are less effective at detecting manoeuvres compared to range and range-rate measurements, which provide quicker and more reliable indications of orbital changes [13]. Therefore, range and range-rate measurements are utilized in the MDF, meaning that the number of degrees of freedom in the χ-squared test is set to $n = 2$.

The χ^2 cumulative distribution function is computed using the regularized lower incomplete gamma function $P(a, x)$ with a being half of the degrees of freedom and x being half of the squared MD. $P(a, x)$ is defined as the ratio between the lower incomplete gamma function and the complete gamma function :

$$P(a,x) = \frac{\int_0^x t^{a-1}e^{-t}dt}{\int_0^\infty t^{a-1}e^{-t}dt} \tag{9}$$

In hypothesis testing, and using the χ-square distribution, it is usual to compute $P(\frac{n}{2}, \frac{MD^2}{2})$ which is the probability that a variable distributed according to the χ-square distribution with n degrees of freedom exceeds the value of MD^2. The problem here faced is the one of computing the confidence level given an experimentally determined χ-square value and this involves the direct computation of the cumulative distribution function [14].

The predicted covariance P_k^- is inflated until $PR_{md} < 0.5$. The larger the manoeuvre, the larger the inflation. This enables the filter to not diverge and to keep tracking the target despite an un-modeled manoeuvre. Radar data is processed after the time of the manoeuvre until the sequential filter converges again, that is, when the uncertainty in the state estimate returns to a nominal ballistic condition.

The MDF is detailed in the Algorithm 1. The variables $\hat{\boldsymbol{x}}_k$ and $\boldsymbol{P}_k$ represent the estimated state and covariance at step k of the filter. $\boldsymbol{\Phi}_{k-1,k}$ represents the State Transition Matrix (STM) between steps $k-1$ and k. The superscripts $-$ and $+$ indicate whether the quantity corresponds to the prediction from step $k-1$ $(-)$, or to the update after incorporating the observation at step k $(+)$. Additionally, $\boldsymbol{y}_k$ is the actual observation, $\boldsymbol{R}_k$ is the sensor noise covariance and $\boldsymbol{Q}_k$ is the process noise covariance.

Algorithm 1 Manoeuvre Detection Filter (MDF)

Inputs: $\hat{\boldsymbol{x}}_{k-1}^+$ and $\boldsymbol{P}_{k-1}^+$

Outputs: $\hat{\boldsymbol{x}}_k^+$ and $\boldsymbol{P}_k^+$

(1) Propagate the state until next attributable epoch (step k)

(2) Compute $\boldsymbol{\Phi}_{k-1,k}$ between $k-1$ and k

(3) Propagate the state error covariance

$$\boldsymbol{P}_k^- = \boldsymbol{\Phi}_{k-1,k}\boldsymbol{P}_{k-1}^+\boldsymbol{\Phi}_{k-1,k}^T + \boldsymbol{Q}_k, \quad \boldsymbol{P}_0^+ = \boldsymbol{P}_0$$

(4) Compute the theoretical attributable $\hat{\boldsymbol{y}}_k$ with the observation function $\boldsymbol{h}$ and the partials matrix $\boldsymbol{H}_k$:

$$\hat{\boldsymbol{y}}_k = \boldsymbol{h}(\hat{\boldsymbol{x}}_k^-), \quad \boldsymbol{H}_k = \left.\frac{\partial \boldsymbol{h}}{\partial \boldsymbol{x}}\right|_{\hat{\boldsymbol{x}}_k^-}$$

(5) Compute the total covariance of the observation $\boldsymbol{S}_k$

$$\boldsymbol{S}_k = \boldsymbol{H}_k\boldsymbol{P}_k^-\boldsymbol{H}_k^T + \boldsymbol{R}_k \tag{10}$$

(6) Compute the Mahalanobis distance MD of range and range-rate residuals

$$\mathrm{MD}_k = \sqrt{(\boldsymbol{y}_k - \hat{\boldsymbol{y}}_k)\boldsymbol{S}_k^{-1}(\boldsymbol{y}_k - \hat{\boldsymbol{y}}_k)^T} \tag{11}$$

(7) Compute the probability of manoeuvre

$$PR_{md} = \max\left(0, 2(\chi^2(\mathrm{MD}_k{}^2, n) - 0.5\right) \tag{12}$$

If $PR_{md} > 0.5$, multiply $\boldsymbol{P}_k^-$ by 2 and come back to step (5). Otherwise, continue.

(8) Compute the Kalman gain and update the estimated state and covariance

$$\boldsymbol{K}_k = \boldsymbol{P}_k^-\boldsymbol{H}_k^T\boldsymbol{S}_k^{-1} \tag{13}$$

$$\hat{\boldsymbol{x}}_k^+ = \hat{\boldsymbol{x}}_k^- + \boldsymbol{K}_k(\boldsymbol{y}_k - \hat{\boldsymbol{y}}_k) \tag{14}$$

$$\boldsymbol{P}_k^+ = (\boldsymbol{I} - \boldsymbol{K}_k\boldsymbol{H}_k)\boldsymbol{P}_k^-(\boldsymbol{I} - \boldsymbol{K}_k\boldsymbol{H}_k)^T + \boldsymbol{K}_k\boldsymbol{R}_k\boldsymbol{K}_k^T \tag{15}$$

The process noise covariance $\boldsymbol{Q}_k$ is null when we deal with simulated observations where the dynamical model is identical for both the generated tracks and the estimation filter. However, the State Noise Compensation (SNC) technique can be used when dealing with real observations. In addition, a backward smoothing process is applied to refine the post-manoeuvre region, where covariance inflation has been applied. The McReynolds consistency test [15] can also be leveraged for manoeuvre detection.

3.2. Use Cases

The subject of the tests is based on Sentinel 3-A, which is on a low-eccentricity, near-polar, sun-synchronous orbit. Both the orbit and manoeuvre history are publicly available and taken from [16] and [17] respectively. The time window considered for the simulation starts in September 1^{st}, 2018, 10:30 and ends in September 12^{th}, 2018, 00:00.

Table 6. S3A satellite data.

Epoch	September 1^{st}, 2018, 10:30 (UTC)
Position vector	$[-2301.83, 1156.13, 6694.98]^T$ km
Velocity vector	$[-4.27, 5.60, -2.43]^T$ km/s

The altitude of the orbit is 814.5 km with an inclination of 98.65 deg.

Station 1 will be used as the only source of measurement tracking data. The latter will be simulated and used with approximately 4 to 5 tracks per day. The simulated measurement noise is detailed in Tab. 7.

Table 7. Simulated radar sensor measurement noise.

σ_{Az}	σ_{El}	σ_R	$\sigma_{\dot{R}}$
0.4 deg	0.4 deg	20 m	650 mm/s

Station 1 is inspired by the the Spanish S3T Surveillance Radar (S3TSR) that is a radar system developed by Indra within a project funded by Spanish Administration and technically managed by ESA [18].

The manoeuvre history of Sentinel-3A has been retrieved from the *International Laser Ranging Service* (ILRS) public data [17]. The most typical manoeuvres are either impulsive, with duration in the order of a few seconds, or long, with duration between 12 and 15 minutes. Impulsive manoeuvres are usually in the along-track component while long manoeuvres have a major cross-track component and a minor along-track component. Some manoeuvres with higher Δv in the along-track direction and of longer duration occur exceptionally (e.g., on 07/01/2022). An in-track impulsive manoeuvre case and an out-of-plane long manoeuvre case will be considered in this section.

The simulation begins with a high covariance, reflecting the initial uncertainty in the state. As the filter processes subsequent tracks, the covariance gradually decreases, indicating improved confidence in the estimated state. In both cases, the fourth track was intentionally selected as an outlier to assess its impact on the MDF. This track corresponds to a vertical pass, where the elevation reaches 90 deg and therefore the azimuth exhibits discontinuities.

3.2.1. Impulsive manoeuvre

The MDF is put to the test with an impulsive low-thrust manoeuvre with a major in-track component.

Table 8. Impulsive manoeuvre characteristics.

Epoch t_M	September 5^{th}, 2018 19:21:10 (UTC)
Components (RTN)	[0.19472, -3.05837, 0.02038] mm/s
Magnitude	3.06463 mm/s
Arrival time of next track	$t_M + 1h$

This in-track manoeuvre is taken from a real Sentinel-3A [17] that occurred on February 27^{th}, 2019, 09:15:37. It typically corresponds to an orbit raising operation needed to correct the semi-major axis decay caused by the atmospheric drag [19].

As shown in Fig. 8, impulsive manoeuvres are not detected immediately but typically after 3 to 4 tracks. Due to their low Δv nature, the filter struggles to identify anomalies from the first post-manoeuvre track, making immediate detection challenging. Nonetheless, immediate detection is not critical for this method. The filter can still effectively converge and handle the manoeuvre, ensuring accurate tracking despite the delayed detection. Additionally, the outlier is detected as a manoeuvre and incorrectly inflates the covariance; nonetheless, it is not critical for the filter that gets the covariance reduced from the subsequent tracks.

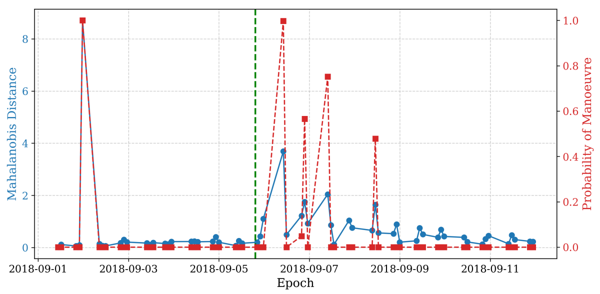

Figure 8. Mahalanobis distance MD and PR_{md} for impulsive manoeuvre test case.

The filter residuals are depicted in Fig. 9 During filter initialization, the disparity between pre-fit and post-fit residuals is expected. However, as the filter progresses, it converges. From the manoeuvre epoch onward, range and range-rate pre-fit and post-fit residuals begin to diverge. In this case of an impulsive low-thrust manoeuvre, convergence is slower, but the covariance inflation mechanism still enables the filter to adapt over time, eventually leading both pre-fit and post-fit residuals to align, as expected.

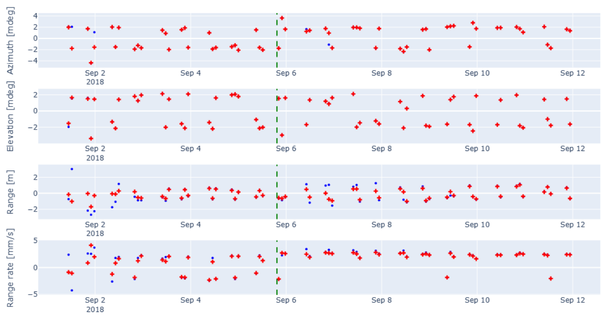

Figure 9. Filter residuals for impulsive manoeuvre test case. Symlog scale. Pre-fit residuals (blue circle). Post-fit residuals (red cross). Green dashed vertical line correspond to real manoeuvre epoch.

While it is often possible to maintain tracking of an object undergoing a low-thrust manoeuvre, the primary difficulty lies in precisely detecting that the manoeuvre has occurred and distinguishing it from unmodeled non-conservative perturbation effects [20]. They may go undetected or detected with a significant delay if residuals are used as the primary means for detection, like here. Fortunately, the detection of the manoeuvre is not crucial for catalogue maintenance and this is the reason why a Manoeuvre Detection Filter is well suited for this scenario.

An alternative to detect small manoeuvres is to perform a smoothing process using the Rauch-Tung-Striebel (RTS) fixed-interval smoother [21]. The McReynolds filter-smoother consistency test [15] is initially thought for determining if the process noise is correctly tuned when un-modeled perturbations are present, like in the case of real observations. Therefore, an unknown manoeuvre might leverage this consistency statistic and provide an additional metric that can be used to detect the manoeuvre, provided that the filter/smoother process is properly "tuned" to the specific object and its associated data [20].

The backward smoothing process is applied to refine the post-manoeuvre region, where covariance inflation has been introduced to accommodate the uncertainty induced by the manoeuvre. This process enhances state estimation by incorporating future observations, allowing for a more accurate reconstruction of the trajectory. The smoothing procedure terminates when the consistency test fails, effectively serving as a redundant detection metric. The manoeuvre detection threshold is set such that a manoeuvre is identified when the absolute value of any element exceeds the predefined threshold $|R_i| > 4.5$ [22].

The failure of the consistency test often occurs near the actual manoeuvre onset epoch, making it a valuable indicator of state deviations. However, this approach is most effective when the manoeuvre takes place while the satellite is within the field of view of a ground station, where frequent observations improve detection accuracy [22]. Despite this limitation, the estimated manoeuvre epoch remains sufficiently close to the true event, providing a useful reference for a subsequent filter-smoother iteration. These iterations should enable further refinement of the state estimate, improving the accuracy of the trajectory reconstruction. Additionally, they should bring the detected manoeuvre epoch closer to the actual event, enhancing the reliability of the detection process.

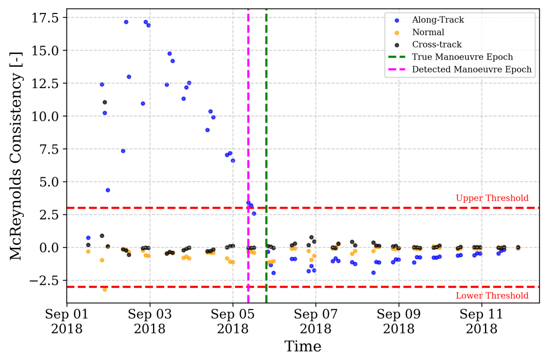

Figure 10. McReynolds consistency for impulsive manoeuvre test case.

Additionally, the consistency test provides a strong indication of a potential manoeuvre occurring in the tangential direction, reinforcing its role as a valuable diagnostic tool. By identifying discrepancies between the estimated and actual trajectory, this test serves as an effective means of post-event confirmation, allowing to retrospectively validate manoeuvre occurrences with a high degree of confidence.

Moreover, this information can be leveraged for pattern-of-life analyses as it offers insights into recurrent manoeuvre behaviors and operational strategies of the satellite. By systematically monitoring these consistency failures over time, it becomes possible to infer manoeuvring trends and anticipate future orbital adjustments. This capability is particularly crucial for catalogue maintenance and conjunction assessment, ensuring a more accurate and long-term orbit maintenance.

3.2.2. Long manoeuvre

The MDF is put to the test with a long hybrid manoeuvre with a major cross-track and a minor in-track component. The simulated manoeuvre takes place on September 5, 2018, 19:21:10 (UTC), with a duration of 13.82 minutes and a magnitude of 2.13 m/s in out-of-plane direction. This manoeuvre is taken from a real Sentinel-3A [17] that occurred on March 13^{th}, 2019, 08:15:15. It typically corresponds to inclination drift corrections mainly caused by the luni-solar perturbation.

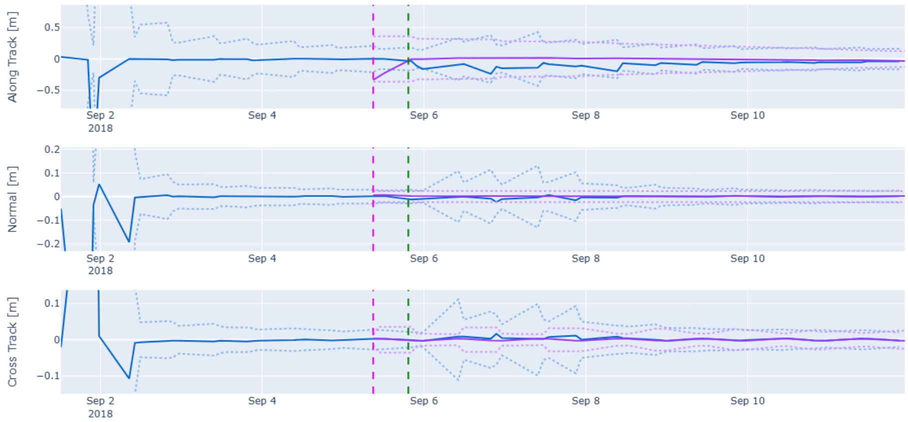

Figure 11. Orbital difference (TNW) between the estimated/smoothed trajectory and the ground truth. 3σ covariance envelope. Impulsive manoeuvre test case. Green dashed vertical line corresponds to the real manoeuvre epoch and pink dashed line corresponds to the detected manoeuvre epoch by the smoother. EKF (blue) vs RTS (violet).

During filter initialization, the disparity between pre-fit and post-fit residuals is expected. However, as the filter progresses, it converges effectively. From the manoeuvre epoch onward, range and range-rate pre-fit and post-fit residuals begin to diverge. In this case, the covariance inflation mechanism ensures the filter to quickly converge to the true trajectory after detecting the manoeuvre, allowing that both pre-fit and post-fit residuals ultimately align, as expected. Again, a backward smoothing process is applied to refine the post-manoeuvre region, where the covariance has been inflated. The smoothing procedure terminates when the consistency exceeds $|R_i| > 4.5$.

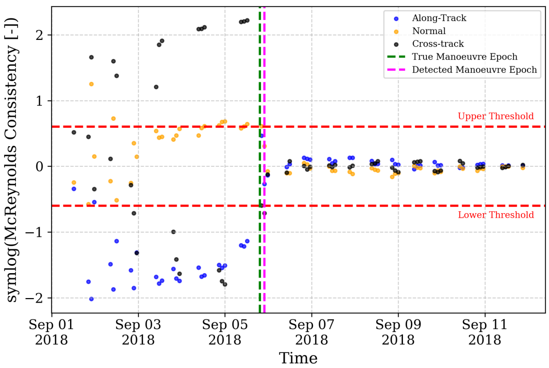

Figure 12. McReynolds consistency for long manoeuvre test case. Symlog scale.

In Fig. 12 the results for McReynolds consistency for this case can be found. As observed again, this additional detection method successfully identifies the manoeuvre at the post-manoeuvre epoch. Furthermore, it provides valuable insight into the manoeuvre Δv components, suggesting the presence of both a dominant cross-track and minor in-track contributions, which accurately reflects the actual manoeuvre characteristics. This further validates the method's ability to not only detect manoeuvres but also infer key dynamical properties, enhancing its applicability for catalogue maintenance. The final orbital differences are displayed in Fig. 13.

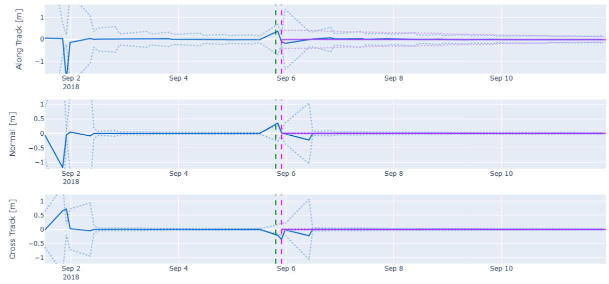

Figure 13. Orbital difference (TNW) between the estimated/smoothed trajectory and the ground truth. 3σ covariance envelope. Long manoeuvre test case. Green dashed vertical line corresponds to the real manoeuvre epoch and pink dashed line corresponds to the detected manoeuvre epoch by the smoother. EKF (blue) vs RTS (violet).

4. MULTIPLE HYPOTHESIS TRACKING WITH MANOEUVRE ESTIMATION

To address the survey problem previously defined, we introduce a Multiple Hypotheses Tracking (MHT) methodology that jointly tackles the track-to-orbit (T2O) association and manoeuvre estimation problems. The proposed methodology is based on the work [23] and provides a robust framework for associating uncorrelated tracks with known catalogued objects by explicitly accounting for dynamic changes induced by significant manoeuvres. This approach systematically evaluates multiple hypotheses to determine the most likely source object of the observed tracks while inferring the manoeuvre parameters, specifically the epoch as well as the magnitude and direction of the applied Δv. By incorporating these dynamic changes into the association process, this methodology enhances the effectiveness of SSA systems in maintaining accurate and continuously updated catalogues while minimizing the likelihood of miscorrelations.

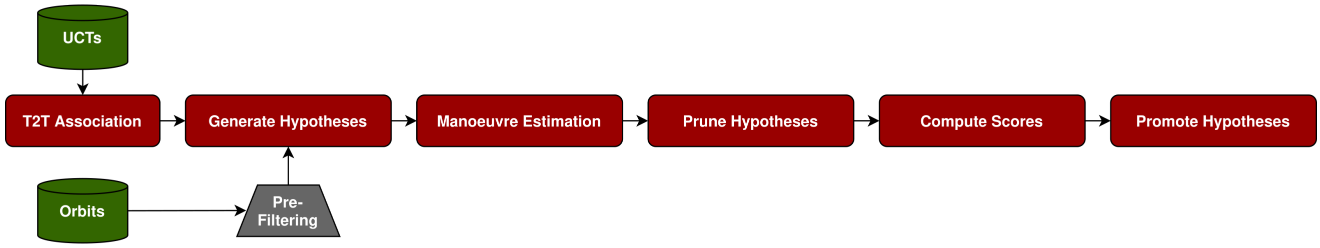

Figure 14. Track-to-Orbit with manoeuvres processing chain.

4.1. Methodology

This section provides an overview of the proposed methodology, highlighting its key features and outlining the main processing steps.

The first step in catalogue maintenance is to determine whether an observed track corresponds to an already catalogued space object. This is done during the T2O process that attempts to associate new observations with the predicted positions of already catalogued objects. Here, the likelihood of an association is determined by evaluating the observational residuals between the actual observations and their predicted counterparts. If the correlation is successful, the track is confirmed to have originated from an existing object, and its data is subsequently used to update the object's trajectory through an orbit determination. However, this process may fail if the track originates from an uncatalogued object or if a known object has altered its trajectory, possibly due to a manoeuvre.

The objective of the proposed methodology is to process these uncorrelated tracks (UCT), where the conventional T2O algorithm failed due to a manoeuvre. This approach follows the processing steps outlined in Fig. 14, which will be described in more detail in the subsequent paragraphs.

1. *Track-to-Track (T2T) Association*: First, UCTs are grouped into clusters that are likely to originate from the same space object. This is achieved using multi-target, multi-sensor algorithms that systematically assess observational similarity and orbital consistency among individual tracks, facilitating robust and reliable associations. [9] [10]

2. *Generate Hypotheses*: In this context, a hypothesis represents a potential association between a set of previously uncorrelated tracks and a known catalogued space object. Specifically, it seeks to determine whether a plausible manoeuvre can account for the observed deviation from the expected trajectory. This is achieved by estimating the manoeuvre parameters necessary to align the observed post-manoeuvre tracks with the predicted orbit of the candidate satellite. Theoretically, the number of generated hypotheses L is equivalent to $L = n \cdot m$, where n refers to the number of catalogued objects and m to the number of T2T associations (i.e., groups of tracks originating from the same object). Since there are currently more than forty thousand catalogued objects, the number of manoeuvre estimations grows rapidly when employing a brute-force approach. To mitigate this computational burden, the observed object's Keplerian elements are approximated based on the T2T associations and subsequently compared with those of all catalogued objects. If the differences in Keplerian elements exceed predefined thresholds, the corresponding objects are eliminated from the list of hypotheses, significantly reducing computation time. Depending on the orbital regime, additional filtering criteria, such as geocentric longitude for GEO objects, can be incorporated to further enhance computational efficiency.

3. *Manoeuvre estimation*: Once the potential candidates are filtered, the manoeuvre estimation for each remaining hypothesis is triggered. During this process, the manoeuvre parameters, including its magnitude, direction, and execution time are estimated using a weighted least-squares approach. This is accomplished by fitting post-manoeuvre observations to the object's pre-manoeuvre orbit, utilizing an iterative approach to refine these estimates until they closely align with the observed data. Here, the post-manoeuvre tracks refer to the set of tracks grouped through T2T, while the pre-manoeuvre orbit corresponds to the candidate object's catalogued orbit. For a more comprehensive description of the underlying manoeuvre estimation algorithms, refer to [4] and [5].

4. *Prune Hypotheses*: After completion of the manoeuvre estimation for each hypothesis, unlikely candidates are eliminated. For this pruning process, two metrics are used: the estimated Δv and the WRMS of the residuals. Here, the WRMS quantifies how closely the predicted post-manoeuvre orbit aligns with the actual sensor observations, indicating the accuracy of the estimated manoeuvre. Excessively large WRMS values therefore indicate that the investigated tracks are unlikely to have originated from the candidate object. Similarly, if the Δv exceeds a predefined threshold, the manoeuvre is deemed unfeasible, leading to the elimination of the corresponding hypothesis from the list of poten-

tial candidates.

5. *Compute Scores*: The most likely candidate is not selected solely based on the hypothesis with the lowest Δv. Instead, the selection process incorporates two distinct scoring metrics to ensure a more robust and accurate determination. The WRMS and Δv of the hypotheses that remain after the pruning process are first normalised using a min-max normalisation according to

$$X_i^* = \frac{X_{max} - X_i}{X_{max} - X_{min}}. \tag{16}$$

being, X, the variable to be normalize. This is done to ease comparison between values of different orders of magnitude, scaling them to a range between 0 and 1. The normalized values, combined with their respective weights $W_{\Delta v}$ and W_{WRMS} are then used to compute the weighted sum score as follows:

$$S_i = W_{\Delta v} \cdot \Delta v_i + W_{WRMS} \cdot WRMS_i. \tag{17}$$

The weights are constrained to sum to one and can be determined either through Principal Component Analysis or manual tuning to optimize performance. As a second metric, the SoftMax score, P_i, is used to transform the weighted sum scores of competing hypotheses into normalized probabilities. It quantifies the confidence level for each hypothesis, enhancing the decision-making process by rewarding higher scores and penalising lower ones. The hypothesis with the highest SoftMax score is promoted, attempting to resolve ambiguities among closely competing options. The score can be computed based on Eq. 18, where β refers to the ratio between the maximum and minimum values of the weighted sum S_i.

$$P_i = \frac{e^{\beta S_i}}{\sum_i e^{\beta S_i}} \tag{18}$$

6. *Promote Hypotheses*: Finally, after scoring all remaining hypotheses, the best-performing ones are selected for promotion. To qualify for promotion, they must surpass a predefined threshold. The promoted hypotheses represent the candidates with the highest likelihood of being the source objects from which the tracks originated.

4.2. Use Cases

To evaluate the proposed methodology, simulated observation scenarios are designed, incorporating both nominal satellite tracking and manoeuvre events. The satellites considered in this study are modelled based on Starlink satellites, reflecting their orbital elements as well as key properties such as mass, area, and drag coefficients. However, unlike actual Starlink satellites, the simulated manoeuvres are modelled as impulsive, meaning the Δv is applied instantaneously. Additionally, only manoeuvres with thrust applied in the along-track direction are considered. Four distinct satellites are considered, with their Keplerian elements listed in Tab. 9, differing only in their true anomaly $\Delta\Theta$. The suffixes A1, A2, A3, and A4 in the results correspond to the respective Starlink satellites.

For the sake of simplicity, all observations in this study are simulated using the spanish radar described in Section 2.2 as reference, incorporating its location and sensor noise characteristics.

Table 9. Keplerian elements of the observed satellites.

a	e	i	Ω	ω
$6928.14km$	$1.5 \cdot 10^{-4}$	$53°$	$188°$	$10°$

As a baseline, a case is constructed based on the assumption of $\Delta v = 10$ cm/s and a phasing of $\Delta\Theta = 15°$ between each of the satellites. For this reference case, the initial T2T process achieves a 100% success rate, meaning that all post-manoeuvre tracks are correctly grouped with the remaining tracks of the same object. Subsequently, each set of grouped tracks is successfully associated with its corresponding source object, while accurately recovering the manoeuvre parameters during the estimation. However, real-world scenarios are often significantly more challenging due to factors such as the close proximity of satellites or high manoeuvre magnitudes. Additionally, adverse observation conditions, resulting from high sensor noise or limited track availability for a given object, can further degrade association performance. To examine the impact of these factors, several test cases are investigated. For additional test cases, an analysis of the entropy scoring metric, and a more comprehensive evaluation of the presented studies, refer to [23].

Many satellites, particularly those in constellations, are often clustered into groups where multiple satellites share the same orbital plane, differing only in their true anomaly. As a result, distinguishing between individual satellites based on noisy observations can be challenging. Therefore, it is crucial to assess how satellite proximity affects the performance of association algorithms.

Now, considering a more challenging scenario with $\Delta\Theta = 5°$, 88% of the post-manoeuvre tracks are successfully grouped through T2T association. Fig. 15 depicts a subset of the generated hypotheses along with the corresponding relative errors of the estimated Δv and manoeuvre epoch. The heatmap colouring represents the scores assigned to each hypothesis, with higher scores indicating greater likelihood, while the hypotheses ultimately promoted are highlighted with red borders. It is evident that hypotheses with larger errors, particularly in Δv, receive significantly lower scores, as indicated by the purple colouring. Tab. 10 provides a detailed breakdown of these results, including the score S, the classification as a true positive TP, and the relative errors of the estimated manoeuvre magnitudes and epochs for each of the satellites. Although the manoeuvre epoch for A4 is estimated to be one orbital period earlier than the actual manoeuvre occurrence, all tracks are successfully associated with the respective candidate object. Furthermore, since the primary objective is to reliably predict the post-manoeuvre orbit rather than precisely infer the manoeuvre parameters themselves, this deviation does not adversely impact the cataloguing process.

Table 10. Statistics of the performance metrics for promoted hypotheses in the test case with a satellite spacing of $\Delta\Theta = 5°$

	$\boldsymbol{S}$	$\varepsilon_{\Delta v}[\%]$	$\varepsilon_{t_M}[min]$	**WRMS**	**TP**
A1	0.96	3.64	-1.06	1.025	✓
A2	0.964	4.02	-1.77	0.965	✓
A3	0.967	1.1	1.85	1.091	✓
A4	0.955	-0.17	-88.64	1.274	✓

A clear relationship between satellite separation and algorithm accuracy emerges when the spacing is reduced to $\Delta\Theta = 1°$. In this scenario, only 62% of post-manoeuvre tracks are successfully correlated during T2T. However, after the manoeuvre estimation, all of these tracks are correctly associated with their respective objects while avoiding false positives through the use of the SoftMax score. When satellites are relatively distant, the association process remains highly reliable, yielding accurately estimated manoeuvre with minimal ambiguity. However, as satellite separation decreases, distinguishing individual objects becomes increasingly challenging. This is reflected in higher WRMS values and reduced confidence in the generated hypotheses

When considering a spacing of $\Delta\Theta = 5°$ while doubling the Δv to 20 cm/s, the results for T2T association, manoeuvre estimation and subsequent scoring remain nearly identical to those observed in the case with $\Delta v = 10$ cm/s. This suggests that the association algorithms are less sensitive to manoeuvre magnitude compared to changes in satellite proximity. However, when analysing the combined effect of a close spacing of $\Delta\Theta = 1°$ and an increased Δv of 20 cm/s, the results presented in Tab. 11 are obtained. All of the promoted hypotheses exhibit low errors, with both WRMS and scoring values close to the optimum. As expected, the T2T association rate further declines due to the increased manoeuvre magnitude, leading to a reduced T2T success rate of only 50% and a notable increase in false positives during the T2T process. However, with a single exception, all miscorrelated tracks are successfully discarded after the manoeuvre estimation, during the WRMS and Δv pruning stage.

In the following scenario, the impact of doubled radar noise on association performance is examined. Using the previously discussed case as a baseline, this scenario introduces additional complexity due to close satellite proximity, high manoeuvre magnitude, and increased sensor noise. While the T2T association rate remains around 62%, the accuracy of T2O correlation and manoeuvre estimation is significantly degraded, leading to eight false positives and multiple ambiguous cases. To identify and discard these false positives, it is essential to prune the hypotheses and subsequently employ the SoftMax score. As shown in Tab. 12, the estimation errors are considerably larger than in the previous cases. The WRMS on the other hand remains around one, as the increased noise is accounted for when computing the residuals. Despite these challenges, the pipeline successfully promotes most of the optimal hypotheses, albeit with a notable reduction in overall accuracy compared to less noisy scenarios.

Table 11. Statistics of the performance metrics for promoted hypotheses in the test case with a satellite spacing of $\Delta\Theta = 1°$ and $\Delta v = 20\ cm/s$.

	$\boldsymbol{S}$	$\varepsilon_{\Delta v}[\%]$	$\varepsilon_{t_M}[min]$	**WRMS**	**TP**
A1	0.996	0.62	-0.36	1.025	✓
A2	0.992	4.3	-2.54	0.975	✓
A3	0.991	1.5	2.18	1.069	✓
A4	0.99	0.78	1.24	1.099	✓

Table 12. Statistics of the performance metrics for promoted hypotheses in the test case with a satellite spacing of $\Delta\Theta = 1°$, $\Delta v = 20\ cm/s$ and doubled radar noise.

	$\boldsymbol{S}$	$\varepsilon_{\Delta v}[\%]$	$\varepsilon_{t_M}[min]$	**WRMS**	**TP**
A1	0.959	17.26	-109.97	1.166	✓
A2	0.994	2.42	-0.25	1.067	✓
A3	0.998	3.05	1.5	0.991	✓
A4	0.976	9.09	26.02	1.147	✓

Finally, we evaluate the impact of using fewer tracks during the association on its performance. In all previously presented cases, each T2T association consisted of four distinct tracks, whereas in the following results, only three tracks were used. Reducing the number of tracks inherently limits the information available in the T2T association, leading to degraded manoeuvre estimation and a subsequent decline in the accuracy of correlations with known objects. For this scenario, the T2T association rate is approximately 54%, with the association results detailed in Tab. 13. Compared to the previously analysed case with four tracks, the number of false positives increases by nearly 60% due to the heightened uncertainty introduced when relying on only three tracks. As the manoeuvre estimation process relies on the residuals of observations to accurately determine manoeuvre parameters, the reduction in available information increases uncertainty, leading to less precise estimation. Despite this, the correlation process remains effective, successfully promoting the correct hypotheses. However, the reduction in track count introduces significant errors, par-

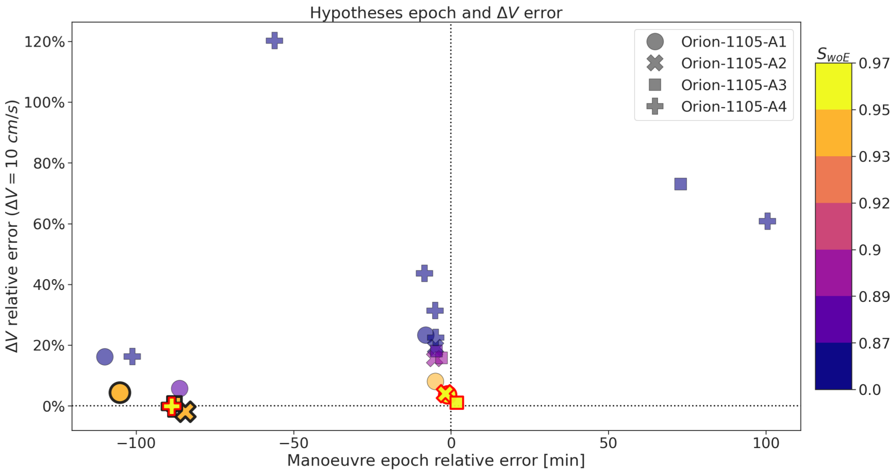

Figure 15. Hypotheses with the corresponding relative errors of the estimated Δv and manoeuvre epochs t_M, for the test case with a satellite spacing of $\Delta\Theta = 5°$.

ticularly in the estimated manoeuvre epoch and magnitude. A comparison between Tab. 11 and the promoted hypotheses in the current scenario clearly demonstrates this degradation, with the most pronounced discrepancy observed in the manoeuvre epoch estimation.

Table 13. Statistics of the performance metrics for promoted hypotheses in the test case with a satellite spacing of $\Delta\Theta = 1°$, $\Delta v = 20\ cm/s$ and three tracks per T2T.

	S	$\varepsilon_{\Delta v}[\%]$	$\varepsilon_{t_M}[min]$	WRMS	TP
A1	0.992	0.35	-0.45	1.116	✓
A2	0.994	3.36	2.02	0.919	✓
A3	0.99	3.35	19.89	1.028	✓
A4	0.946	22.8	-109.42	1.207	✓

5. CONCLUSIONS

This paper has presented three promising strategies developed to improve cataloguing capabilities of tracking and survey. All of them have in common the development of simplified algorithms for propagation to fit operational requirements and their focus on the cataloguing processes. A series of use cases have been defined to prove their validity, specifying the assumptions and their limitations. Thus, future works will be conducted aiming to cover the identified missing points, which include: the use of real instead of simulated data for validation, expand the considered test cases with different scenarios and orbital regimes, and adapt the methodologies to more complex conditions such as low-thrust manoeuvring, which was covered in the sequential filtering method but not in the rest. Nevertheless, the use cases presented in this paper show the potential of these methodologies to improve the reliability, availability and accuracy of the space object catalogues. In fact, these approaches have served as a starting point for the development of more complex solutions within different projects, which try to solve these gaps and are expected to be implemented in the future in operational services.

ACKNOWLEDGEMENTS

This paper is a compilation of the MSc theses of: Emmanuele Nasini at Politecnico di Milano (Section 2), Álvaro Flórez González at Politecnico di Milano (Section 3) and Antonio Malaver at University of Seville (Section 4). These research works have been developed thanks to the cooperation between GMV and the universities.

The authors would like to acknowledge GMV team boosting the projects: D. Escobar-Antón, J. Fernández-Sánchez, C. Fernandez, J. Berzosa-Molina, L. Bao, M. A. Muñoz de la Torre, R. Pastor-Muela, M. Fernández-Usón, S. Lara-Espinosa, M. Salas-Lasala, S. Setty, R. Ströbel, S. Metz, F. Gámez-Losada, B. Bija, K. Stock, M. Callejón-Cantero, O. I. García-Lobo, D. Skoulidou and H. Kaur Mann.

REFERENCES

1. Space Debris Office, Space Environment Statistics, ESA, 06 03 2025. [Online]. [Accessed 11 03 2025]. `https://sdup.esoc.esa.int/discosweb/statistics/`
2. Kelso TS., Celestrak SATCAT, 06 03 2025. [Online]. [Accessed 11 03 2025] `https://celestrak.org/pub/satcat.csv`
3. Montenbruck O. and Gill E., (2000). Satellite Orbits: Models, Methods and Applications, Springer Verlag Berlin Heidelberg, p. 369.

4. Pastor A., Escribano G., Sanjurjo-Rivo M., and Escobar D., (2022) Satellite maneuver detection and estimation with optical survey observations," *The Journal of the Astronautical Sciences* , pp. 879–917. DOI: 10.1007/s40295-022-00311-5.

5. Porcelli L., Pastor A., Cano A., Escribano G., Sanjurjo-Rivo M., Escobar D. and Di Lizia P., (2022) Satellite maneuver detection and estimation with radar survey observations, *Acta Astronautica*

6. Mahalanobis, P. C. (1936), On the generalised distance in statistics *Proceedings of the National Institute of Sciences of India*, **Vol. 2, No. 1** p. 49-55

7. Reihs B., Vananti A., Schildkneckt T., Siminski J., Flohrer T., (2021). Application of attributables to the correlation of surveillance radar measurements *Acta Astronautica* **182**, 399–415.

8. Montilla J.M., Sanchez J.C., Vazquez R., Galan-Vioque J., Rey Benayas J., Siminski J., (2022) Manoeuvre detection in low earth orbit with radar data, *Advances in Space Research*.

9. Pastor A., Sanjurjo-Rivo M., Escobar D., (2022) Track-to-track association methodology for operational surveillance scenarios with radar observations, *CEAS Space Journal*

10. Pastor A., Sanjurjo-Rivo M., Escobar D., (2020), Track-to-track association for space object cataloguing with optical survey data, *71st International Astronautical Congress (IAC)*

11. Liu A., Xu X., Xiong Y. and Yu S., (2024). Maneuver strategies of Starlink satellite based on SpaceX-released ephemeris, *Advances in Space Research*, **74**(7), 3157–3169. DOI: https://doi.org/10.1016/j.asr.2024.06.038

12. Wiesel, W. E., (2003). *Modern Orbit Determination*, Aphelion Press, Beavercreek, OH, USA.

13. Abbot, R. I., Wallace, T. P., (2007). Decision Support in Space Situational Awareness, *Lincoln Laboratory Journal*, **16**(2), 297–335.

14. Gil, A., Segura, J., Temme, N. M., (2015). GammaCHI: A package for the inversion and computation of the gamma and chi-square cumulative distribution functions, *Computer Physics Communications*, **191**, 132-139.

15. McReynolds S.R. (1984). Editing data using sequential smoothing techniques for discrete systems, *AIAA*.

16. ESA Copernicus, (2025). Sentinel 3-A Orbit Description. URL: https://sentiwiki.copernicus.eu/web/s3-mission

17. International Laser Ranging Service (ILRS), (2025). Satellite Maneuvers. URL: https://ilrs.cddis.eosdis.nasa.gov/data_and_products/predictions/maneuver.html

18. Gomez, R., Salmerón, J. M.-V., Besso, P., et al., (2019). Initial operations of the breakthrough Spanish Space Surveillance and Tracking Radar (S3TSR) in the European context, *1st ESA NEO and Debris Detection Conference*.

19. Aguilar Taboada, D., de Juana Gamo, J. M., Righetti, P. L., (2018). Sentinel-3 orbit control strategy, *18th Australian Aerospace Congress*.

20. Kelecy, T., Jah, M., (2010). Detection and orbit determination of a satellite executing low thrust maneuvers, *Acta Astronautica*, **66**(5), 798-809.

21. Rauch H.E., Tung F., Striebel C. T., (1965). Maximum likelihood estimates of linear dynamic systems, *AIAA*.

22. Goff, G. M., (2015). *Orbit Estimation of Non-Cooperative Maneuvering Spacecraft*, PhD Thesis, Air Force Institute of Technology.

23. Malaver A., (2024). Track-to-orbit association problem for non-collaborative targets in the presence of manoeuvres, *Depósito de Investigación de la Universidad de Sevilla*