AAS 19-845

ROBUST PARTICLE FILTER FOR SPACE OBJECTS TRACKING UNDER SEVERE UNCERTAINTY

Cristian Greco[*], Lorenzo Gentile[†], Massimiliano Vasile[‡], Edmondo Minisci[§], and Thomas Bartz-Beielstein[¶]

This paper presents a robust particle filter approach able to handle a set-valued specification of the probability measures modelling the uncertainty structure of tracking problems. This method returns robust bounds on a quantity of interest compatibly with the infinite number of uncertain distributions specified. The importance particles are drawn and propagated only once, and the bound computation is realised by inexpensively tuning the importance weights. Furthermore, the uncertainty propagation is realised efficiently by employing an intrusive polynomial algebra technique. The developed method is finally applied to the computation of a debris-satellite collision probability in a scenario characterised by severe uncertainty.

INTRODUCTION AND MOTIVATIONS

Standard Bayesian state estimation techniques for space objects tracking and collision avoidance require the definition of a single precise probability measure, either explicitly or implicitly. Indeed, precise probability distributions are employed to model uncertainty both in the prior knowledge of state and in the received observations. As a consequence, the computed posterior distribution may be greatly sensitive to the specification of such measures.

However, in many real-world applications, it may be difficult, or likely impossible, to characterise the true precise distribution, since it would require a perfect knowledge of the factors which cause uncertainty in the first place, e.g. uncertain dynamical parameters and sensors noise. In particular, this issue is of major importance for the always growing number of low-budget satellites or space debris, objects which can rely only on a limited number of often inaccurate observations. Furthermore, the majority of methods relies on parametric distributions and simplifying assumptions to keep the analysis computationally tractable. As an example, the renowned Kalman Filter is based on the assumption (or approximation for nonlinear variants) of Gaussian distributions describing the state prior, the dynamical process noise, and the observation likelihood. Different methods exist

[*]PhD Candidate, Department of Mechanical and Aerospace Engineering, University of Strathclyde, 75 Montrose Street, G1 1XJ, Glasgow, United Kingdom.

[†]PhD Candidate, Institute for Data Science, Engineering and Analytics, TH Köln, Steinmüllerallee 1, 51643, Gummersbach, Germany.

[‡]Professor, Department of Mechanical and Aerospace Engineering, University of Strathclyde, 75 Montrose Street, G1 1XJ, Glasgow, United Kingdom.

[§]Senior Lecturer, Department of Mechanical and Aerospace Engineering, University of Strathclyde, 75 Montrose Street, G1 1XJ, Glasgow, United Kingdom.

[¶]Professor, Institute for Data Science, Engineering and Analytics, TH Köln, Steinmüllerallee 1, 51643, Gummersbach, Germany.

to deal with more complex distributions, such as particle filters or Gaussian mixture ones, but they still rely on precise and very informative distributions to model the involved uncertainties. Finally, the fact that these estimation techniques have proven to be optimal or efficient primarily depends on the assumptions made in their problem statement, rather than their accuracy in describing the true uncertainty structure of certain dynamical systems.

To overcome these limitations, this work presents a novel method for state estimation of dynamical systems under an imprecise specification of the probability distributions modelling the problem uncertainty.[1] This approach allows one to compute robust bounds on a quantity of interest, and therefore an interval of equally likely values, while a traditional technique would return a single point estimate sensitive to the prior and likelihood choice. Furthermore, since it will be shown that the filter evaluation stage is decoupled from the bound computation routine, the new filtering scheme can be used also for efficient sensitivity analysis, i.e. to study how the filter output is affected by changes in input parameters of the involved distributions.

The developed method is tested on the scenario of tracking a non-operational debris in a potential collision trajectory with an operational satellite. The quantity of interest to compute robust bounds of is the probability of collision between the two space objects. The results of this approach are particularly relevant for space surveillance, where a specific probability threshold is used to flag up potential collisions.

The remainder of the paper is structured as follows. The next section deals with the filtering problem formulation and solution approach. First, the case of standard precise uncertainty specification is discussed and the general Particle Filter algorithm described. Successively, a generalisation of the filtering formulation to accommodate further epistemic uncertainty is developed, and then the newly developed Robust Particle Filter is presented. The subsequent section discusses the application of such method to the computation of collision probability in the presence of severe uncertainty, and it shows the results of the developed approach in comparison to a standard filter. Finally, the paper is concluded with the final remarks in the last section.

FILTERING FORMULATION AND ROBUST SOLUTION APPROACH

The state-space model for the state estimation problem addressed in this research is a continuous-discrete one.[2] The system state evolves according to a time-continuous ordinary differential equation, whereas indirect observation $\mathbf{y}_k$ are collected at discrete instances of time. Specifically, the state-space model is formulated as

$$\begin{cases} \dot{\mathbf{x}}(t) = \mathbf{f}(t, \mathbf{x}(t), \mathbf{d}) & \text{(1a)} \\ \mathbf{y}_k \quad = \mathbf{h}(t_k, \mathbf{x}(t_k), \boldsymbol{\varepsilon}) \qquad \text{for } k = 1, \ldots, l\,. & \text{(1b)} \end{cases}$$

In Equation (1a), $\mathbf{x}$ is the system continuous state, $\mathbf{f}$ represents the functional relationship of the equations of motion, t is the independent variable (usually time), and $\mathbf{d}$ are static model parameters. In Equation (1b), $\mathbf{y}_k$ are observations collected at discrete time t_k, $\mathbf{h}$ is the generally nonlinear relation between the state and the observation value, and $\boldsymbol{\varepsilon}$ is a generic noise affecting the measurement realisation.

If the system initial condition $\mathbf{x}(t_0) = \mathbf{x}_0$ and the dynamical model parameters $\mathbf{d}$ were known perfectly, there would be no need of measurements, as the equations of motion could be (usually numerically) integrated to obtain the system evolution in time. However, in real-life scenarios uncertainty is always involved in such systems, and measurements are needed to refine the state

knowledge at a later time. In this paper, we will consider the case of uncertainty, modelled as a random variable, affecting:

- the initial condition $\mathbf{x}_0 \sim p(\mathbf{x}_0)$;
- the static parameters $\mathbf{d} \sim p(\mathbf{d})$, which are nuisance ones and should not be estimated;
- the measurement realisation $\mathbf{y}_k \sim p(\mathbf{y}_k|\mathbf{x}_k)$, as a result of the noise ε, described by a conditional distribution.

Hence, the standard continuous-discrete state-space model in Equation (1) can be reformulated in a probabilistic fashion[3] to explicitly describe the uncertain nature of system as

$$\begin{cases} \mathbf{x}_0 \sim p(\mathbf{x}_0) & \text{(2a)} \\ \mathbf{x}_t \sim p(\mathbf{x}_t \,|\, \mathbf{x}_\tau) \qquad \text{with } \tau \leq t & \text{(2b)} \\ \mathbf{y}_k \sim p(\mathbf{y}_k \,|\, \mathbf{x}_k) \qquad \text{for} \quad k = 1, \dots, l \,, & \text{(2c)} \end{cases}$$

where the transition probability $p(\mathbf{x}_t \,|\, \mathbf{x}_\tau)$ describes the system dynamics as a continuous-time Markov chain, as resulting from the uncertainty on $\mathbf{d}$.

Precise Filtering

Given a precise specification of the densities in Equation (2), the complete solution of the filtering problem of general state estimation is the posterior distribution of the state conditional to the previously received observations, which in a sequential fashion is stated by Bayesian inference as

$$p(\mathbf{x}_k \,|\, \mathbf{y}_{1:k}) = \frac{p(\mathbf{y}_k \,|\, \mathbf{x}_k)\, p(\mathbf{x}_k \,|\, \mathbf{y}_{1:k-1})}{p(\mathbf{y}_k \,|\, \mathbf{y}_{1:k-1})} \,, \tag{3}$$

with $t_1 < t_2 < \cdots < t_k$. Generally, this posterior distribution is very expensive to compute, and its complete knowledge provides greater information than needed for most practical applications. Therefore, in Bayesian inference, the filtering problem can be reduced to the computation of the expectation of a generic function $\phi(\mathbf{x}_k)$

$$\mathbb{E}_p\big[\phi(\mathbf{x}_k) \,|\, \mathbf{y}_{1:k}\big] = \int \phi(\mathbf{x}_k)\, p(\mathbf{x}_k \,|\, \mathbf{y}_{1:k}) d\mathbf{x}_k \,, \tag{4}$$

with respect to the posterior distribution.

Once a method is available to compute such expectation, every quantity of interest can be estimated by tuning the generic function ϕ. As relevant examples, the expected value of the state itself conditional to the observations should be computed using $\phi(\mathbf{x}_k) = \mathbf{x}_k$, whereas the probability of a generic event C should be computed using the indicator function of that event $\phi(\mathbf{x}_k) = \mathbb{I}_C(\mathbf{x}_k)$. Similarly, the variance, confidence regions, and other statistics can be computed by plugging the appropriate function.

When no specific assumption or parameterisation is imposed on the distributions $p(\cdot)$ in Equation (2), Equation (4) has no closed-form solution and numerical techniques are required. Among them, the particle filter is a sequential Monte Carlo method that approximates the posterior density as a discrete one by using samples $\mathbf{x}_k^{(i)}$ drawn from a proposal importance distribution as

$$\pi(\mathbf{x}_k \,|\, \mathbf{x}_{k-1}, \mathbf{y}_{1:k}) \,, \tag{5}$$

here assumed as Markovian for numerical and theoretical convenience. Equation (4) can then be approximated as

$$\mathbb{E}_p\big[\phi(\mathbf{x}_k) \,|\, \mathbf{y}_{1:k}\big] \approx \sum_{i=1}^{N} \hat{w}_k^{(i)} \phi(\mathbf{x}_k^{(i)}) \,, \tag{6}$$

where $\hat{w}_k^{(i)} = w_k^{(i)} / \sum_{j=1}^{N} w_k^{(j)}$ are normalised weights which are sequentially computed as

$$w_k^{(i)} = w_{k-1}^{(i)} \frac{p(\mathbf{y}_k \,|\, \mathbf{x}_k^{(i)})\; p(\mathbf{x}_k^{(i)} \,|\, \mathbf{x}_{k-1}^{(i)})}{\pi(\mathbf{x}_k^{(i)} \,|\, \mathbf{x}_{k-1}^{(i)}, \mathbf{y}_{1:k})} \,. \tag{7}$$

Hence, once the samples $\mathbf{x}_{k-1}^{(i)}$ have been propagated $\mathbf{x}_k^{(i)}$, which is by far the most expensive step in aerospace applications, the filtering inference is solved by the inexpensive evaluations of the function ϕ and the densities in Equations (2) and (5).

One common choice for the importance distribution is to use the transition probability of the system as in Equation (2b):

$$\pi(\mathbf{x}_k \,|\, \mathbf{x}_{k-1}, \mathbf{y}_{1:k}) = p(\mathbf{x}_k | \, \mathbf{x}_{k-1}) \,. \tag{8}$$

This choice is advantageous as generating samples from the transition probability is generally a straightforward task. Furthermore, by substituting Equation (8) in Equation (7), the weight update expression simplifies to

$$w_k^{(i)} = w_{k-1}^{(i)} \; p(\mathbf{y}_k \,|\, \mathbf{x}_k^{(i)}) \,. \tag{9}$$

On the other hand, this importance distribution does not take into account the last received observation $\mathbf{y}_k$ in generating the new samples. This often leads to the *degeneracy problem* when most of the particles have zero (or nearly zero) weights as a consequence of being outside (or in an extremely unlikely subspace) of the observation likelihood support. To face this issue, *resampling* is a step that could be employed to replace degenerate particles with more likely ones. This is achieved by sampling new particles from the posterior distribution, approximated after the update step as

$$p(\mathbf{x}_k \,|\, \mathbf{y}_{1:k}) \approx \sum_{i=1}^{N} \hat{w}_k^{(i)} \; \delta(\mathbf{x}_k - \mathbf{x}_k^{(i)}) \,, \tag{10}$$

where δ is the Dirac's delta function.[3] Assorted algorithms exist to perform the resampling,[4] but since this step always increase the estimate variance, low variance methods should be preferred.

From here, the general *Particle Filter* algorithm is reported in Algorithm 1 as it will serve as the precise basis over which the advancements presented in the next section will be developed. From the computed weights and samples, the expectation of a generic function can be computed via Equation (6).

When the transition probability is used as importance distribution (as in Equation (8)), the resulting Particle Filter is known as *Bootstrap Filter*.[5]

Imprecise Filtering

In the setting of severe uncertainty, the probability distributions $p(\cdot)$ are not assumed to be known precisely, but they are only specified within parameterised sets of probability measures. This for-

Algorithm 1 Algorithmic scheme for the Particle Filter.

Given:

· the state model in Equation (2)

· the proposal distribution $\pi(\mathbf{x}_k \,|\, \mathbf{x}_{k-1}, \mathbf{y}_{1:k})$

1: Draw N particles from the initial distribution $p(\mathbf{x}_0)$ and set the weights to be equal
$\mathbf{x}_0^{(i)} \sim p(\mathbf{x}_0)\,, \;\; \hat{w}_0^{(i)} = 1/N$

2: **for** $k = 1 : l$ **do**

3: Draw one particle $\mathbf{x}_k^{(i)}$ for each $\mathbf{x}_{k-1}^{(i)}$ from the proposal distribution
$\mathbf{x}_k^{(i)} \sim \pi(\mathbf{x}_k \,|\, \mathbf{x}_{k-1}^{(i)}, \mathbf{y}_{1:k})$

4: Update the weights with Equation (7) and normalise them to unity
$w_k^{(i)} = \hat{w}_{k-1}^{(i)}\, p(\mathbf{y}_k \,|\, \mathbf{x}_k^{(i)})\, p(\mathbf{x}_k^{(i)} \,|\, \mathbf{x}_{k-1}^{(i)}) / \pi(\mathbf{x}_k^{(i)} \,|\, \mathbf{x}_{k-1}^{(i)}, \mathbf{y}_{1:k})\,, \;\; \hat{w}_k^{(i)} = w_k^{(i)} / \sum_{j=1}^{N} w_k^{(j)}$

5: Perform resampling and set the samples to be equal
$\mathbf{x}_k^{(i)} \sim \sum_{i=1}^{N} \hat{w}_k^{(i)}\, \delta(\mathbf{x}_k - \mathbf{x}_k^{(i)})\,, \;\; \hat{w}_k^{(i)} = 1/N$

6: **end for**

mulation will hold for all the uncertainties involved:

$$\begin{aligned} p(\mathbf{x}_0) &\in \mathcal{P}_{\mathbf{x}_0} \\ p(\mathbf{d}) &\in \mathcal{P}_{\mathbf{d}} \\ p(\mathbf{y}_k \,|\, \mathbf{x}_k) &\in \mathcal{P}_{\mathbf{y}_k|\mathbf{x}_k}\,. \end{aligned} \tag{11}$$

Within these imprecise sets of distributions, no judgement is made about their likeliness of being the true one. This definition allows one to model more faithfully uncertain scenarios in which information is too scarce to specify a single distribution.

In the setting of state estimation, the probabilistic continuous-discrete filtering problem in Equation (2) is adapted to account for the new epistemic uncertainty component

$$\begin{cases} \mathbf{x}_0 \sim \quad p(\mathbf{x}_0) \quad \in \mathcal{P}_{\mathbf{x}_0} & & \text{(12a)} \\ \mathbf{x}_t \sim p(\mathbf{x}_t \,|\, \mathbf{x}_\tau) \;\in \mathcal{P}_{\mathbf{x}_t|\mathbf{x}_\tau} & \text{with } \tau \leq t & \text{(12b)} \\ \mathbf{y}_k \sim p(\mathbf{y}_k \,|\, \mathbf{x}_k) \in \mathcal{P}_{\mathbf{y}_k|\mathbf{x}_k} & \text{for} \quad k = 1, \ldots, l\,, & \text{(12c)} \end{cases}$$

where the transition probability and its corresponding imprecise set in Equation (12b) result from the uncertainty on the static parameters $\mathbf{d}$ which influence the system dynamical evolution.

This uncertainty specification induces an infinite number of posterior distributions $p(\mathbf{x}_k \,|\, \mathbf{y}_{1:k})$ to plausibly exist compatibly with each precise distribution in the imprecise sets. Hence, in this development, a robust estimation routine is formulated for computing lower and upper bounds on the expectation of the quantity of interest conditional on the observations, as

$$\underline{\mathbb{E}}\big[\phi(\mathbf{x}_k) \,|\, \mathbf{y}_{1:k}\big] = \inf_{p \in \mathcal{P}_k} \mathbb{E}_p\big[\phi(\mathbf{x}_k) \,|\, \mathbf{y}_{1:k}\big] \tag{13a}$$

$$\overline{\mathbb{E}}\big[\phi(\mathbf{x}_k) \,|\, \mathbf{y}_{1:k}\big] = \sup_{p \in \mathcal{P}_k} \mathbb{E}_p\big[\phi(\mathbf{x}_k) \,|\, \mathbf{y}_{1:k}\big]\,. \tag{13b}$$

These lower and upper expectations express the tight bounds on the expectation of the quantity of interest as resulting from the imprecise specification of uncertainty. Clearly, if the imprecise sets

were composed of a single distribution, Equations (13a) and (13b) would coincide and collapse in Equation (4).

These bounds can be computed by a direct optimisation over parameterised imprecise sets by running a precise filtering method, e.g. the Bootstrap filter, for each candidate precise distribution to evaluate. In this case, however, the optimisation would be highly inefficient as the expectation $\mathbb{E}_p$ would be used as a black-box cost function, and therefore the computations performed for one filter run would be discarded once $\mathbb{E}_p$ is computed.

Hence, we introduce a filtering approach which is based on the precomputation of the most expensive operations involved in the Bootstrap filter. By taking further advantage of the *importance sampling* technique, it is possible to efficiently compute the expectation of interest for new precise distributions by correction weights in a fraction of the original time.

The proposed *Robust Particle Filter* (RPF) approach is described in Algorithm 2. First, a *Precomputation* step is carried out to speed up the successive optimisation process. Three extra proposal distributions $\pi(\mathbf{x}_0)$, $\pi(\mathbf{x}_k|\mathbf{x}_{k-1})$ and $\pi(\mathbf{y}_k|\mathbf{x}_k)$ are selected such that their support encloses the ones of the distributions in respectively $\mathcal{P}_0$, $\mathcal{P}_{\mathbf{x}_k|\mathbf{x}_{k-1}}$ and $\mathcal{P}_{\mathbf{y}_k|\mathbf{x}_k}$. Hence, the precise Particle Filter in Algorithm 1 is run with these distributions as prior, transition and likelihood, with the standard proposal distribution $\pi(\mathbf{x}_k|\mathbf{x}_{k-1}, \mathbf{y}_{1:k})$. From this computation, the sets of initial samples $\{\mathbf{x}_0^{(1)}, \ldots, \mathbf{x}_0^{(N)}\}$ and propagated ones $\{\mathbf{x}_k^{(1)}, \ldots, \mathbf{x}_k^{(N)}\}$ are stored to avoid further propagation in the next filter calls. A generic optimisation statement is reported in the algorithm description for the *Bound Computation* stage. The optimiser to practically solve this step should be selected once a specific state estimation problem is formulated, hence exploiting at maximum the prior knowledge about the objective index landscape, which highly depends on the dynamics, observation models and uncertainty specifications. Whichever optimisation method is selected, the *Filter Evaluation* of the expectation of interest, i.e. the very solution of the filtering problem, is made efficient by exploiting the performed precomputations. Contrarily to the precise Particle Filter, the initial particles are not sampled anymore from the prior distribution directly, but rather the precomputed samples $\{\mathbf{x}_0^{(1)}, \ldots, \mathbf{x}_0^{(N)}\}$ are employed, and their weights adjusted as the ratio between the evaluated prior $p(\mathbf{x}_0)$ and the proposal initial distribution $\pi(\mathbf{x}_0)$ so that the new distribution is targeted. Then, a for-loop over the observation times is carried out to update the importance weights of the *precomputed* propagated samples given the newly received measurements with Equation (7). Finally, the expectation $\mathbb{E}_p\big[\phi(\mathbf{x}_l) \,|\, \mathbf{y}_{1:l}\big]$ is evaluated via Equation (6) with the *precomputed* final samples $\mathbf{x}_l^{(i)}$ and the updated weights $\hat{w}_l^{(i)}$.

The main advantage of this filtering scheme is that the prediction step from the proposal distribution (see step 3 of Algorithm 1) is removed from the filter evaluation thanks to the precomputations with the proposal distributions. This results in a substantial saving of computational time as the numerical propagation of particles through the dynamical system in Equation (1a) is supplanted by an inexpensive weight re-scaling in the update equation. This improvement is critical to computing efficiently the lower and upper expectations by optimisation, as the objective index $\mathbb{E}_p$ can be inexpensively evaluated for different candidate distributions, only requiring a fraction of the time than the one a precise Particle Filter would require.

The resampling step has been removed from the filter evaluation as well when comparing it to the precise Particle Filter. There is no intrinsic or theoretical reason why the new robust filter should not have a resampling step, but this development choice is motivated by a twofold reasoning. First, a resampling step is already carried out in the precise filter evaluation within the precomputation

Algorithm 2 Algorithmic scheme for the Robust Particle Filter.

Given:

· the state model in Equation (12)
· the proposal distribution $\pi(\mathbf{x}_k \,|\, \mathbf{x}_{k-1}, \mathbf{y}_{1:k})$
· the initial proposal distribution $\pi(\mathbf{x}_0)$
· the transition proposal distribution $\pi(\mathbf{x}_k \,|\, \mathbf{x}_{k-1})$
· the likelihood proposal distribution $\pi(\mathbf{y}_k \,|\, \mathbf{x}_k)$

Precomputation

1: Run the precise Particle Filter (Algorithm 1) with :

· $\pi(\mathbf{x}_0)$ as prior
· $\pi(\mathbf{x}_k \,|\, \mathbf{x}_{k-1})$ as transition
· $\pi(\mathbf{y}_k \,|\, \mathbf{x}_k)$ as likelihood
· $\pi(\mathbf{x}_k \,|\, \mathbf{x}_{k-1}, \mathbf{y}_{1:k})$ as proposal

and save:

· the set of initial samples $\{\mathbf{x}_0^{(1)}, \ldots, \mathbf{x}_0^{(N)}\}$
· the sets of propagated samples $\{\mathbf{x}_k^{(1)}, \ldots, \mathbf{x}_k^{(N)}\}$ at time t_k, for $k = 1, \ldots, l$

Bound Computation

2: Compute the bounds by direct optimisation over the imprecise sets

$$\underline{\mathbb{E}}\big[\phi(\mathbf{x}_l) \,|\, \mathbf{y}_{1:l}\big] = \inf_{p \in \mathcal{P}} \mathbb{E}_p\big[\phi(\mathbf{x}_l) \,|\, \mathbf{y}_{1:l}\big]$$
$$\overline{\mathbb{E}}\big[\phi(\mathbf{x}_l) \,|\, \mathbf{y}_{1:l}\big] = \sup_{p \in \mathcal{P}} \mathbb{E}_p\big[\phi(\mathbf{x}_l) \,|\, \mathbf{y}_{1:l}\big]$$

where $p \in \mathcal{P}$ indicates any triple $\{p(\mathbf{x}_0), p(\mathbf{x}_k|\mathbf{x}_{k-1}), p(\mathbf{y}_k|\mathbf{x}_k)\}$ in the imprecise set

Filter Evaluation

$\{p(\mathbf{x}_0), p(\mathbf{x}_k|\mathbf{x}_{k-1}), p(\mathbf{y}_k|\mathbf{x}_k)\} \rightarrow \mathbb{E}_p\big[\phi(\mathbf{x}_l) \,|\, \mathbf{y}_{1:l}\big]$

3: Set the weights of the *precomputed* initial samples $\mathbf{x}_0^{(i)}$ and scale them to unity
$w_0^{(i)} = p(\mathbf{x}_0^{(i)})/\pi(\mathbf{x}_0^{(i)})\,, \;\; \hat{w}_0^{(i)} = w_0^{(i)} / \sum_{j=1}^N w_0^{(j)}$

4: **for** $k = 1 : l$ **do**

5: Update the weights of the *precomputed* samples $\mathbf{x}_k^{(i)}$ and $\mathbf{x}_{k-1}^{(i)}$ and normalise them to unity
$w_k^{(i)} = \hat{w}_{k-1}^{(i)}\, p(\mathbf{y}_k \,|\, \mathbf{x}_k^{(i)})\; p(\mathbf{x}_k^{(i)} \,|\, \mathbf{x}_{k-1}^{(i)})/\pi(\mathbf{x}_k^{(i)} \,|\, \mathbf{x}_{k-1}^{(i)}, \mathbf{y}_{1:k})\,, \;\; \hat{w}_k^{(i)} = w_k^{(i)} / \sum_{j=1}^N w_k^{(j)}$

6: **end for**

7: Evaluate the expectation $\mathbb{E}_p$ with the *precomputed* samples $\mathbf{x}_l^{(i)}$ and the updated weights $\hat{w}_l^{(i)}$
$\mathbb{E}_p\big[\phi(\mathbf{x}_l) \,|\, \mathbf{y}_{1:l}\big] = \sum_{i=1}^N \hat{w}_l^{(i)} \phi(\mathbf{x}_l^{(i)})$

stage. Hence, unless the imprecise set $\mathcal{P}_{\mathbf{y}_k|\mathbf{x}_k}$ defines a collection of extremely different likelihoods, already the majority of the samples are concentrated in the distributions' support. Second, a new resampling in the *Filter Evaluation* would imply the generation of new particles (or duplicate ones

with different dynamical noise) which would require new dynamical propagations, and therefore a computational performance decrement.

TEST CASE

The developed approach is applied to the tracking of a space debris in a potential collision orbit with an operational well-tracked satellite. The quantity of interest to compute robust bounds of is the collision probability. A collision is assumed to happen when the distance between the two objects is smaller than a given threshold. Hence, the probability of collision is evaluated by computing the expectation of the indicator function of the satellite-debris distance to be less than or equal to the chosen threshold.

SOCRATES (Satellite Orbital Conjunction Reports Assessing Threatening Encounters in Space),[6] an online service which provides twice-daily reports on the most likely collision events based on NORAD two-line elements (TLEs) and the SGP4 propagator, is employed to select two space objects with low range at the close approach, whose elements are summarised in Table 1. The first object is an operational satellite, NORAD ID 35684, whose ephemerides are assumed to be known with good accuracy. The second object, NORAD ID 17296, is a non-operational rocket body in orbit since 1987, and uncertainty is considered in its ephemerides. The reference time of closest approach (TCA) reported is the date 15-Jul-2019 at 21:54:40 UTC.

Table 1: Spacecraft orbital elements at reference epoch from NORAD TLEs.

NORAD ID	EPOCH [UTC]	a [km]	e [-]	i [deg]	Ω [deg]	ω [deg]	θ [deg]
35684	09-Jul-2019 22:50:30	7009.95	4.7e-3	98.23	155.74	194.40	165.72
17296	10-Jul-2019 04:07:40	6992.96	4.0e-3	82.57	40.59	211.66	203.54

The motion of the body is described in Cartesian components in an Earth-centered inertial reference frame. The dynamical model employed (Equation (1a)) includes the following components:[7] the gravitational force using the EGM96 geopotential model up to degree and order 4; atmospheric drag according to Jacchia-Gill model; third-body disturbances due to the Moon and Sun gravitational attraction; solar radiation pressure (SRP) with a conical shadow model.

To face the requirement of a possibly high number of particles in a multi-dimensional nonlinear scenario as tracking of space objects in Low-Earth Orbit (LEO), a generalised intrusive polynomial expansion[8] over extended sets is employed for uncertainty propagation. This method replaces the numerical integration of the equations of motion for different initial conditions with inexpensive polynomial evaluations. The intrusive approach is implemented in the SMART-UQ[9] (Strathclyde Mechanical and Aerospace Research Toolbox for Uncertainty Quantification) C++ library, which enables the propagation of polynomials through any dynamics thanks to overload operators of the standard double operations. This approach has been applied for uncertainty propagation in a number of aerospace applications,[10,11] specifically employing the Chebyshev basis which has shown to exhibit superior global convergence and robustness.

The polynomial expansion is used to efficiently propagate samples to the observation times and the reference TCA. However, small deviations in the initial conditions at the epoch time lead to substantial differences in the final state at the reference TCA due to the dynamical nonlinearities in LEO, and the rather long propagation time. Hence, the actual TCA for an uncertainty realisation generally differs from the reference TCA. To compute the actual TCA, the GSL GNU implemen-

tation of the golden-section algorithm[12] is employed to find the precise time which minimises the distance of closest approach (DCA) between the two satellites. The reference TCA is provided as initial guess and a numerical integrator is used to propagate the state at the required evaluation time from the polynomially propagated state. This combination of polynomial propagation and subsequent optimisation routine is still considerably faster than a direct numerical propagation with event stoppage because the actual TCA is generally close to the reference one, and therefore the numerical propagation in the optimisation requires very few integration steps.

Initial Ephemeris Uncertainty

After retrieving the Cartesian state from the TLEs in Table 1, covariance errors are considered in the non-operational satellite initial state. The value of such uncertainty, taken from ESA guidelines[13] for inclinations larger than 60 deg, is summarised in Table 2 in radial, transversal and normal components.

Table 2: $\mathbf{1}\sigma$ position (r) and velocity (v) uncertainty of TLEs for orbits with $\mathbf{e < 0.1}$, $\mathbf{i > 60}$ deg, perigee altitude $\leq \mathbf{800km}$, in radial (U), transversal (V), and normal components (W).

$1\sigma_{r_U}$ [m]	$1\sigma_{r_V}$ [m]	$1\sigma_{r_W}$ [m]	$1\sigma_{v_U}$ [mm/s]	$1\sigma_{v_V}$ [mm/s]	$1\sigma_{v_W}$ [mm/s]
104	556	139	559	110	148

The covariance matrix resulting from these standard deviations is rotated to $\sigma^2_{\mathbf{x}_0}$ in the covariance in the Cartesian inertial frame. From here, the importance initial distribution is defined as a normal distribution

$$\pi(\mathbf{x}_0) = \mathcal{N}\left(\mu_{\mathbf{x}_0}, \sigma^2_{\mathbf{x}_0}\right) \tag{14}$$

where $\mu_{\mathbf{x}_0}$ is the Cartesian state retrieved from the object TLE.

An initial Monte Carlo uncertain propagation is carried out without any observation by propagating 10^6 samples drawn from the distribution in Equation (14). The resulting DCA distribution is reconstructed from the propagated samples as an Empirical Cumulative Distribution Function (ECDF)

$$F_{DCA}(dca \leq \delta_{dca}) = \frac{1}{N}\sum_{i=1}^{N} \mathbb{I}_{DCA(\mathbf{x}^{(i)}_{TCA}) \leq \delta_{dca}} \,, \tag{15}$$

where

$$DCA(\mathbf{x}^{(i)}_{TCA}) = \|\mathbf{r}^{(i)}_{s2}(TCA) - \mathbf{r}_{s1}(TCA)\|_2 \tag{16}$$

is the function which returns the Euclidean distance between the position vectors of the propagated sample of the debris and the operational satellite at the actual TCA. The resulting ECDF is plotted in Figure 1. For a threshold distance $\delta_{DCA} = 500$ m (a conservative threshold for most satellites and debris), the collision probability compatible with $\pi(\mathbf{x}_0)$ in Equation (14) is

$$F_{DCA}(dca \leq 500 \text{ m}) = 0.4105\% \,.$$

In the following, epistemic uncertainty is considered to affect the initial distributions, with the

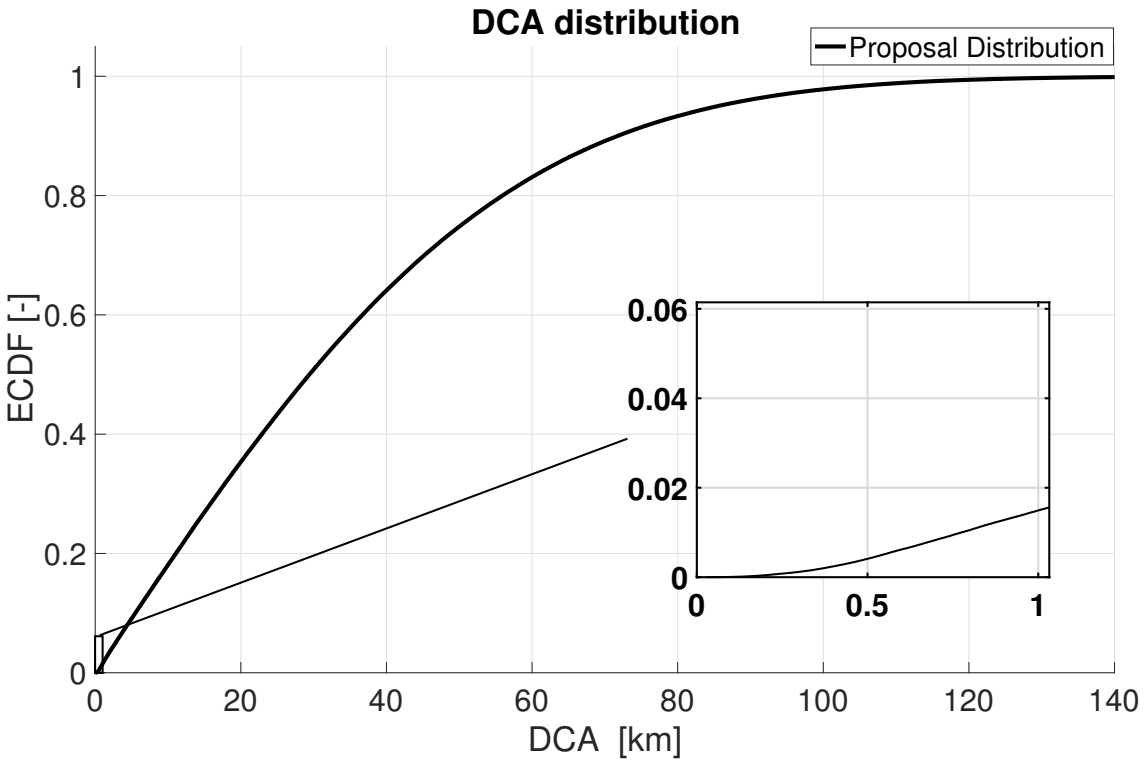

Figure 1: Empirical Cumulative Distribution Function for Distance of Closest Approach between debris and operational satellite as resulting from initial uncertainty specification in Equation (14).

imprecise set $\mathcal{P}_{\mathbf{x}_0}$ parameterised as

$$\begin{aligned}\mathcal{P}_{\mathbf{x}_0} = \{\, p(\mathbf{x}_0) :\ & p(\mathbf{x}_0) = \mathcal{N}\Big(\mu_{\mathbf{x}_0}, \widetilde{\sigma}^2_{\mathbf{x}_0}\Big), \\ & \widetilde{\sigma}^2_{\mathbf{x}_0} = \text{diag}(\, \lambda_{\mathbf{x}_{0-1}}\, \sigma^2_{\mathbf{x}_0}(1{:}3, 1{:}3),\, \lambda_{\mathbf{x}_{0-2}}\, \sigma^2_{\mathbf{x}_0}(4{:}6, 4{:}6)), \\ & \lambda_{\mathbf{x}_{0-1}} \in [0.33^2, 1.5^2],\, \lambda_{\mathbf{x}_{0-2}} \in [0.33^2, 1.5^2] \,\}\,,\end{aligned} \tag{17}$$

where $\sigma^2_{\mathbf{x}_0}(1{:}3, 1{:}3)$ and $\sigma^2_{\mathbf{x}_0}(4{:}6, 4{:}6)$ indicate respectively the position block and the velocity block of the covariance matrix $\sigma^2_{\mathbf{x}_0}$, and the operator diag indicates a block-diagonal matrix. Therefore, the set $\mathcal{P}_{\mathbf{x}_0}$ is parameterised using two multipliers $\lambda_{\mathbf{x}_{0-1}}$ and $\lambda_{\mathbf{x}_{0-2}}$ which scale the covariance matrix, reducing the initial uncertainty for multipliers < 1, or increasing the initial uncertainty for multipliers > 1. The multiplier range $[0.33^2, 1.5^2]$ means that the standard deviations of the initial state may be reduced to 33% or increased by 50% of their reference value in Table 2.

We can now run Algorithm 2 with no observations using $\pi(\mathbf{x}_0)$ as proposal prior and compute the bounds of $F_{DCA}(dca \leq \delta_{dca})$ for the probabilities compatible with the imprecise set in Equation (17). To compute the probability of the DCA being less than or equal to δ_{dca}, the auxiliary function ϕ to employ is

$$\phi(\mathbf{x}_{TCA}) = \mathbb{I}_{DCA(\mathbf{x}_{TCA}) \leq \delta_{dca}}(\mathbf{x}_{TCA})\,, \tag{18}$$

since the expectation of the indicator function of an event is equal to the probability of that event, i.e.

$$F_{DCA}(dca \leq \delta_{dca}) = \mathbb{E}_{\mathbf{x}_{TCA}}\left[\mathbb{I}_{DCA(\mathbf{x}_{TCA}) \leq \delta_{dca}}(\mathbf{x}_{TCA})\right]\,. \tag{19}$$

The *Bound Computation* stage of Algorithm 2 is realised by first a Design of Experiments (DOE) plan employing 20 Halton samples for $\lambda_{\mathbf{x}_{0-1}}$ and $\lambda_{\mathbf{x}_{0-2}}$ in their definition set, and then by a local interior-point optimiser (MATLAB's *fmincon*) starting once from the lowest and once from the highest probability value found among the samples of the initial DOE.

The ECDFs resulting from the bounds obtained by the developed method are plotted in Figure 2, together with the ones resulting from Halton samples used for the DOE, from extra Halton samples computed after the optimisation to check that the bounds actually enclose all the set, and the initial proposal distribution for comparison with Figure 1.

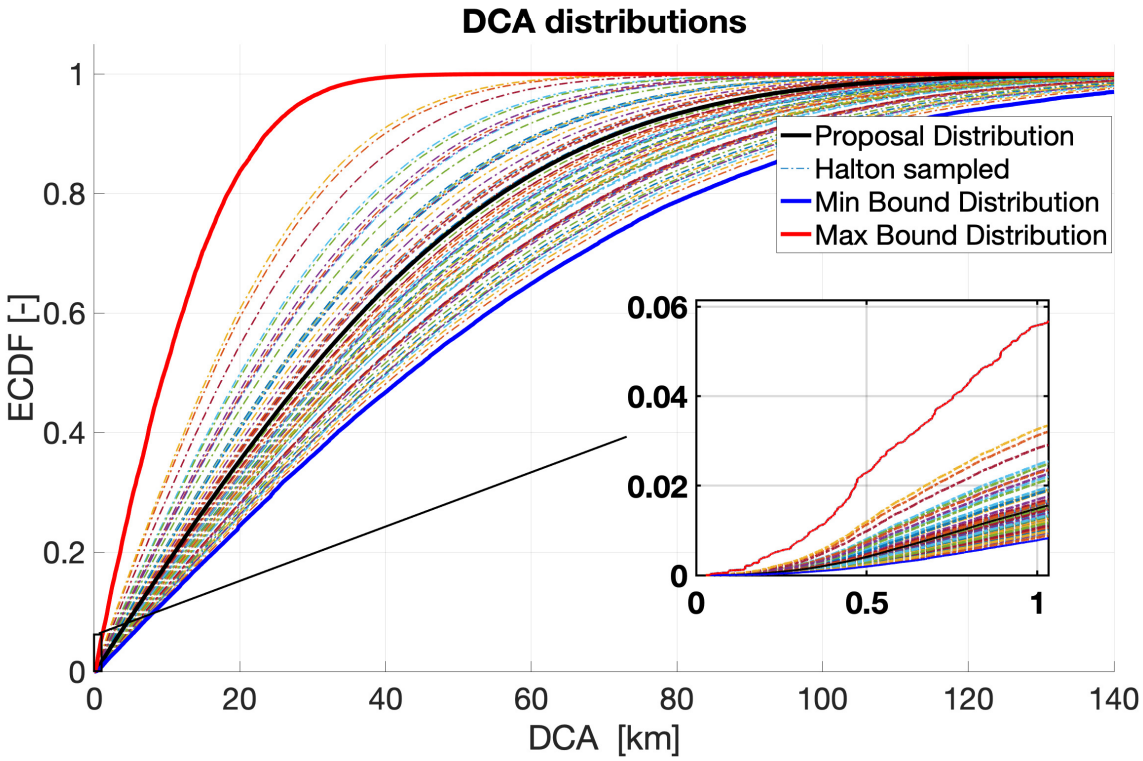

Figure 2: Empirical Cumulative Distribution Functions for Distance of Closest Approach between debris and operational satellite as resulting from initial uncertainty specification in Equation (17). The ECDFs correspond to the bounds found by the developed method, different Halton samples, and the proposal prior distribution.

It is straightforward to see that the bounds computed with the developed approach do enclose the ECDFs corresponding to the imprecise set $\mathcal{P}_{\mathbf{x}_0}$. It is interesting to see how the cumulative functions differ significantly, and therefore the probability value for a specific threshold, even for relatively small modifications in the initial distribution $p(\mathbf{x}_0)$ standard deviations. Indeed, the impact probability for $\delta_{DCA} = 500$ m, which is now an interval value, is

$$F_{DCA}(dca \leq 500 \text{ m}) = [\,0.1942\%\,,\, 2.2945\%\,]\,,$$

pointing out the substantial sensitivity of the quantity of interest to the parameters of the prior distribution. The multipliers corresponding to these lower and upper bounds are reported in Table 3. In this case without observations, the parameters generating the lower distribution coincide with the

Table 3: Multipliers which correspond to lower and upper collision probability for case with no observations.

Bound	$\lambda_{\mathbf{x}_{0-1}}$	$\lambda_{\mathbf{x}_{0-2}}$
Lower	1.5^2	1.5^2
Upper	0.33^2	0.33^2

upper bounds of the multiplier definition in Equation (17), whereas the upper distribution comes from the lower bounds of the multipliers' range. These outcomes are in line with what expected since the mean of the prior distribution $\mu_{\mathbf{x}_0}$ corresponds to a collision trajectory. Therefore, the smaller the covariance the more probability is concentrated in uncertain regions which will eventually lead to a collision occurrence, whereas the larger the covariance the less likely those regions are. The linearity of the objective index with respect to the free variables justifies the employment of a local gradient-based solver.

In the context of space situational awareness, some collision probabilities in the range of 0.1942% to 2.2945% (or similar ones for different scenarios) may and some may not trigger an alarm, just for different standard deviations used to define the prior. This result further highlights the usefulness of methods considering also epistemic uncertainty versus purely precise approaches.

Filtering with Observations

Now we will consider the filtering case with received observations which will reshape the propagated uncertainty at intermediate times. Two observation windows are assumed between the satellite epoch and the date of closest approach. The dates of these measurements are 14-Jul-2019 at 08:00:00 UTC, and 15-Jul-2019 at 04:00:00 UTC. The indirect quantities measured (Equation (1b)) are the debris azimuth and elevation with respect to the equatorial plane. The ideal measurements are generated using the debris reference trajectory, i.e. with initial conditions $\mu_{\mathbf{x}_0}$, which is one of the collision trajectories of the space debris. To mimic measurement errors, the noisy measurements are drawn from the distribution

$$\pi(\mathbf{y}_k \,|\, \mathbf{x}_k) = \mathcal{N}\Big(\mathbf{h}(t_k, \mathbf{x}_k, \mathbf{0}), \sigma^2_{\mathbf{y}_k}\Big)\,, \tag{20}$$

where $\mathbf{h}(t_k, \mathbf{x}_k, \mathbf{0})$ is the ideal azimuth an elevation observation model with zero noise, and $\sigma^2_{\mathbf{y}_k}$ is the diagonal covariance resulting from the standard deviations specified in Table 4. The resulting noisy observations are named as $\bar{\mathbf{y}}_1$ and $\bar{\mathbf{y}}_2$.

Table 4: 1σ azimuth (az) and elevation (el) uncertainty for noisy measurements of debris.

$1\sigma_{az}$ [deg]	$1\sigma_{el}$ [deg]
0.1	0.1

For reference and precomputation, the precise Bootstrap Filter, i.e. Algorithm 1 with Equation (8) and with multinomial resampling, is run with the prior specified in Equation (14) and likelihood in Equation (20). The distribution resulting from the Bootstrap Filter with 10^6 samples is plotted in Figure 3.

By comparing Figure 3 and Figure 1 we can notice how the inclusion of observations leads to an

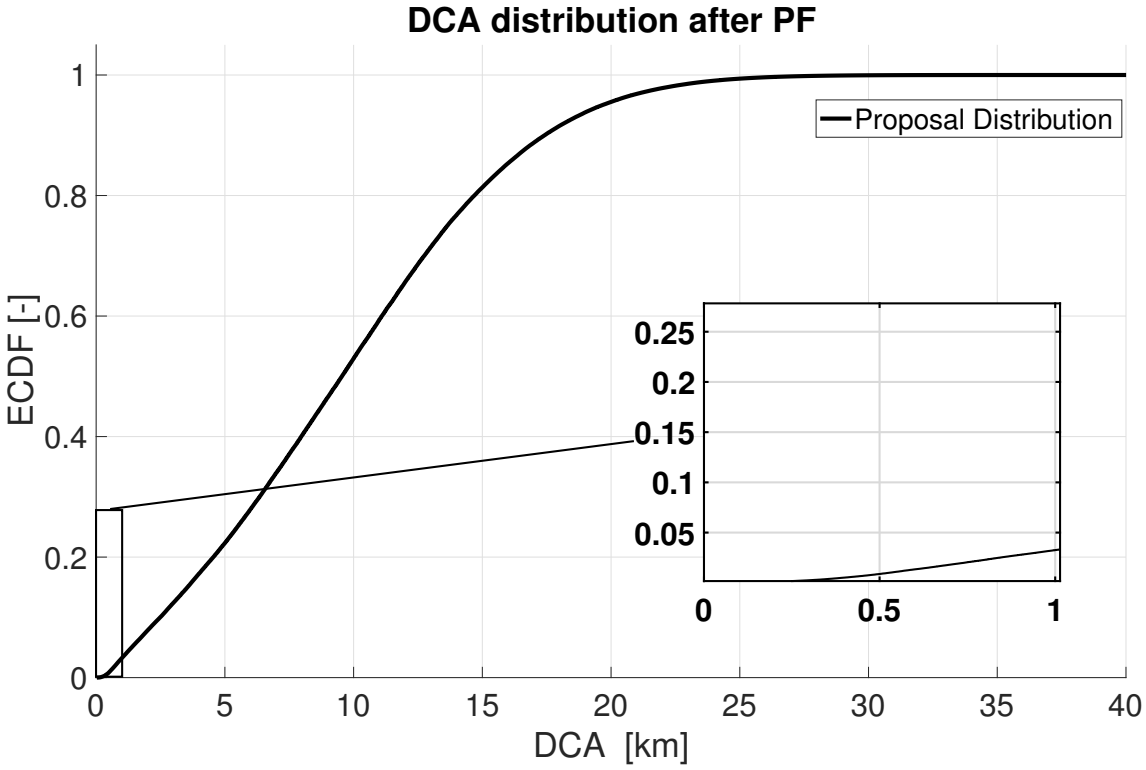

Figure 3: Empirical Cumulative Distribution Function for Distance of Closest Approach between debris and operational satellite as resulting from Bootstrap Filter with two observations, with prior as in Equation (14) and likelihood as in Equation (20).

increase in the probability of collision for any threshold, with the value for the 500 m equal to

$$F_{DCA}(dca \leq 500 \text{ m}) = 0.877\% \, .$$

Indeed, since the observations are taken relatively close to the collision trajectory, the state uncertainty is shifted closer to regions which more likely will result in the collision event.

Epistemic uncertainty will now be considered in both the prior distribution, using the same set definition as in Equation (17), and the observation likelihood with the imprecise set $\mathcal{P}_{\mathbf{y}_k|\mathbf{x}_k}$ parameterised as

$$\begin{aligned} \mathcal{P}_{\hat{\mathbf{y}}_k|\mathbf{x}_k} = \{\, p(\hat{\mathbf{y}}_k|\mathbf{x}_k) :\ & p(\hat{\mathbf{y}}_k|\mathbf{x}_k) = \mathcal{N}\Big(\hat{\mathbf{y}}_k, \widetilde{\sigma}^2_{\mathbf{y}_k}\Big), \\ & \widetilde{\sigma}^2_{\mathbf{y}_k} = \text{diag}(\, \lambda_{\mathbf{y}_{k-1}}\, \sigma^2_{\mathbf{y}_k}(1,1),\, \lambda_{\mathbf{y}_{k-2}}\, \sigma^2_{\mathbf{y}_k}(2,2)), \\ & \lambda_{\mathbf{y}_{k-1}} \in [0.33^2, 1.5^2],\, \lambda_{\mathbf{y}_{k-2}} \in [0.33^2, 1.5^2] \,\} \, , \end{aligned} \tag{21}$$

where $\sigma^2_{\mathbf{y}_k}(1,1)$ and $\sigma^2_{\mathbf{y}_k}(2,2)$ indicate respectively the azimuth and elevation variance values of the reference covariance matrix $\sigma^2_{\mathbf{y}_k}$ as resulting from the standard deviations in Table 4. Similarly to the imprecise prior definition, the imprecise likelihood set $\mathcal{P}_{\hat{\mathbf{y}}_k|\mathbf{x}_k}$ is parameterised using two multipliers per observation.

Therefore, in the full case of severe uncertainty, 6 parameters are affected by epistemic uncertainty: two multipliers, $\lambda_{\mathbf{x}_{0-1}}$ and $\lambda_{\mathbf{x}_{0-2}}$, for the prior distribution; four multipliers, $\lambda_{\mathbf{y}_{k-1}}$ and

$\lambda_{\mathbf{y}_{k-2}}$ for $k = 1, 2$, for the observation likelihoods at the two observation times. These 6 multipliers are the free variables for the *Bound Computation* routine.

Given these parameterisations for $\mathcal{P}_{\mathbf{x}_0}$ and $\mathcal{P}_{\hat{\mathbf{y}}_k|\mathbf{x}_k}$, the newly developed Robust Particle Filter is run with proposal prior as in Equation (14) and proposal likelihood as in Equation (20). The transition probability will be used as importance distribution like it is done in the Bootstrap filter.

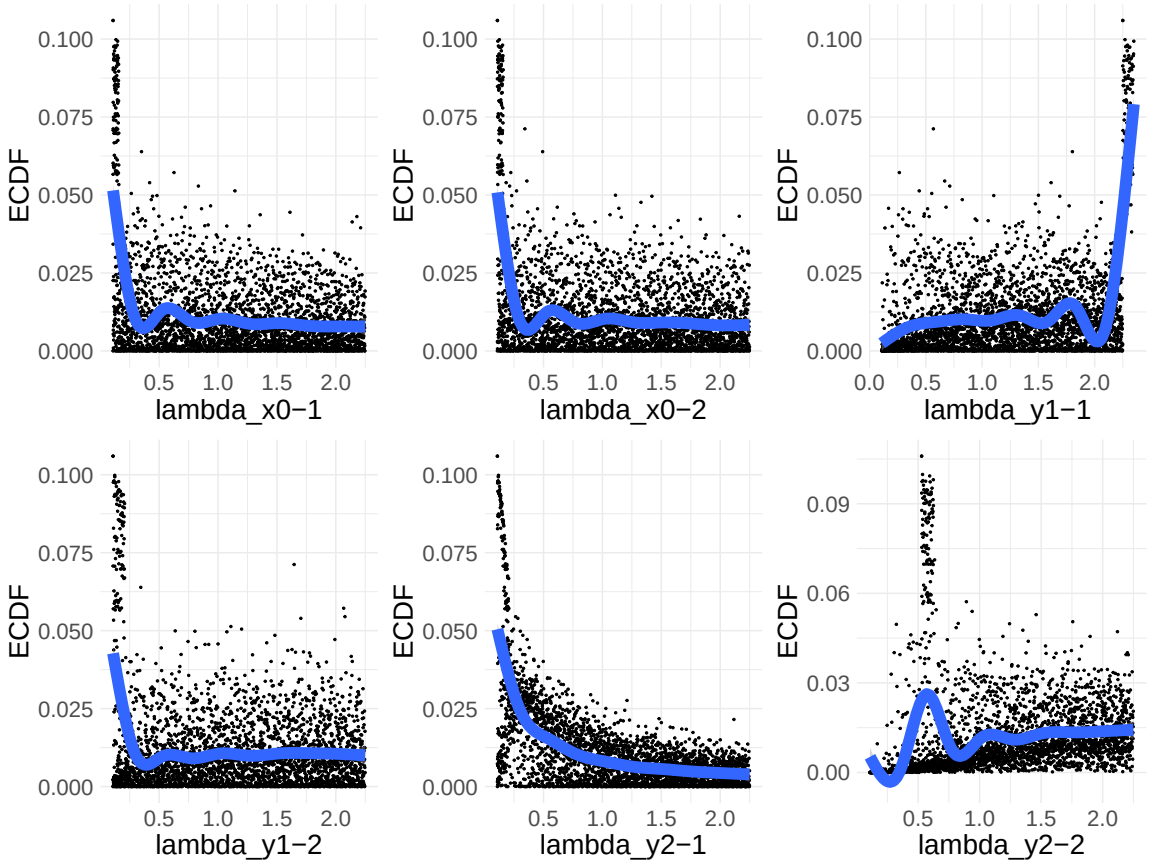

Figure 4: Correlations between the free variables and the ECDF. The blue solid line represents the locally estimated scatterplot smoothing (LOESS). It is evident that the ECFD does not behaves linearly in respect of any of the considered free variables.

In this case, the *Bound Computation* stage relies on a sophisticated global optimisation routine. In fact, the identification of the bounds is hardened by the higher number of parameters to be considered and their unknown interactions. For these reasons, the Adaptive multi-population inflationary differential evolution algorithm (MP-AIDEA)[14,15] has been employed. MP-AIDEA consists in a multi-population hybridisation of Differential Evolution[16] with the restarting procedure of monotonic basin hopping;[17] it implements both a local restart in the neighbourhood of a local minimum and a global restart in the whole search space. In this framework, the aforementioned optimisation strategy has been chosen because of its robustness and efficiency extensively shown in literature.[14]

In this case, it can be seen from Figure 4 that there is no strong linear dependency between the free variables and the collision probability. Furthermore, one can see from Figure 5 the prominent role assumed by the interactions between the free variables in determining the collision probability.

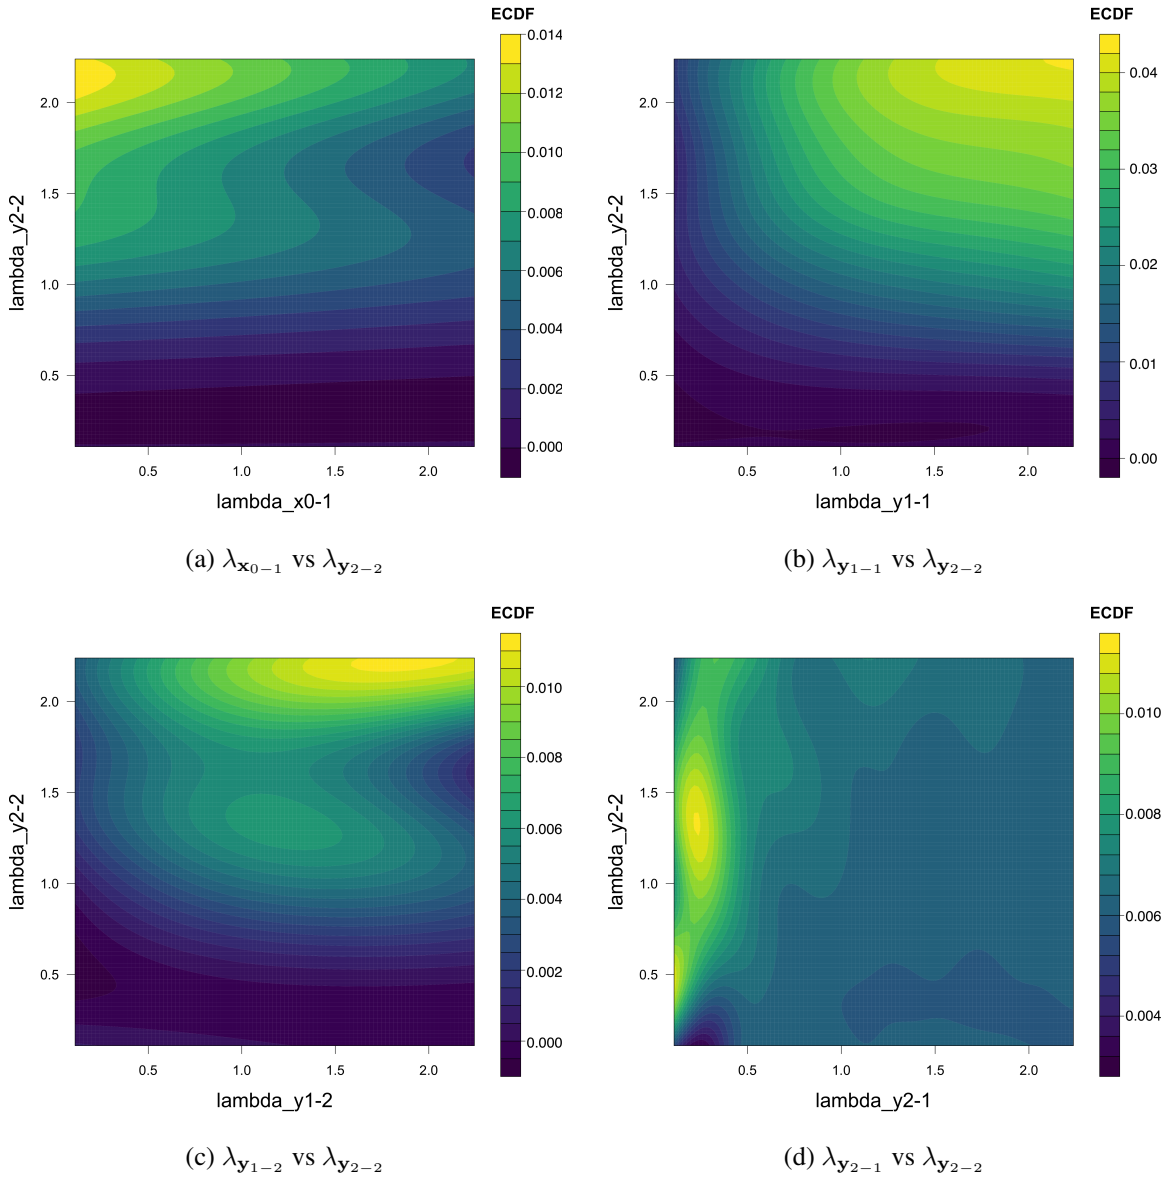

(a) $\lambda_{\mathbf{x}_{0-1}}$ vs $\lambda_{\mathbf{y}_{2-2}}$

(b) $\lambda_{\mathbf{y}_{1-1}}$ vs $\lambda_{\mathbf{y}_{2-2}}$

(c) $\lambda_{\mathbf{y}_{1-2}}$ vs $\lambda_{\mathbf{y}_{2-2}}$

(d) $\lambda_{\mathbf{y}_{2-1}}$ vs $\lambda_{\mathbf{y}_{2-2}}$

Figure 5: Visualizations of the collision probability landscape with respect to respect of most impacting variable combinations. For each individual plots, variables that are not shown are fixed to the respective values of the lower bound solution.

These two features make the objective function strongly multi-modal and, consequently, not trivial to optimise. In light of these considerations, the need for a global search strategy that can effectively handle interactions and global search appears evident.

The results of the Robust Particle Filter run with these settings are plotted in Figure 6, again together with a number of random solutions generated trough Halton sequence for an a posteriori validation of the bounds computed.

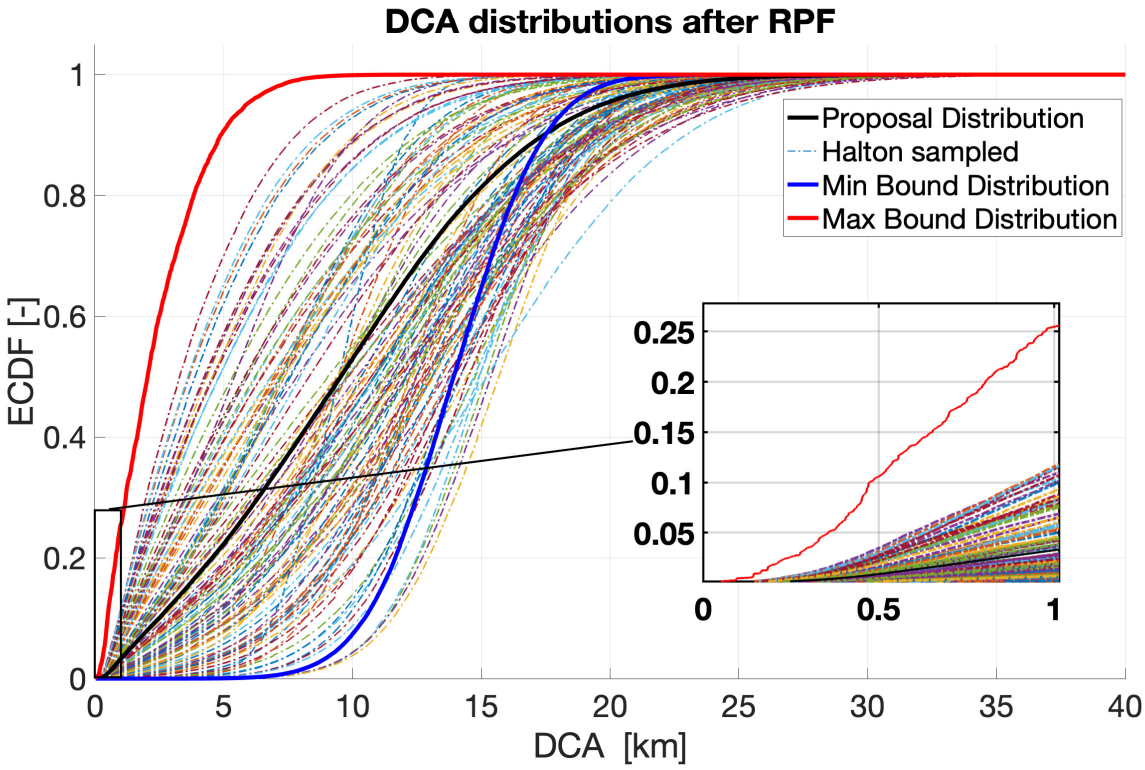

Figure 6: Empirical Cumulative Distribution Function for Distance of Closest Approach between debris and operational satellite as resulting from the Robust Particle Filter with two observations, with prior imprecise set as in Equation (17) and likelihood imprecise set as in Equation (21).

The impact probability for $\delta_{DCA} = 500$ m is now in the range of

$$F_{DCA}(dca \leq 500 \text{ m}) = [\,1.9019e\text{-}11\%\,,\ 10.5966\%\,]\,.$$

The width of this robust probability interval, which is visualised in greater detail in Figure 7, indicates how significantly sensitive this state estimation scenario is to different specifications of the uncertainty density functions. In particular from Figure 7, we can see that for the quantity of interest, i.e. the probability of DCA less than 500 m, the computed lower and upper bounds enclose all the other probabilities compatible with the distributions within the imprecise set. It is important to underline that the distributions corresponding to lower or upper bounds for $\delta_{DCA} = 500$ m, labelled as lower or upper distributions (respectively blue and red lines in the figures), do not necessarily result in lower and upper bounds for the collision probability at different thresholds $\delta_{DCA} \neq 500$ m. Indeed, the lower and upper distributions are not lower and upper envelopes for the ECDFs resulting from the imprecise set, but they are distributions belonging to the set itself bounding the expectation of a specific quantity of interest. From Figure 6 we can notice that indeed the lower distribution does not yield to probability lower bounds for DCA higher than roughly 6-7 km. Besides, we can notice that there is no distribution which is also a lower envelope for the specified imprecise set due to the several ECDFs crossings. On the other hand, the upper distribution also coincides with the upper envelope.

The lower collision probability is extremely close to the theoretical minimum of zero probability, whereas the upper one is significantly higher than the proposal value. For both the bounds, the

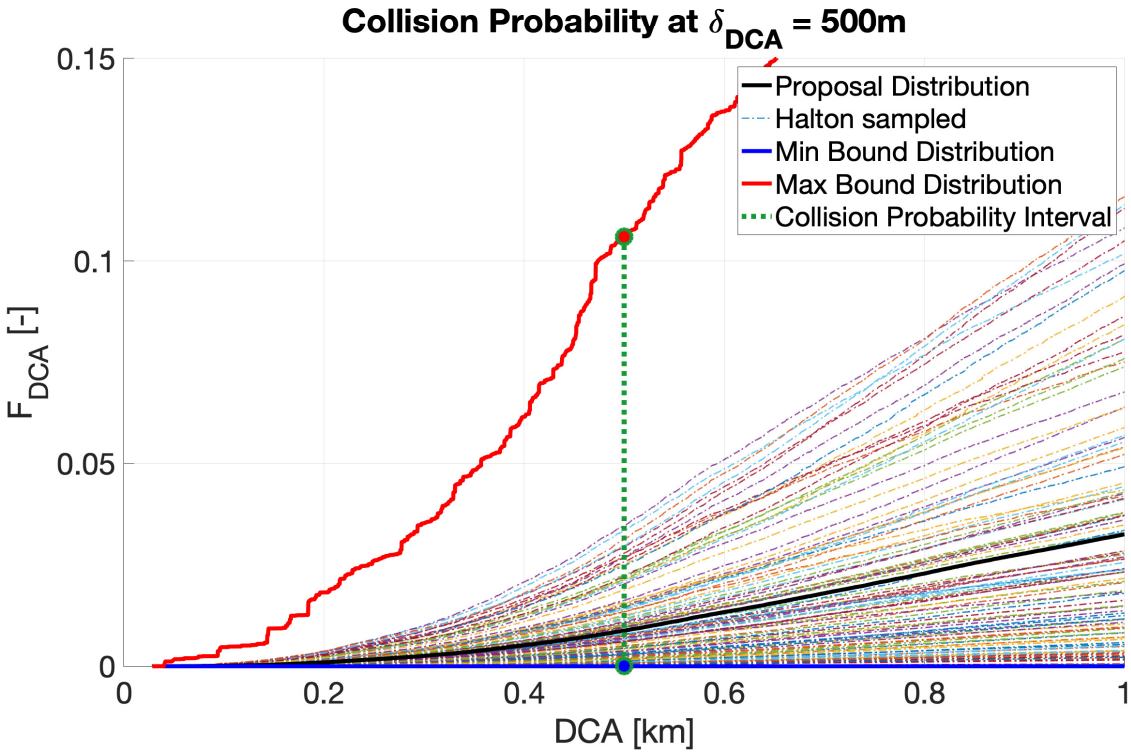

Figure 7: Collision probability interval for set threshold $\delta_{DCA} = 500$ m, and ECDFs for DCA between debris and operational satellite as resulting from the Robust Particle Filter with two observations, with prior imprecise set as in Equation (17) and likelihood imprecise set as in Equation (21).

corresponding multipliers are reported in Table 5.

Table 5: Multipliers which correspond to lower and upper collision probability as computed with the Robust Particle Filter.

Bound	$\lambda_{\mathbf{x}_{0-1}}$	$\lambda_{\mathbf{x}_{0-2}}$	$\lambda_{\mathbf{y}_{1-1}}$	$\lambda_{\mathbf{y}_{1-2}}$	$\lambda_{\mathbf{y}_{2-1}}$	$\lambda_{\mathbf{y}_{2-2}}$
Lower	1.30^2	1.45^2	0.35^2	0.66^2	1.45^2	0.33^2
Upper	0.33^2	0.33^2	1.50^2	0.50^2	0.33^2	0.73^2

For the lower bound, the multipliers $\lambda_{\mathbf{x}_{0-1}}$, $\lambda_{\mathbf{x}_{0-2}}$ and $\lambda_{\mathbf{y}_{2-1}}$ are larger than 1 to dilute the probability mass near collision trajectories, but their optimal value is not found at the upper range. On the contrary, the low values of $\lambda_{\mathbf{y}_{1-1}}$ and $\lambda_{\mathbf{y}_{2-2}}$ help to concentrate the probability around the support of an observation with high measurement error, i.e. in a state-space region which less likely will result in a collision. As for the upper bound, the behaviour is reversed for $\lambda_{\mathbf{x}_{0-1}}$, $\lambda_{\mathbf{x}_{0-2}}$ and $\lambda_{\mathbf{y}_{2-1}}$ to concentrate the probability near collision trajectories, as well as for the multiplier $\lambda_{\mathbf{y}_{1-1}}$, whose large value implies less confidence in the large-error measurement. On the other hand, the values of the multipliers $\lambda_{\mathbf{y}_{1-2}}$ and $\lambda_{\mathbf{y}_{2-2}}$ are not clearly in opposition and therefore not easily explicable. These multipliers don't assume extremes values likely due to the predominant interactions with the other free variables. Hence, both from these values and the plots, which highlights different ECDFs crossings, we can observe how the resulting set of distributions is dominated by nonlinearities, in contrast to the case when no observations are considered (see Figure 2), further justifying the need

of a dedicated optimisation routine. This gets even more apparent for more complex test cases and higher numbers of measurements.

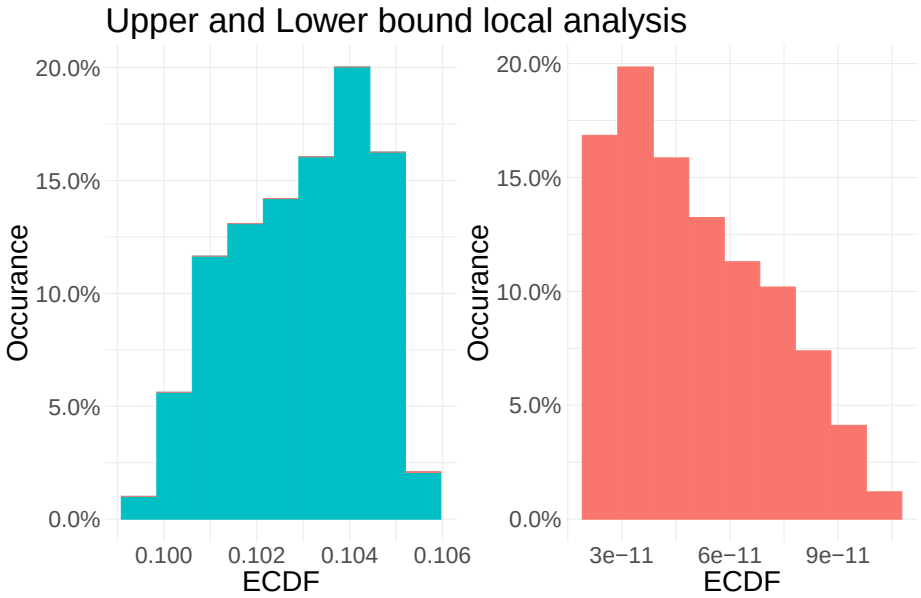

Figure 8: Histogram of the results obtained from the quasi-Montecarlo local analysis. Two clouds of 10^4 points, centred in the optimal solutions, have been generated by means Halton sampling. None of the so generated points exceeds the bounds identified by the optimisation routine.

To verify the quality of the multipliers set found, a local sensitivity analysis trough quasi-Montecarlo has been conduced over the optimal solutions. The surroundings of the solutions leading to lower and upper collision probability have been investigated with 10^4 Halton samples contained in hypercubes of side $1.5^2 \cdot 10^{-3}$ centred in the optimal solutions. The results, depicted in Figure 8, show that none of the points generated exceeds the limits found by the optimisation routine. This suggests that the optimiser correctly identified the (at least local) optimal solutions.

Again, in the context of space surveillance and tracking, this accentuated probability range indicates how sensitive the collision probability computation can be to modelling assumptions in the problem uncertainty structure in the presence of observations.

Performance Assessment

The performance of the developed Robust Particle Filter can be assessed in comparison with the precise Bootstrap Filter in Algorithm 1 on a single *Filter Evaluation* stage. One prior $p(\mathbf{x}_0)$ and one likelihood $p(\mathbf{y}_k \mid \mathbf{x}_k)$ are randomly selected within the imprecise sets. In the Boostrap Filter, the polynomial expansion computed before is used to efficiently propagate the particles. In the Robust Particle Filter, the precomputations run with the proposals in Equations (14) and (20) are exploited, just like it would be done in the *Filter Evaluation* stage. The number of particles is reduced to 10^5 samples. The results of this experiment, run with MATLAB R2018a on a MacBook Pro laptop with 3.5 GHz Intel Core i7 and 8GB of RAM running macOS Mojave v.10.14.5, are reported in Table 6.

This result shows how, once the precomputation stage has been performed, the Robust Particle Filter

Table 6: Time performance comparison between standard Particle Filter and Robust Particle Filter Approach.

Approach	Average Time $[s]$
PF	648.78
RPF	0.259

outperforms the standard PF on the evaluation of new uncertainty distributions dramatically. Since these filter evaluations are employed within an optimisation loop which calls the filter multiple times, the new RPF results crucial to the practical solution of complex scenarios.

CONCLUSIONS

This paper has presented the formulation and developed a solution approach for the filtering problem under epistemic uncertainty affecting the prior, likelihood or model parameters' uncertainty distributions. In the imprecise scenario discussed, the value of expectation of interest is enclosed in an interval. Therefore, lower and upper expectations have been introduced as the robust bounds to be computed as the solution of the imprecise filtering problem. From here, a Robust Particle Filter has been developed to efficiently solve such state estimation problems. This robust filtering approach employs precomputation performed with a standard filter and importance distributions for the prior, likelihood, transition and main proposal, to speed up the successive numerous filter evaluations required by the bounds computation routine, which is faced by optimisation. The main idea of the RPF is to exploit the already propagated samples, drawn from proposal distributions, as much as possible, and evaluate inexpensively the final expectation for new input distributions by re-scaling the particles' importance weights to target the correct distribution.

This approach has been applied to the computation of the collision probability between an operational satellite and a debris in LEO environment in the presence of epistemic uncertainty affecting the prior and likelihood distributions. A generalised polynomial expansion has been employed as surrogate model for efficient uncertainty propagation in order to increase the number of particles in the filter. The results have shown how the RPF is able to efficiently compute robust probability bounds in such scenario. Furthermore, they have also pointed out how the collision probability is severely sensitive to parameters modelling the uncertainty distributions, highlighting the usefulness of and need for robust filtering. Finally, the computational saving is assessed by comparing the filter evaluation stage using the RPF versus a standard PF. For the presented case study, the RPF leads to an improvement of more than 3 orders of magnitude in the computational time, a critical feature since the filter has to be called multiple times within the bound computation routine. In general, the ability of the Robust Particle Filter in delivering more reliable estimates with efficient computations can be critical in a variety of highly uncertain scenarios.

Future work will focus on the integration of efficient resampling techniques to perform within the Filter Evaluation step based on a trade-off between the number of new particles to propagate and the number of degenerate particles. Such further development would allow one to consider even more complex parameterisation for the imprecise sets. Furthermore, theoretical analyses will be performed to characterise convergence properties of the developed Robust Particle Filter.

ACKNOWLEDGEMENTS

This work was funded by the European Commission's H2020 programme, through the H2020-MSCA-ITN-2016 UTOPIAE Marie Curie Innovative Training Network, grant agreement 722734.

REFERENCES

[1] T. Augustin, F. P. A. Coolen, G. De Cooman, and M. Troffaes, *Introduction to Imprecise Probabilities*. John Wiley & Sons, 2014, doi:10.1002/9781118763117.

[2] A. H. Jazwinski, *Stochastic Processes and Filtering Theory*, Vol. 64 of *Mathematics in Science and Engineering*. Bellman R. Academic Press, New York, 1st ed., 1970.

[3] S. Sarkka, *Bayesian Filtering and Smoothing*. Cambridge University Press, New York, 1st ed., 2013.

[4] P. Kozierski, M. Lis, and J. Ziętkiewicz, Resampling in Particle Filtering - Comparison, *Studia z Automatyki i Informatyki*, 2013, pp. 35–64.

[5] N. J. Gordon, D. J. Salmond, and A. F. Smith, Novel approach to nonlinear/non-Gaussian Bayesian state estimation, *IEE proceedings F (radar and signal processing)*, Vol. 140 (2), 1993, pp. 107–113, doi:10.1049/ip-f-2.1993.0015.

[6] SOCRATES - Satellite Orbital Conjunction Reports Assessing Threatening Encounters in Space, http://celestrak.com/SOCRATES/. Accessed: 10-Jul-2019.

[7] O. Montenbruck and E. Gill, *Satellite orbits: models, methods and applications*. Springer Science & Business Media, 1st ed., 2000.

[8] M. Vasile, C. A. Ortega, and A. Riccardi, Set propagation in dynamical systems with generalised polynomial algebra and its computational complexity, *Communications in Nonlinear Science and Numerical Simulation*, Vol. 75, 2019, pp. 22–49, doi:10.1016/j.cnsns.2019.03.019.

[9] SMART-UQ - Strathclyde Mechanical and Aerospace Research Toolbox for Uncertainty Quantification, https://github.com/strath-ace/smart-uq. Accessed: 22-Jul-2019.

[10] A. Riccardi, C. Tardioli, and M. Vasile, An intrusive approach to uncertainty propagation in orbital mechanics based on Tchebycheff polynomial algebra, *Advances in Astronautical Sciences*, 2015, pp. 707–722. https://strathprints.strath.ac.uk/60560/.

[11] C. A. Ortega, R. Serra, A. Riccardi, and M. Vasile, De-orbiting and re-entry analysis with generalised intrusive polynomial expansions, *67th International Astronautical Congress*, 2016. https://strathprints.strath.ac.uk/60596/.

[12] M. Galassi, T. J. Davies, J., B. Gough, G. Jungman, P. Alken, M. Booth, F. Rossi, and R. Ulerich, *GNU Scientific Library*. GNU, 2018. Release 2.5.

[13] H. Klinkrad, J. R. Alarcon, and N. Sanchez, Collision avoidance for operational ESA satellites, *4th European Conference on Space Debris*, Vol. 587, 2005.

[14] M. Di Carlo, M. Vasile, and E. Minisci, Adaptive multi-population inflationary differential evolution, *Soft Computing*, 2019, pp. 1–31.

[15] M. Di Carlo, M. Vasile, and E. Minisci, Multi-population inflationary differential evolution algorithm with Adaptive Local Restart, *2015 IEEE Congress on Evolutionary Computation (CEC)*, IEEE, 2015, pp. 632–639.

[16] K. Price, R. M. Storn, and J. A. Lampinen, *Differential evolution: a practical approach to global optimization*. Springer Science & Business Media, 2006.

[17] D. J. Wales and J. P. Doye, Global optimization by basin-hopping and the lowest energy structures of Lennard-Jones clusters containing up to 110 atoms, *The Journal of Physical Chemistry A*, Vol. 101, No. 28, 1997, pp. 5111–5116.