RODDAS: Robust Orbit Determination and Data Association library for SSA based on GODOT

P. Gago[(1)], L. Porcelli[(1)], A. Cano[(1)], L. Quadri[(1)], A. Pastor[(2)], J. Rubio[(1)], D. Escobar[(1)], Y. Sztamfater[(3)], M. Sanjurjo[(3)], J. Míguez[(3)], A. Narykov[(4)] , L. Devlin[(4)] , S. Maskell[(4)] , L. Vladimirov[(4)], A. García-Fernández[(4)], A. Thompson[(4)], M. Belmekki[(4)], M. Fontana[(4)], A. Phillips[(4)], M. Losacco[(5)], R. Mackenzie[(5)], J. Siminski[(5)]

[(1)] *GMV, Calle Isaac Newton 11, Tres Cantos, 28670, Spain, pau.gago.padreny@gmv.com, lorenzo.porcelli@gmv.com, alcs@gmv.com, lcquadri@gmv.com,* jorge.rubio.anton@gmv.com, *descobar@gmv.com*

[(2)] *GMV, Europaplatz 2 64293 Darmstadt, Germany apastor@gmv.com*

[(3)] *Universidad Carlos III de Madrid, Avenida de la Universidad 30, Leganés, 28911, Spain, ysztamfa@pa.uc3m.es, msanjurj@ing.uc3m.es, jmiguez@ing.uc3m.es*

[(4)] *University of Liverpool, Foundation Building, Brownlow Hill, Liverpool L69 7ZX, United Kingdom, a.narykov@liverpool.ac.uk, lee.devlin@liverpool.ac.uk, s.maskell@liverpool.ac.uk, l.vladimirov@liverpool.ac.uk, angel.garcia-fernandez@liverpool.ac.uk, a.thompson16@liverpool.ac.uk, m.a.a.belmekki@liverpool.ac.uk, marco.fontana@liverpool.ac.uk, a.m.phillips@liverpool.ac.uk*

[(5)] *ESOC, 5 Robert-Bosch Strasse, Darmstadt, Germany, matteo.losacco@ext.esa.int, ruaraidh.mackenzie@esa.int, jan.siminski@esa.int*

ABSTRACT

The current space context heightens the challenges of catalogue build-up and maintenance. The European Space Agency's (ESA) funded project "Robust Orbit Determination for Space Debris", led by GMV in collaboration with the University of Liverpool (UoL) and the University Carlos III de Madrid (UC3M), puts forward novel methodologies to tackle these issues. This paper focuses on the most relevant contributions of the project. The most promising pre-prototyped algorithms are embedded into "RODDAS", a library for Robust Orbit Determination and Data Association using ESA's GODOT as underlying astrodynamics engine and Stone Soup's library as the foundation for the tracking and estimation algorithms.

1 INTRODUCTION

Catalogue build-up and maintenance are fundamental to Space Situational Awareness (SSA) services, ensuring updated databases of all observed space objects, including operational satellites or space debris, with their current and predicted orbit estimates. The accuracy and reliability of the catalogue is crucial for the provision of other SSA services such as conjunction assessment, fragmentation detections, re-entry predictions, and manoeuvre detection, all essential for the safety and sustainability of the space environment.

Space democratisation and the rise of commercial activities, exacerbates this situation. Associating new observations with existing objects becomes ambiguous in overly congested regions, which can lead to poor track association performance. This strains the reliability of the subsequent processes of orbit determination and propagation. Poor catalogue estimation and prediction further deteriorate data association performance, aggravating data scarcity and ultimately impacting SSA services. Therefore, it is needed to analyse and develop robust and reliable methodologies for orbit determination and data association to address current and future challenges of the space environment sustainability.

GMV in collaboration with UoL and UC3M explores novel approaches to mitigate this situation in ESA's funded project "Robust Orbit Determination for Space Debris".

On the one hand, modern extensions to state-of-the-art batch least-squares estimation algorithms designed to improve robustness will be presented. These include the Huber penalty function for least-squares estimation to enhance convergence, time-dependent measurement weight to achieve smoother Lengths of Update Intervals (LUPI), and uncertainty quantification techniques for covariance realism improvement, including Stochastic Consider Parameters and alternative state representations for uncertainty assessment.

On the other hand, the performance of less established algorithms (in the context of space object tracking) is evaluated. Namely, particle and sequential filters. Regarding sequential algorithms, we explore the family of methods that stems from the tuning the hyperparameters of the Iterated Posterior Linearization Filter (IPLF), going one step beyond EKF or sigma-point Kalman filters, and taking advantage of the knowledge of the value of the measurement and of statistical linear regression for an improved update. Furthermore, the benefit of introducing a smoother (IPLS), taking advantage of future known estimates to smooth updates in the past, is assessed. For particle filters, the standard Particle Filter (PF), Ensemble Kalman Filter (EnKF) and

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)

Nested Particle Filter (NPF) are selected for analysis in the very same scenarios as the other algorithms detailed above.

In the context of the multi-target data association problem, novel approaches making use of sequential filters together with Global Nearest Neighbour and Joint Probabilistic Data Association Filter associators are investigated. The performances are compared to those of the track-to-track association methods [21], which is an association algorithm specifically tailored for SSA operations.

All these algorithms are tested in situations representing operational environments.

2 METHODS

This section details the methodologies that have been developed as part of the activity. Those applicable to state and parameter estimation are described in Section 2.1, while those concerning the data association problem are tackled in Section 2.2.

2.1 Robust Orbit Determination

Traditionally, the preferred approaches for state and parameter estimation (also referred to as Orbit Determination, OD) in space have been batch least squares and Kalman filters. Modern extensions to these are sought. Additionally, alternative methods popular in the tracking community, in the realm of sequential and particle filters, are surveyed and their applicability to the space domain is assessed.

2.1.1 Modern extensions to batch least-squares

Four different concepts are pre-prototyped, aiming to improve traditional batch least-squares methods performance in terms of uncertainty quantification, convergence capabilities, state representation to ensure Gaussianity for extended time intervals, and optimum Length of Update Interval (LUPI) through time-dependent weighting of the measurements.

2.1.1.1 Uncertainty quantification: Stochastic Consider Parameters

Complex representations of a space object's orbital state, beyond mean and second moment, are nowadays still not a feasible approach for SSA activities. Therefore, a significant effort shall be placed into Uncertainty Quantification, which reduces to Covariance Realism in the usual context where only up to the second moment of the state Probability Density Function (PDF) is retained.

In LEO, the most significant uncertainty source comes from the difficulty of modelling and predicting Solar activity and the interaction between the space object and Earth's atmosphere [1, 2, 3]. Typical strategies to deal with its aleatoric behaviour, such as stochastic models that result in complex data processing systems and simple Monte Carlo methods with poor computational performance, deem unsuitable for catalogue maintenance activities. Additionally, classical methods such as Consider Parameter theory fail to capture the stochastic nature of atmospheric density uncertainty as it introduces a single fixed variance to model it.

Consequently, the Stochastic Consider Parameter (SCP) theory is proposed, allowing to introduce stochastic time-correlated errors to model the uncertainty in the problem. To that end, the latter is modelled as an auto-regressive function of order 1 AR(1), governed by an unknown noise power and correlation time scale, such that each realization of the stochastic noise is defined as an uncertain consider parameter. Then, the contribution of each stochastic parameter into the estimated covariance is mapped by exploiting the properties of the variational equations.

As is the case in Consider Parameter theory, realistic values for the governing parameters of the stochastic model, in this case noise power and correlation time scale, are not known. Therefore, this methodology is complemented with covariance determination techniques [4, 5], based on the observed distribution of the Mahalanobis distance of the orbital differences (between predicted and estimated orbits) and Empirical Distribution Function (EDF) metrics such as the Cramer-von-Mises (CvM) and the Kolmogorov-Smirnov (KS) distances to determine optimum noise power and correlation time scales that ensure covariance realism.

For a batch estimation process where n_y parameters are being estimated, consider parameter theory allows to include any unaccounted uncertainty (n_c consider parameters) into the estimated covariance:

$$\mathbf{P}_c = \mathbf{P}_n + \mathbf{K}\mathbf{C}\mathbf{K}^T \in \mathbb{R}^{n_y \times n_y}, \tag{1}$$

where $\mathbf{P}_n$ is the so-called noise-only covariance coming from batch estimation, and $\mathbf{K}, \mathbf{C}$ are defined as:

$$\mathbf{K} = \mathbf{P}_n\left(\mathbf{H}_y^T \mathbf{W} \mathbf{H}_c\right) \in \mathbb{R}^{n_y \times n_c}, \tag{2}$$

$$\mathbf{C} = \begin{pmatrix} \sigma_1^2 & \cdots & 0 \\ \vdots & \ddots & \vdots \\ 0 & \cdots & \sigma_{n_c}^2 \end{pmatrix}, \tag{3}$$

where $\mathbf{H}_y \in \mathbb{R}^{n_m \times n_y}$, with n_m being the number of measurements, represents the partials of the measurements with respect to the estimated state, $\boldsymbol{W}$ weights according to the expected accuracy of the measurements and $\mathbf{H}_c \in \mathbb{R}^{n_m \times n_c}$ represents the partials of the measurements with respect to the consider parameters. Thus, the i^{th} row of matrix H_c can be constructed as the product of the partial of the measurement with respect to the state at measurement epoch, times the partial of the state with respect to the consider parameter, which is the so-called sensitivity matrix of the consider parameter, obtained through the variational parameters.

$$\mathbf{H_{ci}} = \frac{\partial \mathbf{h_i}(\boldsymbol{t_i})}{\partial \mathbf{c}} = \frac{\partial \mathbf{h_i}(\boldsymbol{t_i})}{\partial \mathbf{x}(\boldsymbol{t_i})} \cdot \frac{\partial \mathbf{x}(\boldsymbol{t_i})}{\partial \mathbf{c}} \in \mathbb{R}^{1 \times n_c} \tag{4}$$

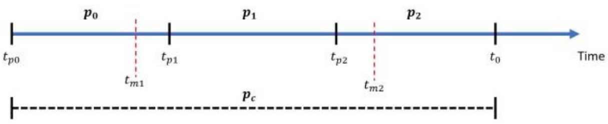

Figure 2-1. Multiple constant consider parameters

If we split a single fixed parameter into multiple parameters through the estimation arc, which while representing the same physical variable are allowed to have different value, as illustrated in Figure 2-1, it can be shown from the properties of the variational equations that the partial derivatives of the state with respect to such parameters can be derived from the partials of a single global parameter p_c:

$$\frac{\partial \mathbf{x}(t)}{\partial p_i} = \begin{cases} \dfrac{\partial \mathbf{x}(t)}{\partial p_c} - \dfrac{\partial \mathbf{x}(t)}{\partial \mathbf{x}(t_{pi})} \dfrac{\partial \mathbf{x}(t_{pi})}{\partial p_c} & t_{pi} < t <= t_{pi+1} \\ \dfrac{\partial \mathbf{x}(t)}{\partial \mathbf{x}(t_{pi+1})} \left(\dfrac{\partial \mathbf{x}(t_{pi+1})}{\partial p_i} \right) & t > t_{pi+1} \\ 0 & t <= t_{pi} \end{cases} \tag{5}$$

where the contribution of each parameter depends on the time-interval of application of the parameter (active from t_{pi} to t_{pi+1}).

Atmospheric uncertainty is then modelled by considering stochastic atmospheric density of the form of

$$\rho(t) = \bar{\rho}(t) + p(t) \tag{6}$$

where the perturbing noise $p(t)$ added to the mean atmospheric density is a zero-mean correlated noise sequence, an auto-regressive function with time correlation of order 1 AR(1):

$$\boldsymbol{p}(\boldsymbol{t_n}) = \boldsymbol{a}(\boldsymbol{n})\boldsymbol{p}(\boldsymbol{t_{n-1}}) + \boldsymbol{u}(\boldsymbol{n}), \tag{7}$$

with

$$\boldsymbol{a}(\boldsymbol{n}) = \frac{\boldsymbol{r_0}(\boldsymbol{t_n})}{\boldsymbol{r_0}(\boldsymbol{t_{n-1}})} \boldsymbol{e}^{-\boldsymbol{\alpha}(\boldsymbol{t_n} - \boldsymbol{t_{n-1}})}, \tag{8}$$

$$\boldsymbol{u}(\boldsymbol{n}) \sim \mathbf{N}\left(\mathbf{0}, \boldsymbol{\sigma_u^2}(\boldsymbol{n})\right), \tag{9}$$

$$\sigma_u^2(n) = r_0(t_n)\left[1 - \frac{r_0(t_n)}{r_0(t_{n-1})} e^{-2\alpha(t_n - t_{n-1})}\right], \text{and} \tag{10}$$

$$p(t_0) = u(0) \sim \mathrm{N}\big(0, r_0(t_0)\big) \tag{11}$$

Accordingly, the noise at each step is related to the previous noise via a factor $a(n)$, and a Gaussian contribution $u(n)$. The correlation strength is controlled by $\alpha = 1/\tau_\alpha$. The variance of the Gaussian term is inversely proportional to the correlation. Thus, constant power noise $(r_0(t_n) = r_0(t_{n-1}))$ is achieved when the correlation time scale tends to infinity, hence being reduced to a constant noise model. On the contrary, a purely Gaussian noise is obtained when the correlation time scale tends to zero.

Therefore, the proposed approach is to divide a global parameter into multiple parameters, applied sequentially and correlated in time according to the stochastic model defined. We define $u(n)$ of Eq. (7) as the stochastic consider parameters, such that:

$$\boldsymbol{p} = \boldsymbol{A}\boldsymbol{u}, \tag{12}$$

where:

$$\boldsymbol{p} = \begin{pmatrix} p_0 \\ \vdots \\ p_n \end{pmatrix}; \boldsymbol{u} = \begin{pmatrix} u_0 \\ \vdots \\ u_n \end{pmatrix}; \qquad \boldsymbol{A} = \begin{pmatrix} 1 & \cdots & 0 \\ \vdots & \ddots & \vdots \\ b_{ij} & \cdots & 1 \end{pmatrix} \in \mathbb{R}^{N \times N},$$

with a total of $N = n + 1$ consider parameters for each of the n time steps in the noise sequence. Finally, $b_{ij} = \prod_{j+1}^{i} a\,(n); \forall\; i > j$ where i and j represent the row and column of the matrix, respectively. Each value of vector u corresponds to a sample of $\boldsymbol{u} \sim (0, \sigma_u^2)$, thus being a suitable choice of consider parameter vector.

Thus, the sensitivity of each consider parameter in Eq. (4) can be formulated as follows:

$$\frac{\partial \mathbf{x}(\boldsymbol{t})}{\partial \mathbf{c}} = \frac{\partial \mathbf{x}(\boldsymbol{t})}{\partial \boldsymbol{u}} = \frac{\partial \mathbf{x}(\boldsymbol{t})}{\partial \boldsymbol{p}} \frac{\partial \boldsymbol{p}}{\partial \boldsymbol{u}} = \frac{\partial \mathbf{x}(\boldsymbol{t})}{\partial \boldsymbol{p}} \mathbf{A} \tag{13}$$

such that for the case of multiple parameters, Eq. (4) reduces to:

$$\frac{\partial \mathbf{x}(\boldsymbol{t_i})}{\partial \mathbf{c}} = \left[\frac{\partial \mathbf{x}(\boldsymbol{t_i})}{\partial \boldsymbol{p_i}}, \cdots, \frac{\partial \mathbf{x}(\boldsymbol{t_i})}{\partial \boldsymbol{p_{n_c}}}\right] \cdot \mathbf{A} \in \mathbb{R}^{\boldsymbol{n_x} \times \boldsymbol{n_c}} \tag{14}$$

and the covariance with SCP formulation takes the terms $\boldsymbol{C}$ and $\boldsymbol{H_c}$ of the form:

$$\boldsymbol{C} == \begin{pmatrix} r_0 & \cdots & \cdots & 0 \\ \vdots & \sigma_u^2 & & \vdots \\ \vdots & & \ddots & \vdots \\ 0 & \cdots & \cdots & \sigma_u^2 \end{pmatrix}, \in \mathbb{R}^{n+1 \times n+1} \tag{15}$$

$$\mathbf{H}_{CA}(m_i, t_m) = \frac{\partial \mathbf{h_i}(t_m)}{\partial \mathbf{x}(t_m)} \cdot \frac{\partial \mathbf{x}(t_m)}{\partial \mathbf{p}} \cdot \boldsymbol{A} \tag{16}$$

2.1.1.2 QtW reference frame

The methodology presented in the previous section aims at improving the realism of the state distribution of our space object, which is eventually represented by a mean state and a covariance. To delay the state's distribution departure from Gaussianity for long propagation arcs due to the highly non-linear dynamics of the space environment, especially in low altitude orbits, the QtW reference frame to represent the object's state is proposed [6].

It consists of a non-linear local orbital, time-dependent transformation based on the QSW local orbital frame, with "Q" pointing towards the centre of the Earth, and "W" following the orbital's angular momentum. It is centred in the space object of interest, and applied, in this case, to a given sample of the space object's distribution, such that:

1. The object's state is frozen at a given epoch of interest t_0, and the QW plane defined
2. A sample drawn from the object's distribution is propagated from t_0 until it crosses the QW plane. The crossing epoch is named t_c
3. The QtW coordinates of the sample are then defined by its coordinates in the QW plane at the crossing epoch and the propagation time needed to reach this plane $t_c - t_0$

Thus, the coordinates of a given sample can be expressed as:

$$\boldsymbol{d}\big(t, \boldsymbol{x}(t), \boldsymbol{x}_{\boldsymbol{r}}(t)\big)\big|_{\mathrm{QtW}} = \begin{pmatrix} \boldsymbol{q} \cdot \big(\boldsymbol{x}(t_c) - \boldsymbol{x}_{\boldsymbol{r}}(t)\big) \\ t_c - t \\ \boldsymbol{w} \cdot \big(\boldsymbol{x}(t_c) - \boldsymbol{x}_{\boldsymbol{r}}(t)\big) \end{pmatrix} \tag{17}$$

where $\boldsymbol{x}_{\boldsymbol{r}}$ and $\boldsymbol{x}$ are the object and sample states, and $\boldsymbol{q}, \boldsymbol{w}$ define the QW local frame. The crossing time t_c can be determined solving $\big(\boldsymbol{x}(t_c) - \boldsymbol{x}_{\boldsymbol{r}}(t)\big) \cdot \boldsymbol{s} = \boldsymbol{0}$, from the knowledge of the normal vector $\boldsymbol{s}$.

Then, to project the covariance onto the QtW frame, under the consideration that, even in an inertial reference this is being computed assuming linear theory, one might assume:

- Linearity in the relative position between the reference point and the QtW sample, i.e. $\Delta\boldsymbol{x} \approx \boldsymbol{x}(t) - \boldsymbol{x}_{\boldsymbol{r}}(t)$.
- Linearity in the time difference between analysis and crossing epoch. If $\Delta t = t_c - t \approx 0$, then the State Transition Matrix (STM) is close to unity $\boldsymbol{\Phi}(t, t_c)|_{\mathrm{xyz}} \approx \boldsymbol{I}$.

Accordingly, the Jacobian for the position of a sample in QtW can be expressed as:

$$\boldsymbol{J}(\boldsymbol{d}(t))|_{\mathrm{xyz}\to\mathrm{QtW}} = \frac{\boldsymbol{d}(t, \boldsymbol{x}, \boldsymbol{x}_{\boldsymbol{r}} + \Delta\boldsymbol{x})|_{\mathrm{QtW}} - \boldsymbol{d}(t, \boldsymbol{x}_{\boldsymbol{c}}, \boldsymbol{x}_{\boldsymbol{r}})|_{\mathrm{QtW}}}{\Delta x} = \frac{1}{\Delta x}\begin{pmatrix} \boldsymbol{q} \cdot \Delta\boldsymbol{x} \\ \Delta t(\boldsymbol{x}_{\boldsymbol{r}} + \Delta\boldsymbol{x}) - \Delta t(\boldsymbol{x}_{\boldsymbol{r}}) \\ \boldsymbol{w} \cdot \Delta\boldsymbol{x} \end{pmatrix} \tag{18}$$

Under the linear assumption $\boldsymbol{x}(t_c) = \boldsymbol{x}_{\boldsymbol{r}}(t) + \boldsymbol{v}_{\boldsymbol{r}}(t)\Delta t$, and using the crossing plane equation:

$$\boldsymbol{J}(\boldsymbol{d}(t))|_{\mathrm{xyz}\to\mathrm{QtW}} == \begin{pmatrix} \boldsymbol{q} \\ \dfrac{-\boldsymbol{s}}{\boldsymbol{v}_{\boldsymbol{r}}(t) \cdot \boldsymbol{s}} \\ \boldsymbol{w} \end{pmatrix} \tag{19}$$

which allows to rotate the covariance of the samples as a function only of the reference state. Thus, combining the QtW transformation with a linear propagation of the covariance, we can evaluate it in such a frame through the following expression:

$$\begin{aligned} &\boldsymbol{P}(t)|_{\mathrm{QtW}} \\ &= \boldsymbol{J}(t)|_{\mathrm{xyz}\to\mathrm{QtW}} \times \boldsymbol{P}(t)|_{\mathrm{XYZ}} \times \boldsymbol{J}(t)|^{T}{}_{\mathrm{xyz}\to\mathrm{QtW}} \\ &= \boldsymbol{J}(t)|_{\mathrm{xyz}\to\mathrm{QtW}} \\ &\times \big(\boldsymbol{\Phi}(t, t_0)|_{\mathrm{xyz}} \boldsymbol{P}(t_0)|_{\mathrm{XYZ}} \boldsymbol{\Phi}(t, t_0)|_{\mathrm{xyz}}\big) \\ &\times \boldsymbol{J}(t)|^{T}{}_{\mathrm{xyz}\to\mathrm{QtW}} \end{aligned} \tag{20}$$

2.1.1.3 Huber penalty parameters

The classical approach to solve the batch least-squares problem comes from applying Newton's method to the second order Taylor expansion of the loss function minimising residuals $\boldsymbol{\rho}$ (Gauss-Newton's method):

$$J(\boldsymbol{x}_0) = \boldsymbol{\rho}^T\boldsymbol{\rho} = (\boldsymbol{z} - \hat{\boldsymbol{z}})^T(\boldsymbol{z} - \hat{\boldsymbol{z}}) \tag{21}$$

where the estimated measurement $\hat{\boldsymbol{z}} = \boldsymbol{h}(\boldsymbol{x}_{\boldsymbol{0}})$ is built from the estimated state $\hat{\boldsymbol{x}}$ and the measurement model $\boldsymbol{h}()$ mapping the estimated state $\boldsymbol{x}_{\boldsymbol{0}}$. In this way, the solution for $\boldsymbol{x}_{\boldsymbol{0}}$ is iterated according to [7]:

$$(\boldsymbol{H}^T\boldsymbol{H})\Delta\hat{\boldsymbol{x}}_0 = \boldsymbol{H}^T\Delta\boldsymbol{z} \tag{22}$$

$\boldsymbol{H}$ is the Jacobian of the measurement model, where second-order information terms are neglected.

This iterative scheme is locally q-linearly convergent if second-order terms are small relative to the first order contribution, which is reasonable if residuals are small.

Nonetheless, if the latter is not the case, the method may not be locally convergent or can take long steps even if in the correct direction. Additionally, the classical least-squares solution is optimal under the assumption of Gaussian noise, which is not the expected case for measurement noise or model errors.

To improve convergence and robustness of the iteration, the expression for the loss function can be reformulated using Wilson's Huber penalty parameters [8]:

$$J(\boldsymbol{x}_0) = \phi(\boldsymbol{z} - \hat{\boldsymbol{z}}) + \lambda\psi\big(\boldsymbol{x} - \boldsymbol{x}_{ref}\big) \tag{23}$$

ϕ is a robust penalty function for the residuals, ψ is a penalty function for the estimate update and λ is the trust-region weight. This way, the loss function transforms into a balance between maximum likelihood of a solution fitting the data $(\boldsymbol{z} - \hat{\boldsymbol{z}})$ and ensuring that linearisation of the dynamics and observation models remain a valid approximation $\boldsymbol{x} - \boldsymbol{x}_{ref}$.

By selecting ϕ and ψ to be convex, the system can be solved by subsequent linearisation after each iteration around the current solution $\boldsymbol{x}_{ref} = \boldsymbol{x}_k$ for iteration k+1. The residual penalty function takes the following form:

$$\phi(\rho_i) = \begin{cases} \rho_i^2 & if\ |\rho_i| \leq M \\ M(2|\rho_i| - M) & if\ |\rho_i| > M \end{cases} \quad (24)$$

where M (defined as the Huber penalty parameter) is a threshold on the size of the residuals, chosen beyond which Gaussian noise no longer holds, and the measurement is heavily penalized. Selecting ψ as a Euclidean norm penalty, taking care to map the value to the expected value of the measurement residuals for consistent consideration in the cost function, one produces:

$$\psi(\boldsymbol{x} - \boldsymbol{x}_{ref}) = \left\| (\boldsymbol{H}^T\boldsymbol{H})^{1/2}(\boldsymbol{x} - \boldsymbol{x}_{ref}) \right\|^2 \quad (25)$$

For a given iteration, constructing a set A with the measurements below the non-linear threshold, M, and B with those above it

$$J(\Delta \boldsymbol{x}_0) = \boldsymbol{\rho}_A \boldsymbol{\rho}_A^T + \sum_{\rho_i \in B} M(2|\rho_i| - M) + \lambda \|\boldsymbol{L}\, \Delta \boldsymbol{x}_0\|^2 \quad (26)$$

where $\boldsymbol{L} = \boldsymbol{H}^T\boldsymbol{H}$.

Equating the gradient of the loss function to zero, the solution can be iterated according to:

$$\left(\boldsymbol{H}_A^T\boldsymbol{H}_A + M\widehat{\boldsymbol{H}}_B^T\widehat{\boldsymbol{H}}_B + \lambda\, \boldsymbol{L}^T\boldsymbol{L}\right)\boldsymbol{\Delta}\widehat{\boldsymbol{x}}_0 = (\boldsymbol{H}_A^T + \widehat{\boldsymbol{H}}_B^T)\boldsymbol{\Delta z} \quad (27)$$

The problem is then a-dimensionalised by additionally weighting with the expected noise of the measurements ρ_i/σ_i.

Concerning the threshold M, it is proposed to use a fixed value. The trust region weight λ is inherited from the Levenverg-Marquardt (LM) method [9], a precursor to Wilson's Huber parameters. The rationale for computing λ follows the rationale of [8]:

- Evaluate predicted improvement: $\hat{\delta} = \sum_i \phi_{hub}(\rho_{i,k-1}) - \sum_i \phi_{hub}(z_i - h_{i,k-1}(\boldsymbol{x}_k))$
- Evaluate observed improvement: $\delta = \sum_i \phi_{hub}(\rho_{i,k-1}) - \sum_i \phi_{hub}(\rho_{i,k})$
- If $\delta \geq \alpha\hat{\delta}$, then $\lambda = \lambda/\beta_{succ}$, else $\lambda = \lambda/\beta_{fail}$.

where values for the factors α, β are to be derived in benchmarking.

2.1.1.4 Time-dependent measurements weight

In batch least-squares method, when selecting an orbit determination interval, or length of update interval (LUPI), one shall balance the number of measurements considered with the ability of the dynamical models to hold valid for extended time intervals.

To determine the optimum LUPI, it is proposed to adopt a time-dependent measurement weight function. It is designed to have unitary value at the epoch of the last measurement t_0, and decay exponentially going backwards following a certain time scale $w_i(t_m) = e^{-\frac{1}{\tau}(t_0 - t_m)}$. This becomes the control variable, interchangeable with the LUPI itself as its value will determine the moment in time in which measurements will have a null contribution. The value of this parameter is determined as follows. First, covariance matrices from Orbit Determination processes are assumed to be realistic after application of Covariance Determination methodologies such as the SCP described in Section 2.1.1.1. Then, the time scale is selected such that it minimises the predicted covariance of the state, while remaining realistic. The weighting matrix for orbit determination, containing expected noise levels for each measurement is modified acordingly:

$$\boldsymbol{W}^*(\boldsymbol{t_m}) = \begin{pmatrix} \sigma_i^{-2} w_i(t_{m_i}) & \cdots & 0 \\ \vdots & \ddots & \vdots \\ 0 & \cdots & \sigma_N^{-2} w_i(t_{m_N}) \end{pmatrix} \quad (28)$$

This affects directly the noise only covariance $\mathbf{P}_n = \left(\mathbf{H}_y^T\mathbf{W}^*\mathbf{H}_y\right)^{-1}$ and impacts covariance realism through the SCP methodology.

2.1.2 Modern sequential filters

The modern sequential filters described in this section are implemented making use of the **Stone Soup's** [10] open-source tracking software library, developed by several institutions in the Five Eyes nations, led by DSTL and including UoL, contributor to this work.

The dynamic and measurement models considered are non-linear function with additive Gaussian noise:

$$p(\, x_k \mid x_{k-1}\,) = \mathcal{N}(x_k; f_{k-1}(x_{k-1}), Q_{k-1}) \quad (29)$$

$$p(\, z_k \mid x_k\,) = \mathcal{N}(z_k; h_k(x_k), R_k) \quad (30)$$

where f_{k-1} and h_k are representing the dynamics and measurement functions and Q_{k-1} and R_k are the covariance matrices of the dynamics and measurement noise respectively.

Additionally, it is considered that the predicted and filtered densities are Gaussian, such that:

$$p(x_k|z_{1:k}) = N(x_k, \overline{x}_{k|k}, P_{k|k}) \quad (31)$$

$$p(x_k|z_{1:k-1}) = N(x_k, \overline{x}_{k|k-1}, P_{k|k-1}) \quad (32)$$

where $\overline{x}_{k|k}$ is the mean and $P_{k|k}$ is the covariance matrix.

Depending on the filter, the prediction step from $k-1$ to k and subsequent update of the estimate is performed differently. Several filters are explored in the activity.

A first explored candidate is the Iterated Extended Kalman Filter (IEKF) [10, 12]. Differently to EKF and sigma-point Kalman filters, it takes advantage of the knowledge of the measurement $\boldsymbol{z}_k$ for the update of the posterior distribution, meaning that in non-linear settings it will be inaccurate relative to the measurement noise covariance [13]. These methods, at each update step, approximate the non-linear measurement function $\boldsymbol{h}_k(x) \approx \boldsymbol{H}_k^+\boldsymbol{x} + \boldsymbol{b}_k^+ + \boldsymbol{e}_k$ where $\boldsymbol{e}_k$ is a zero-mean independent Gaussian noise with covariance matrix $\boldsymbol{\Omega}_k^+$. This choice of $\boldsymbol{h}()$ is the only for which the updated density is exactly Gaussian.

However, the IEKF sets the term $\boldsymbol{\Omega}_k^+$ to zero, which is not necessary. Derivation of optimal values for $\boldsymbol{H}_k^+, \boldsymbol{b}_k^+, \boldsymbol{\Omega}_k^+$ is the foundation of the Iterated Posterior Linearisation Filter (IPLF), which performs statistical linear regression (SLR) with respect to the best available approximation of the posterior (predicted), iterating on it, to achieve a better approximation of it. SLR approximates as linear a non-linear function (measurement model) in the area indicated by a probability density (predicted posterior).

The prediction step of the IPLF is analogous to that of the Unscented Kalman Filter and Cubature Kalman Filter [10], while the update is based on iterated SLR with respect to the best available approximation of the posterior distribution.

Alongside the IPLF, the Iterated Posterior Linearisation Smoother (IPLS) [10, 14, 15], adds additional linearisation of the dynamics and measurement functions in past updates. It makes use of the RTS smoother, a closed-form forward backward recursion to calculate the smoothed densities at each time step with linear (affine) Gaussian measurement and dynamic models. A standard sigma-point RTS smoother is implemented, in which initial values for the smoothed mean and covariance are achieved by applying a forward sigma-point filter, and the proceeding to the backwards smoother.

2.1.3 Particle filters

The particle filters investigated in this activity for application to the space environment are grouped into state tracking and joint parameter estimation and state tracking schemes.

Firstly, a standard particle filter with clipping, in which:

- Resampling steps are taken whenever the effective sample size (ESS) is below a threshold [16].
- The importance weights are clipped, making the N_c largest weights out of the N weights equal to the N_c weight [17].

Secondly, a 2-timescale particle filter, in which the model parameters are not constant. The evolution of this parameters however is slow with respect to state variables. Hence, two timescales are introduced.

Thirdly, an Ensemble Kalman Filter (EnKF) [18], in which particle weights are updated using a Monte Carlo estimate of the Kalman gain and can be used for both state and parameter tracking.

Ultimately, a Nested Particle Filter (NPF) [19, 20] which tackles the recursive computation of the posterior distribution of the parameters. The idea behind it is to approximate sampling by a jittering procedure that introduces a controlled perturbation to the (parameter) particles. Then, the likelihood $p\left(\boldsymbol{z}_k \middle| \bar{\boldsymbol{\theta}}_k^i, \boldsymbol{z}_{1:k-1}\right)$ is approximated by taking one time step of a standard particle filter with clipping conditional on the jittered parameter sample.

2.2 Robust Data Association

2.2.1 Multi-target Data Association

The **Stone Soup's** software library is again used as the foundation for the development of multi-target data association algorithms applied to the space environment.

Considering the notion that there are N_k targets at the k time epoch, with state $X_k = [x_k^1 \dots x_k^{N_k}]$ and $A_k = \left[a_k^1 \dots a_k^{N_k}\right]$ are the data association hypotheses for each target.

Several algorithms are investigated to find the value A_k which maximises $p(Z_k, A_k | Z_{1:k-1})$, where again Z_k refers to observation k.

To that end, the Global Nearest Neighbour (GNN) defines a cost matrix C, of dimensions $N_k \times M_k + 1$, where M_k is the number of measurements in Z_k, with value $C_{i,j} = \log\left(p\left(Z_k, a_k^i = j \middle| Z_{1:k-1}\right)\right)$. The association that maximises the joint probability is then identified using the standard auction iterative algorithm, which scales linearly with the number of targets. Once the optimal value for A_k is identified, the single-target associations that comprise the joint association are assumed to have a probability of unity and all others a probability of zero.

The Joint Probabilistic Data Association Filter (JPDA), on the other hand, considers an (exponentially large) sum over the allowed joint associations, and does so through an implementation of the Efficient Hypothesis Management (EHM) algorithm. EHM assumes access to the same matrix of benefits C and employs two non-iterative processes: firstly the EHM net is constructed, being a data structure describing the exponentially large number of terms in the sum with substantially sub-exponential computation cost and memory requirements; then, it capitalises on this data structure to compute a revised C which is used in the ongoing operation of the multi-target tracking algorithm.

The first process involves several steps. Firstly, it assumes access to the set of measurement hypotheses that each target can be associated to, to then recurse through the targets in reverse order to identify the set of remaining measurement hypothesis which targets after the i^{th} target can associate with.

Secondly, the net is built layer-by-layer. A layer per target with a non-zero number of nodes. Each node is associated with a set of measurement hypotheses that are relevant to targets that have yet to be considered.

Targets are considered in turn. For each target, each node in the layer is considered for each non-zero measurement hypothesis and for the non-detection measurement hypothesis. For each such association, the set of measurements that have been used and are relevant to the targets that are yet to be considered, is computed. If no node exists, it is created and an arc added between the node i^{th} and $(i-1)^{th}$, annotated with the association. Multiple arcs can then exist between a pair of nodes, and the resulting data structure captures all the information that would have existed had one enumerated the exponentially large number of joint data association hypotheses. The performance of the algorithm is clearly impacted by the ordering of the targets.

The second process comprising EHM involves calculating the marginal probabilities that can be passed back to the single-target tracking components. For each target, this requires EHM to calculate a sum over all constituent joint association hypotheses.

3 RESULTS

Several benchmarking scenarios are derived to assess the performance of the aforementioned methodologies. GMV's **SST Sensor Data Simulator (*Ssdsim*)** is used for the simulation of object populations and observation data. GMV's **SST Orbit Determination and Sensor Calibration Software (*Sstod*)** is used to generate reference weighted batch least-squares solutions.

ESA's **GODOT** software library is used as astrodynamics engine for the necessary orbit propagation activities around which the algorithms are implemented.

Several scenarios are simulated. For state estimation, these include nominal OD scenarios, scenarios which include miscorrelated observations, others with data scarcity, with high noise distribution and others with initial state derived from an Initial Orbit Determination (IOD) algorithm. These are simulated both in Loe Earth Orbit (LEO) and Geostationary Orbit (GEO) regimes.

For data association, nominal scenarios in LEO and GEO for 20 objects and with data for 7 and 20 days are simulated together with a fragmentation case in LEO, a cluster of 6 objects in GEO, and a case with nearby objects in GEO.

Firstly, however, results for the pre-prototypes concerning modern extension to batch least-square methods will be presented.

3.1 Modern batch least-squares methods

3.1.1 Uncertainty quantification: Stochastic Consider Parameters

To assess the performance of the SCP pre-prototype, 3 months of data of an object in Sentinel-1 orbit is used to generate a total of 40 different estimated orbits, every 2/3 days depending on data availability. The estimates are propagated from estimation epoch t_0 (last measurement) to +4 and +7 days after, in which an uncorrelated (without any data reuse in common) orbit estimate is also produced. With these two orbits (predicted and estimated), the CvM EDF test metric, assessing the distribution of the Mahalanobis distances between the orbits, according to the predicted covariance, is evaluated and displayed in *Figure 3-1*.

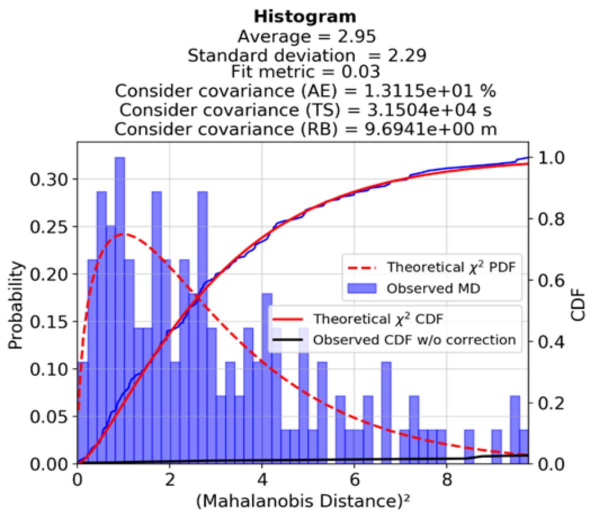

Figure 3-1. Mahalanobis distance distribution of orbital differences between predicted and estimated orbits between t_0+4 to t_0+7, after SCP update for covariance realism.

The observed close match between the observed and the theoretical 3 DoF χ^2 distribution is confirmed with a CvM EDF Fit metric of 0.03. The estimated SCP are in line with the noise introduced in the simulated data, reaching a drag model error power noise of 13.1%, with a correlation time scale of ~0.5 days. The presence of the black line (without correction) highlights the importance of accounting for model errors in the estimation of the covariance.

Additionally, classical covariance containment analysis reveals a 97.14% containment, very much in line with the theoretical containment for a 3 DoF Gaussian (97.1%). Sample containment, in green, is illustrated in *Figure 3-2*.

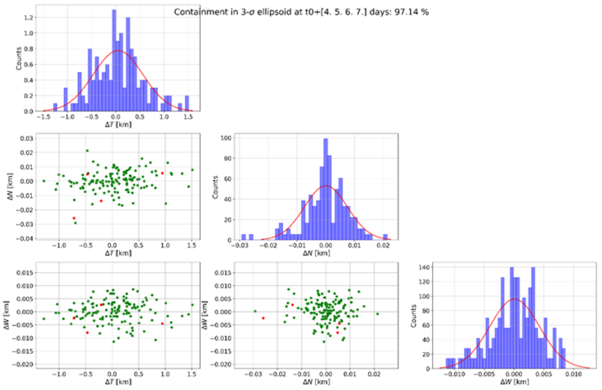

Figure 3-2. 3σ covariance containment of the optimised Mahalanobis distance at various analysis epochs in TNW

3.1.2 QtW reference frame

Three different objects, at 300, 500 and 800km of altitude are chosen to assess the performance of the QtW frame. For each, the state and drag coefficient are Monte Carlo sampled (the latter with a 20% standard deviation). High fidelity dynamics are used to propagate each sample.

At the final epoch, the reconstructed covariance matrices in inertial (GCRF), TNW and QtW frames are evaluated, together with the state's particle distribution.

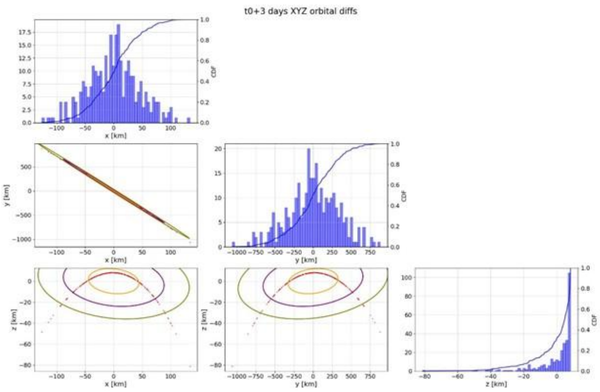

Figure 3-3. MC samples dispersion in GCRF frame after 3 days of propagation. Low altitude object.

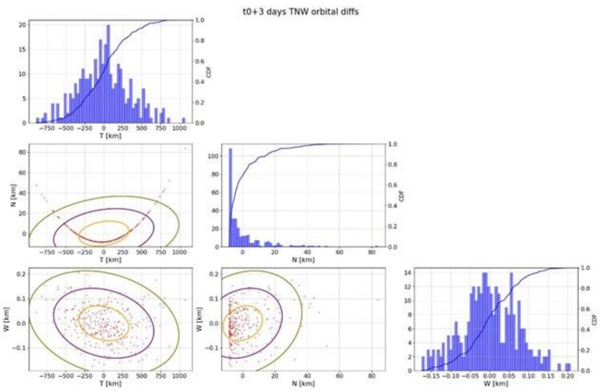

Figure 3-4. MC samples dispersion in TNW frame after 3 days of propagation. Low altitude object.

For the lowest altitude object, samples in GCRF and TNW frame, together with the reconstructed covariance after 3 days of propagation are illustrated in *Figure 3-3* and *Figure 3-4* respectively:

Instead, in QtW frame, along-track dispersion is translated into time dispersion, allowing to retain a Gaussian representation of the state's distribution after the same interval of propagation, as illustrated in *Figure 3-5*.

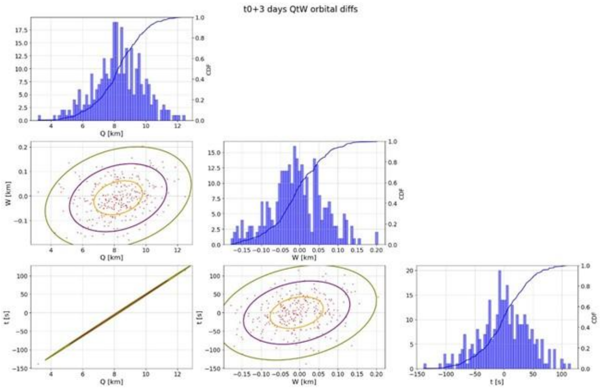

Figure 3-5. MC samples dispersion in QtW frame after 3 days of propagation. Low altitude object.

For higher altitude orbits, the benefit of resorting to the QtW representation is delayed, as the impact of drag in the dynamics is reduced. While for the 800 km the onset of non-linearities does not happen in typical propagation intervals, the QtW maintains Gaussianity after 5 days of propagation, while this is not the case for GCRF and TNW projections.

It is also observed, as anticipated in [6], that the accuracy of the projection to QtW, relies on the validity of Eq. (19) for its Jacobian, which holds as long as linearity in the QtW is maintained. This effect is illustrated in *Figure 3-6*, which represents the QtW samples and reconstructed covariance of a low altitude space object with large initial uncertainty after 3 days of propagation, together with a linearly propagated initial covariance projected by means of the Jacobian of the QtW transformation.

It is seen that while the samples maintain Gaussianity in QtW (which do not in TNW or GCRF), the covariance transformation to QtW does not perfectly fit the reconstructed Monte Carlo covariance, showcasing the limits of the underlying assumptions for the computation of the Jacobian.

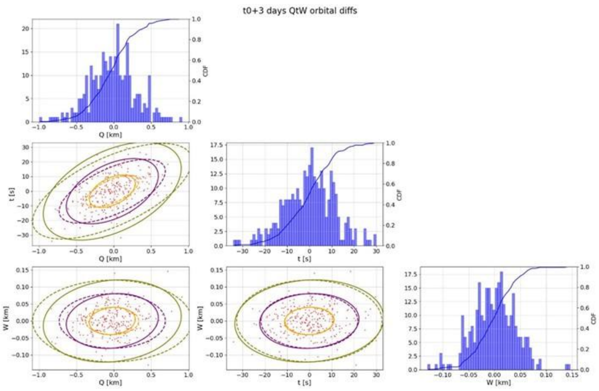

Figure 3-6. QtW samples after 3 days of propagation. The solid ellipsoids represent the Monte-Carlo reconstructed covariance. Dashed ellipsoids represent the linearly propagated initial covariance rotated with the Jacobian of the QtW transformation.

3.1.3 Huber penalty parameters

The results in this section aim to verify the convergence and robustness properties of the Huber penalty parameters pre-prototype alongside with the determination of suitable values for the Huber parameter M and of the trust region parameter λ. As initial values for these parameters, M is set to start from 5.0, considering that for a weighted non-linear least squares problem, in an ideal case with knowledge of the true trajectory, the measurement's residuals would display as a Gaussian distribution with standard deviation equal to the expected sigma of the measurements. For a normal distribution, 99.73% would fall within the 3σ boundary, which in weighted terms translates to a value of 3.0. To account for possible model errors, it is initially set to 5.0. For the trust region parameter λ, a rule of thumb is to select an initial value of $\lambda \in [0, 0.01]$ [9], with runs considering fixed and adaptive evolution of such parameter.

In reasonably good orbit determination conditions, the Huber penalty function does not show any clear improvement with respect to solutions obtained with Gauss-Newton (GN) and Levenberg-Marquardt (LM), as illustrated in *Figure 3-7.*

However, for an unfavourable orbit determination situation, in which the initial guess has been perturbed by 0.5º in true anomaly, *Figure 3-8* shows that GN does not converge, and LM is taking 20 iterations. With little to no damping (λ equal to the lowest considered values of 0 and 0.0001), the Huber penalty function does not converge either. For a fixed λ of 0.01 it does converge, while it does not for its adaptive counterpart. In both cases, the weighted RMS (WRMS) of the measurement residuals is diverging.

The best performance is obtained with a value of 0.001, in which both fixed and adaptive converge after 10 iterations.

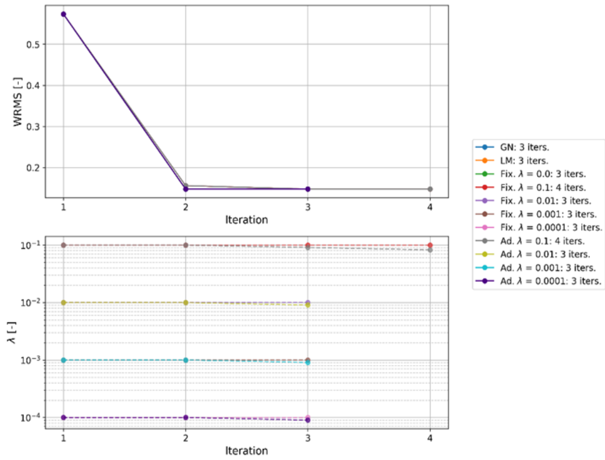

Figure 3-7. GN, LM and Huber penalty parameters comparison in terms of WRMS, together with adaptive λ evolution for a LEO nominal OD case.

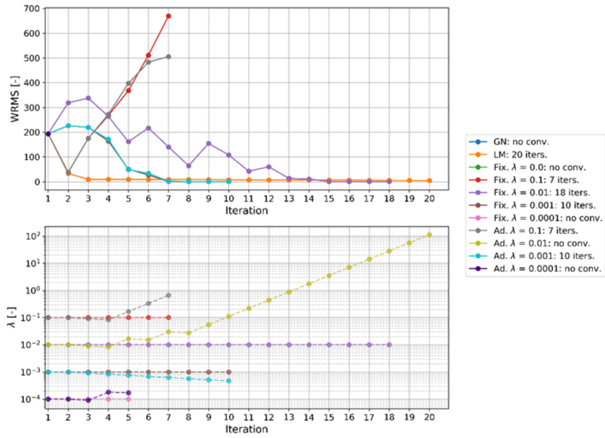

Figure 3-8. GN, LM and Huber penalty parameters comparison in terms of WRMS, together with adaptive λ evolution for a LEO unfavourable, perturbed initial guess OD case.

Additionally, if one considers high noise measurements for a radar case in LEO: 100 metres in range, 1 m/s in range-rate and 2º for angular measurements, the results in *Figure 3-9* are obtained. Both GN and LM show excellent convergence behaviour after only 3 iterations. For the Huber penalty function, a value of $\lambda = 0.1$ is not a suitable choice, with delayed convergence. For smaller values, the adaptive scheme continuously shrinks the trust region, but even by doing so, it displays a slower rate of convergence than GN and LM.

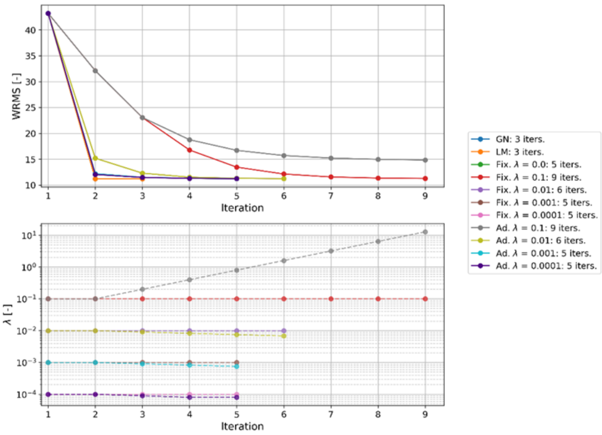

Figure 3-9. GN, LM and Huber penalty parameters comparison in terms of WRMS, together with adaptive λ evolution for a LEO high noise case.

3.1.4 Time-dependent measurements weight

The idea behind the assessment of the time-dependent measurements weight pre-prototype is to determine which correlation scale can provide the best benefits for OD in combination with realism of the output covariances.

To that end, a dynamical model, including stochastic noise perturbations is sought such that the RMS of the orbital differences between the simulated and estimated orbits are of the same order of magnitude as the measurement's accuracy. This is achieved by selecting a dynamical model for the simulated orbit that is different to the OD dynamical model, as is the case in a real scenario. Then, a parametric analysis is carried out for different scales of the time dependent weight for a total of 300 runs:

- Routine OD determination and propagation processes are run, each for different time-dependent measurement correlation time scale.
- SCP is applied to produce realistic covariances.
- Covariance realism metrics and covariance size of the realistic covariances are assessed to determine the optimum weight scale that allows to minimise the covariance size while still being realistic.

This is performed for a space object at 600 km altitude, with an applied stochastic density noise with a 40% power and correlation of 1 day, and dynamical models set up as illustrated in Table 3-1.

Table 3-1. Dynamical model differences between simulated and estimated orbits for the time-dependent measurements weight correlation time scale experiments

Dynamics	Simulated orbit	OD model
Geopotential	128x128	64x64
Third bodies	Sun&Moon	Sun&Moon
J2&Moon interaction	YES	NO
Solid Tides	YES	YES
Ocean tides	YES	NO
Drag constant area	Constant area	Constant area
SRP constant Area	Constant area	Constant area
ATMO model	NRLMSISE-00	NRLMSISE-00

Table 3-2 and Table 3-3 show the results of the stochastic density noise parameters estimation and covariance sizes for different measurement weights correlation scale. Values of CvM lower than 1.16 indicate that the tested distribution matches the χ^2 distribution with a 0.01 confidence level or larger. A similar indicator of the level of realism if the covariance containment, which must be compared against the expected containment of a multivariate distribution of the same Degrees of Freedom (DoF). In this case, position covariance was analysed, thus aiming for a theoretical containment of 97.1% for a 3σ level.

Table 3-2. Covariance realism analysis for parametric analysis of time-dependent measurements weight

τ_{meas} (days)	Analysis epochs	σ_{atm} (%)	τ_{atm} (days)	CvM metric
1	t0+2-6	24,63	0,53	0,4
2	t0+2-6	32,4	1,06	0,59
3	t0+2-6	34,3	1,86	0,89
5	t0+2-6	36,2	1,35	0,96
8	t0+2-6	37,55	1,12	1,06

Overall, the Covariance Determination methodology using Stochastic Consider Parameters can estimate density noise parameters that achieve realism, as confirmed with all CvM metrics being lower than 1.16, and containment levels similar to the expected theoretical behaviour even under such perturbed simulations.

Table 3-3. Covariance containment and size analysis for parametric analysis of time-dependent measurements weight

τ_{meas} (days)	containment 3σ (%)	Covariance trace (m)
1	94,19	1228
2	95,44	1204
3	93,93	1580
5	93,36	1760
8	93,12	1869

However, the accuracy of the estimated parameters (σ_{atm}, τ_{atm}) is affected by the choice of τ_{meas}. For very low measurement weight scales, the assigned measurement weight is below 0.1 for measurements before 2.5 days of the estimation epoch. This leads to very few measurements being considered in the batch estimation, and a degradation of the estimation accuracy. Thus, the nominal covariance is larger than the accuracy of the measurements, and the Covariance Determination methodology converges to density noise parameters that are below the introduced perturbations. On the contrary, larger weight scales for the measurements allow for a better estimation of the density noise parameters.

Nonetheless, it is seen that as the time scale increases, too many measurements are considered in the batch estimation. This leads to a lack of accuracy in the estimation, since the dynamics at 600 km are not well captured with too long measurements arc.

All these factors affect directly the size of the covariances, while still being realistic due to the Covariance Determination methodology. It is seen how the minimum covariance size is achieved for a measurement weight scale of 2 days. With this correlation scale, the applied measurement weight is less than 0.1 for measurements that are 4.75 days before the estimation, which is consistent with the fixed LUPI of around 5-6 days that is applied operationally.

3.2 Modern Sequential filters and Particle filters

Solutions for the UKF, IPLF and IPLS sequential filters are displayed to only some of the scenarios described in the preamble of Section 3. Regarding particle filters, solutions for the ENKF with (B) and without parameter estimation are shown. Other particle filters are discarded due to computational performance. ***Sstod***'s batch least-squares reference solution is also included. Additionally, some of the batch solutions have some of the proposed modern extensions applied.

Table 3-4: Nominal OD in LEO with data of several radars, 90 days, 2 objects

RMS vs reference	Covariance realism
Modern Batch least-squares	
Object 1: 152.94 m Object 2: 12.43 m Object 1 (Huber): 99.13 m Object 2 (Huber): 12.08 m	Using SCP. **Object 1:** - CvM: 0.11 - 3σ containment: 95,69 % **Object 2:** - CvM: 0.03 - 3σ containment: 97,14 %
Modern Sequential filters	
Object 1 (UKF): 74.87 m Object 2 (UKF): 58.13 m Object 1 (IPLF): 71.54 m Object 2 (IPLF): 57.92 m Object 1 (IPLS): 17.05 m Object 2 (IPLS): 8.19 m	Object 1 (UKF): - CvM: 0.304 - 3σ containment: 43.18 % Object 2 (UKF): - CvM: 0.188 - 3σ containment: 74.41 % Object 1 (IPLF): - CvM: 0.326 - 3σ containment: 20.45 % Object 2 (IPLF): - CvM: 0.195 - 3σ containment: 76.74 % Object 1 (IPLS): - CvM: 0.356 - 3σ containment: 0.0 % Object 2 (IPLS): - CvM: 0.25 - 3σ containment: 65.11 %
Particle filters	
Object 1 (ENKF): 25.31 m Object 2 (ENKF): 66.74 m Object 1 (ENKF(B)): 36.45 m Object 2 (ENKF(B)): 33.33 m	Object 1 (ENKF): - CvM: 0.05 - 3σ containment: 96 % Object 2 (ENKF): - CvM: 0.26 - 3σ containment: 100 % Object 1 (ENKF(B)): - CvM: 0.01 - 3σ containment: 100 % Object 2 (ENKF(B)): - CvM: 0.01 - 3σ containment: 100 %

For two LEO objects, at 400 (Object 1) and 800 (Object 2) km of altitude, Table 3-4 displays the RMS of orbital differences and the covariance realism analysis. The batch solution processes the complete 90 days interval in batches of 4 and 7 days, with a sliding window of 1 and 2 days. For the lowest altitude object, the batch solution displays a large RMS, subject to higher dynamics non-linearities, and with a data rejection percentage that reaches 50% for some of the ODs, and a mean value of 12%, whereas no measurements are rejected by the sequential and particle filters. This is not the case for the higher altitude object, where the classical batch solution improves UKF, IPLF and particle filters. RMS is slightly improved by the Huber pre-prototype.

For sequential filters, prominent deviations occur with the UKF after long periods without observations, which the IPLF does not exhibit due to its iterative operation. While application of the smoother by the IPLS does improve noticeably the RMS, it affects negatively covariance containment, where the UKF and IPLF already perform modestly. Particle filters and application of the SCP pre-prototype do show good containment and realism.

Both sequential and particle filters performance are known to be extremely sensitive to process noise definition. A simple parametric study in Table 3-5 confirms this. While the RMS is somehow stable, CvM first improves to then degrade, and containment suddenly drops to zero as smaller values of the process noise are explored.

Table 3-5: IPLS sensitivity to process noise

Process noise σ [m^2s^{-3}]	RMSE [m]	CvM	Containment
5e-09	8.05	0.308	1.00
5e-10	7.53	0.270	1.00
5e-11	6.57	0.151	1.00
5e-12	4.73	0.349	0.05
5e-13	3.48	0.357	0.00

For an object in GEO, for measurement data for 90 days, Table 3-6 shows its performance in terms of RMS vs reference orbit and measurement residuals. The batch solution processes the dataset in intervals of 20 days with a sliding window of 2 days. No measurements are rejected at all. Residuals are aligned with the expected noise from the sensor. UKF and IPLF performances are aligned with the batch solution, while the IPLS reduces RMS by one order of magnitude, despite maintaining similar residual performance.

Particle filters display a slightly poorer performance when parameters are not being estimated, while when parameters are being estimated in EnKF(B), the algorithm fails to attain parameter stability.

Table 3-6: Nominal OD in GEO with data of a telescope network, 90 days, 1 object

RMS vs reference	Measurements residuals (mean, RMS)
Modern Batch least-squares	
Classical: 95.93 m Huber: 95.93 m	Classical: - Right ascension: 0.0, 0.155 mdeg - Declination: 0.0, 0.140 mdeg Huber: - Right ascension: 0.0, 0.155 mdeg - Declination: 0.0, 0.140 mdeg
Modern Sequential filters	
UKF: 162.3 m IPLF: 162.3 m IPLS: 23.41 m	UKF: - Right Ascension: -0.0 mdeg, 0.14 mdeg - Declination: 0.0 mm/s, 0.14 mm/s IPLF: - Right Ascension: -0.0 mdeg, 0.14 mdeg - Declination: 0.0 mm/s, 0.14 mm/s IPLS: - Right Ascension: 0.0 mdeg, 0.14 mdeg - Declination: -0.0 mm/s, 0.14 mm/s
Particle filters	
ENKF 292.01 m ENKF(B): 885.57 m	ENKF: - Right ascension: 0.065 mdeg, 0.156 mdeg - Declination: -0.010 mdeg, 0.132 mdeg ENKF(B): - Right ascension: 0.322 mdeg, 0.563 mdeg - Declination: -0.094 mdeg, 0.177 mdeg

Table 3-7 shows the performance when an OD for a LEO object is initialised by a solution coming from an Initial Orbit Determination (IOD) method. Not all solutions are displayed for residuals as they display a very similar behaviour for each type of algorithm. Batch methods display improved performance over sequential and particle filters. Furthermore, a radius of convergence analysis is performed for this case, in which the initial guess is incrementally perturbed until the algorithm fails to converge. Batch methods are able to deal with solutions that fall up to 11 km, extended to 18 km with the Huber penalty pre-prototype. However, accuracy of the estimate quickly degrades beyond 8 and 5 km of perturbation respectively (i.e. quicker for the Huber methodology). Sequential filters can no longer converge after a 4 km perturbation, with RMS also exhibiting a degrading performance. Particle filters also fail to show a stable behaviour, which would dissipate the initial error, following a trend similar to the particle filters. The typical frequency of data in this domain may prove to be too low.

Table 3-7: OD in LEO with data, with the initial state from an IOD

RMS vs reference	Measurements residuals (mean, RMS)
Modern Batch least-squares	
Classical: 12.81 m Huber: 13.59 m	Classical: - Azimuth: -38.502, 1097.078 mdeg - Elevation: 11.144, 162.751 mdeg - Range: -0.564, 12.234 m - Range-rate: -363.568, 4662.013 mm/s
Modern Sequential filters	
UKF: 42.73 m IPLF: 42.73 m IPLS: 35.12 m	IPLS: - Azimuth: -64.89, 1839.66 mdeg - Elevation: -0.3, 221.91 mdeg - One-way Range: -0.28, 11.97 m - One-way Range-rate: -164.78, 5229.08 mm/s
Particle filters	
ENKF 91.14m ENKF(B): 120.84 m	ENKF(B): - Azimuth: -544.45 mdeg, 2073.18 mdeg - Elevation: -83.39 mdeg, 264.22 mdeg - Range: 12.83 m, 29.39 m - Range-rate: 1686.39 mm/s, 9322.90 mm/s

The ability of the algorithm to deal with miscorrelated data is illustrated in Table 3-8. The dataset includes a 20% of observations of a close-by object, with a difference of 0.5º in inclination. On average, the batch solutions reject 8% of the measurements, whereas sequential filters reject 12.24% throughout the complete dataset processing. Particle filters fail to converge under these simulation conditions.

Table 3-8: OD in LEO with radar data, including miscorrelated observations, 15 days

RMS vs reference	Data Rejection %
Modern Batch least-squares	
137.64 m	8.21 %
Modern Sequential filters	
UKF: 31.31 m IPLF: 26.54 m IPLS: 14.24 m	UKF: 12.24 % IPLF/s: 12.24 %

3.3 Robust Data Association

Results for three different scenarios will be presented: real measurements during 1 week for 20 objects in LEO from 2 radars, simulation of 20 objects from 4 survey telescopes in GEO for 1 month, and simulated measurements for 2 near-by objects (500 metres along-track separation) over the course of a week.

For reference, the same cases are analysed by means of the track-to-track algorithm implemented in ***Sstod.***

Table 3-9 below reports some statistics regarding the number of tracks which have been validated (i.e., promoted by the algorithm and associated to an object) and the non-validated ones, along with the number of detected and missed objects. Only 37% of the tracks were correctly associated, while 80% of the objects detected. These are relatively low values due to the challenge that real data usually poses. Additionally, the radars providing the data have a high expected noise sigma.

Table 3-9: Track and object statistics for data association in LEO with real data for 20 objects, track-to-track algorithm

Total tracks	194	**Total objects**	20
Validated tracks	88	**Detected objects**	16
Non-validated tracks	106	**Missed objects**	4

The algorithm makes use of a Figure of Merit (FoM) as the driver for association scoring and promotion. This value depends on the quality of a orbit determination process (goodness of fit) performed with the tracks in a given association hypothesis, The distribution of the Figure of Merit (FoM) of associations with 4 or more tracks distributed in terms of the estimated perigee is shown in *Figure 3-10*. True positives for objects between 7000 and 7200 km of perigee have a FoM 0.1, whereas false positives range around a value of 1.0. This indicates that a better sensor characterisation in terms of expected noise sigma could lower the number of false positives since the FoM would increase.

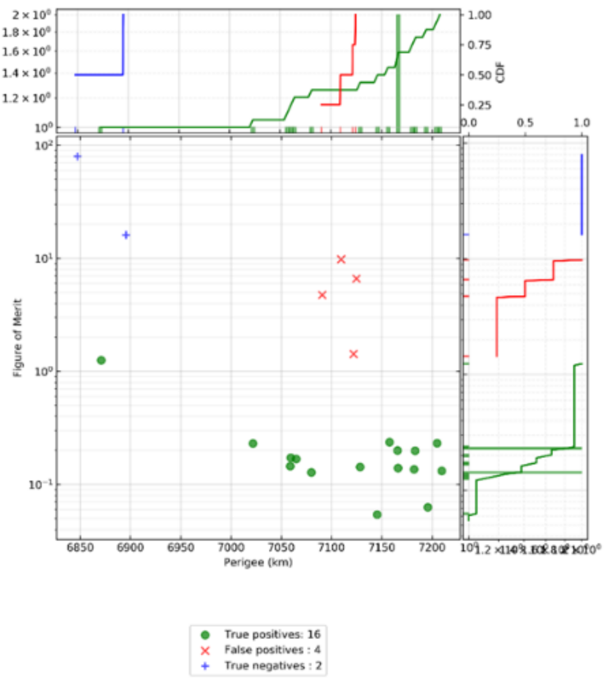

Figure 3-10. Distribution of the FoM according to perigee altitude. LEO, 20 objects, real data

In terms of multi-target trackers, the datasets have been approached by the GNN and JPDA trackers. These algorithms require a prior for the objects that are to be tracked, which is not the usual case in the space environment. Simple track initiator methods were tested, to no avail. Future phases of the project will complement with suitable IOD algorithms. Thus, initiated tracks eventually break down.

For the first of the GEO cases, with simulated data for 20 objects, the track-to-track algorithm correctly associates 92.4% of the tracks and detects all 20 objects, with improved performance in a better characterised environment and sensors with smaller expected noise sigma. GNN and JPDA are in this case supplied with priors for the tracked objects. GNN generates a total of 42 trajectories on the 20 targets, all of which contain measurements from exclusively one target, and no targets are missed. Out of the 20 objects, only 2 targets are continuously followed by a single estimated trajectory, whereas the other objects generate between 2-3 trajectories, due to the inability to associate new detections to existing trajectories. On the contrary, while due to the working principles of the JPDA algorithm, values for the suggested metrics cannot be produced, it generates near-perfect association results: 22 trajectories for 20 targets, without missing any target, and generating a single trajectory per target except for two, which were however properly terminated eventually. JPDA's superior performance to GNN is illustrated in *Figure 3-11*, evaluating the GOSPA penalisation metric [22]. Inspection reveals that aside from localisation errors, there are significant "missed" and "false" target contributions.

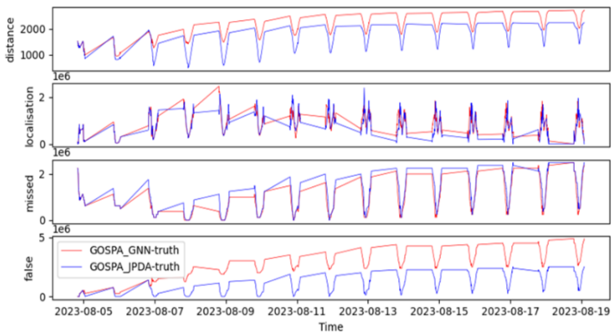

Figure 3-11. GOSPA metric for DA in GEO with multi-target trackers for clearly distinct objects

Finally, a case of two objects in close conjunction in GEO, separated by 500 metres in along-track. In this case, the track-to-track algorithm fails. Out of 18 tracks, 16 are validated, however these are only part of the false positive associations, and no object is detected. Data scarcity (only 1 track per object per sensor per night). On the other hand, both GNN and JPDA generate exactly two trajectories, with nearly identical performance.

4 CONCLUSIONS

Batch methods have proven again to be adequate to deal with state estimation in the space environment. Modern extensions have been investigated and deemed promising. The Huber penalty function is not costly compared to the classic approach and is seen to increase robustness in the case of measurement outliers. Still, it is recommended to implement an outlier rejection mechanism alongside it. The time-weighted measurement function benefits the estimates (state and covariance) by introducing more measurements into the analysis interval even if with lowered weight, while it is challenging to determine the best time scale in a general manner. The SCP pre-prototype is seen to be able to estimate the system's uncertainty and correct accordingly the estimated and predicted covariance. Finally, the QtW transformation proves to be a successful approach to maintain Gaussianity for extended time intervals of propagation, which offers a promising performance coupled to the SCP.

In terms of modern sequential filters, they have proved solid performance in the benchmarking scenarios. While the IPLF offers little improvement over the UKF, the IPLS does provide significant enhancements, but at the cost of re-processing the data. In terms of state and covariance realism, the process noise is seen to play a key role, as is well known, and is typically not easy to assess.

The particle filters investigated in the activity have been reduced to EnKF with and without parameter estimation, where others such as the standard particle filter are discarded due to performance limitations. As for sequential filters, these have been found sensitive to the intensity of the process noise. The EnKF with no dynamical parameters estimation has been seen to perform better, showcasing a more stable behaviour. In general, the particle filters are seen to suffer from model and data mismatch.

In terms of data association, the track-to-track algorithm in GMV's ***Sstod*** software has proved effective. It has demonstrated good precision, sensitivity and accuracy, except for a close conjunction case in GEO. The GNN and JPDA sequential multi-target trackers have also shown a promising performance, while they should be accompanied by suitable trajectory initiator methods, as they rely on a prior. Additionally, proper trajectory termination techniques need also to be incorporated. JPDA has proven superior to GNN and both have dealt successfully with the close conjunction GEO case.

The IPLF and IPLS filters for state estimation and JPDA algorithm for multi-target tracking are selected for implementation in a prototype library, RODDAS.

REFERENCES

1. Schiemenz F, Utzmann J, Kayal H (2019). Least squares orbit estimation including atmospheric density uncertainty consideration. *Advances in Space Research* 63(12):3916–3935.

2. J. Emmert, H. Warren, A. Segerman, J. Byers, and J. Picone. Propagation of atmospheric density errors to satellite orbits. *Advances in Space Research*, vol. 59, pp. 147–165, jan 2017

3. L. Sagnieres and I. Sharf. Uncertainty characterization of atmospheric density models for orbit prediction of space debris. *Proceedings of the 7th European Conference on Space Debris*, (Darmstadt, Germany), ESA Space Debris Office, 2017

4. Cano, A., Pastor, A., Escobar, D., Míguez, J., Sanjurjo-Rivo, M. Covariance determination for improving uncertainty realism in orbit determination and propagation. *Advances in Space Research* (2022).

5. Cano, A., Pastor, A., Fernández, S., Míguez, J., Sanjurjo-Rivo, M., Escobar, D. Improving orbital uncertainty realism through covariance determination in GEO. *The Journal of the Astronautical Sciences* 69(5) (2022).

6. Laurens, S., Dolado, J., Cavallaro, G., Jouisse, M., Seimandi, P., 2018. Towards the maintenance of Gaussianity on state vector uncertainty propagation, in: *69th International Astronautical Congress*. id: IAC-18,A6,9,3,x42351.

7. B. D. Tapley, B. E. Schutz, G. H. Born, Statistical Orbit Determination, *Elsevier Academic Press*, 2004.

8. B. Wilson, "Robust orbit determination via convex optimization," *Project Report, EE364B*, Stanford University, 2014.

9. D. W. Marquardt, An algorithm for least-squares estimation of nonlinear parameters. *J. Soc. Ind. Appl. Math.* 1963.

10. Hiscocks, S., Barr, J., Perree, N., Wright, J., Pritchett, H., Rosoman, O., Harris, M., Gorman, R., Pike, S., Carniglia, P., Vladimirov, L., & Oakes, B. (2023). Stone Soup: No Longer Just an Appetiser *26th International Conference on Information Fusion.*

11. Simo Särkkä and Lennart Svensson (2023). Bayesian Filtering and Smoothing, Second Edition. *Cambridge University Press*

12. B. Bell and F. Cathey, "The iterated Kalman filter update as a Gauss-Newton method," *IEEE Trans. Autom*. Control, vol. 38, no. 2, pp. 294–297, Feb. 1993.

13. Morelande, M. R. & García-Fernández, A. F., Analysis of Kalman filter approximations for nonlinear measurements, *IEEE Transactions on Signal Processing*, 2013, 61, 5477-5484

14. Kevin P. Murphy, "Probabilistic Machine Learning: Advanced Topics", *MIT Press*, 2023.

15. A F. Garcia-Fernandez, L. Svensson, and S. Sarkka, "Iterated posterior linearisation smoother," *IEEE Trans. Automat. Control*, vol. 62, no. 4, pp. 2056–2063, Apr. 2017.

16. A. Doucet, S. Godsill, and C. Andrieu. On sequential Monte Carlo Sampling methods for Bayesian filtering. *Statistics and Computing*, 10(3):197–208, 2000.

17. Koblents, Eugenia, and Joaquín Míguez. "A population Monte Carlo scheme with transformed weights and its application to stochastic kinetic models." *Statistics and Computing* 25.2, 2015

18. Pérez-Vieites, Sara, Inés P. Mariño, and Joaquín Míguez. "Probabilistic scheme for joint parameter estimation and state prediction in complex dynamical systems." *Physical Review E 98.6* (2018): 063305.

19. D. Crisan and J. Miguez. Uniform convergence over time of a nested particle filtering scheme for recursive parameter estimation in state–space Markov models. *Advances in Applied Probability*, 49(4):1170–1200, 2017.

20. D. Crisan and J. Miguez. Nested particle filters for online parameter estimation in discrete-time state-space Markov models. Bernoulli, 2018

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

22. A. S. Rahmathullah, Á. F. García-Fernández and L. Svensson, "Generalized optimal sub-pattern assignment metric," 2017 *20th International Conference on Information Fusion*.