Robust Strategies for Incorporating Parameter Uncertainty in Constrained Admissible Regions

Thomas L. Dearing, Jacob D. Griesbach
ARKA Group - Stratagem
Piyush Mehta
University of West Virginia

ABSTRACT

In furtherance of IARPAs SINTRA debris detection and tracking program, this work examines alternative strategies for incorporating parameter uncertainty in the Constrained Admissible Region (CAR) generation algorithm to address practical issues encountered in large scale Multi-Hypothesis Filter applications. This work first reproduces and expands upon the differential inflation approach developed by Worthy et al. to include both prevalent CAR constraint bounds on orbital eccentricity and semi-major axis. Strategies are developed to reduce intrinsic numerical ill-conditioning of this approach under these constraints, addressing key shortcomings reported in the original literature. To examine the higher order moments of the parameter-uncertainty inflated CAR, this work also adapts the Unscented Transform (UT) to measure nonlinear distortions in the CAR boundary using synchronized parametrizations of the solution contours. The accuracy and numerical efficiency of the above approaches are then compared against Monte-Carlo simulations in both the LEO and GEO orbital regimes. This comparison yields the following key conclusions: (1) the size and shape of the CAR is non-negligibly affected by the inclusion of parameter error in both LEO and GEO regimes, (2) the UT approach reproduces the nonlinear inflation of the CAR more accurately than the differential approach, but is more computationally intensive, and (3) results from these two approaches mainly diverge near saddle points in the constraint functions. These results highlight the importance of parameter error in admissible region generation and present multiple approaches to combine the respective robustness and computational efficiency benefits of the Probabilistic Admissible Region and CAR approaches in practical applications.[1]

1. INTRODUCTION

A fundamental first step for any Kalman filter based Initial Orbit Determination (IOD) strategy is the generation of an initial orbit estimate used for Kalman Filter initialization: a hypothesis necessarily based on sparse and imprecise information. Unlike the IOD approaches which batch-fit large quantities of pre-associated measurements [8], a *Multi-Hypothesis Filter* (MHF) refines the orbital estimate continuously, using the initial detection to generate a family of initial orbit hypotheses that sample the measurement uncertainty space [5]. This set of initial hypotheses must be carefully chosen, sampling the state space both finely and comprehensively enough to guarantee that the true state lies within at least one hypotheses convergence radius (as governed by the underlying Kalman filter). However, the computational impact of these hypotheses is also considerable: each must be propagated and compared against subsequent observations and new hypotheses must be generated whenever the MHF encounters a new unassociated observation. Thus, the MHFs hypothesis set must simultaneously be large enough to contain the true solution, but small enough to remain computationally tractable and scalable while processing thousands of Resident Space Objects (RSOs).

Initializing the MHF hypothesis set becomes particularly challenging when considering Electro-Optical (EO) measurements. Such observations are extracted from short arcs (streaks) in the image plane and nominally provide Right-Ascension/Declination angle pairings, with angle-rates either measured natively or inferred by differencing successive measurements. However, angles and angle-rates only encompass 4 of the 6 Degrees of Freedom (DOF) necessary to compute a full position/velocity orbital state. That is, the range and range-rate space is entirely unobservable over

---

[1] This research is based upon work supported in part by the Office of the Director of National Intelligence (ODNI), Intelligence Advanced Research Projects Activity (IARPA), via 2023-23060200005. The views and conclusions contained herein are those of the authors and should not be interpreted as necessarily representing the official policies, either expressed or implied, of ODNI, IARPA, or the U.S. Government. The U.S. Government is authorized to reproduce and distribute reprints for governmental purposes notwithstanding any copyright annotation therein.

these arcs. To capture the remaining 2 DOF using a finite MHF hypothesis set, the feasible range/range-rate space must be refined to a finite subset or Admissible Region (AR). Two major approaches for defining an AR currently exist in the literature. The first approach is called the Constrained AR (CAR) and uses *a-priori* orbital constraints to include only the orbital families of interest: for example, orbits within an assumed semi-major axis range [2, 5]. This approach rapidly converts a set of observation parameters and orbital constraints into a compact region that can be easily subsampled to return hypotheses for the MHF. The second approach is called the Probabilistic AR (PAR) and is constructed using Gaussian particle clouds to directly estimate the measurement uncertainty distribution in the range/range-rate space [3]. Critically, this inclusion of measurement uncertainty in the construction of the PAR makes it notably larger and more accurate than the equivalent CAR, with the associated distribution also enabling users to sample the PAR more strategically. However, the large number of samples required to directly and accurately estimate the PAR distribution make it extremely computationally intensive. As a result, practical selection between these two approaches often simplifies to the applications prioritization of computational efficiency and robustness.

In furtherance of IARPAs SINTRA debris detection and tracking program, this work examines alternative strategies for incorporating parameter uncertainty in the CAR generation algorithm to address practical issues encountered in large scale MHF applications. Specifically, this work first reproduces and expands upon the Differential Inflation (DI) approach developed by [9] to include both prevalent CAR constraint bounds on orbital eccentricity and semi-major axis. Strategies are developed to reduce intrinsic numerical ill-conditioning of this approach under these constraints, addressing key shortcomings reported in [9]. To examine the higher order moments of the parameter-uncertainty inflated CAR, this work also adapts the Unscented Transform (UT) to measure nonlinear distortions in the CAR boundary using synchronized parameterizations of the solution contours. The accuracy and numerical efficiency of the above approaches are then compared against Monte-Carlo simulations in both the LEO and GEO orbital regimes. This comparison yields the following key conclusions:

1. The size and shape of the CAR is *massively* expanded by the inclusion of standard EO observation error in both the GEO and LEO regimes (3-10×),
2. In most cases, both the DI and UT approaches accurately model the CARs inflation ($\leq 0.25\%$ relative error),
3. Both approaches become unreliable near saddle points in the constraint surfaces (where the CAR boundary is more bimodal than Gaussian. In this case, the UT method better estimates the CAR's (discontinuous) expansion.

These results highlight the importance of including observation parameter error when generating the admissible region and present multiple approaches to combine the respective robustness and computational efficiency benefits of the PAR and CAR approaches in practical applications.

The paper is organized as follows: Section 2 introduces the EO Constrained Admissible Region, including observation parameters and prevalent orbital constraints, Section 3 presents existing (differential), novel (unscented), and reference truth (Monte-Carlo) approaches for modeling and mapping observation parameter errors into a CAR boundary contour distribution, Section 4 presents numerical comparisons of the efficiency and accuracy of these approaches for representative RSO observations at GEO and LEO, and Section 5 summarizes the paper with concluding remarks.

2. PRELIMINARIES

This section reviews fundamental definitions and existing approaches for constructing the CAR for EO measurements. Additionally, this section reviews definitions for the two prevalent CAR constraints bounding the feasible orbit Semi-Major Axis (SMA) and eccentricity.

2.1 Fundamentals of Electro-Optical Sensors

As discussed above, EO Sensors represent a critical source for detections of RSO's and form a key component of the U.S. Space Surveillance Network [6]. Fundamentally, EO sensors passively capture and process images to determine the position and velocity of a detected RSO streak in the sensor's pixel space. Using the sensor's orientation and coordinate frame, the position and velocity in pixel space are then projected onto the celestial sphere to yield the Right Ascension α and Declination δ, as well as the respective angular rates $\dot{\alpha}, \dot{\delta}$. The remaining degrees of freedom orthogonal to the image plane, the range ρ and range-rate $\dot{\rho}$, are unobservable within a single detection.

Without loss of generality, let the observing station and RSO positions in the ECI frame be given by the vectors $\boldsymbol{q}, \boldsymbol{r} \in \mathbb{R}^3$ respectively, with the relative (observed) position of the RSO given by the relative vector $\boldsymbol{p} = \boldsymbol{r} - \boldsymbol{q}$. The celestial observation frame for the sensor is then given by the basis[2] vectors

$$\hat{\boldsymbol{p}} := \begin{bmatrix} \cos(\alpha)\cos(\delta) \\ \sin(\alpha)\cos(\delta) \\ \sin(\delta) \end{bmatrix}, \quad \hat{\boldsymbol{\alpha}} := \cos(\delta) \begin{bmatrix} -\sin(\alpha) \\ \cos(\alpha) \\ 0 \end{bmatrix}, \quad \hat{\boldsymbol{\delta}} := \begin{bmatrix} -\cos(\alpha)\sin(\delta) \\ -\sin(\alpha)\sin(\delta) \\ \cos(\delta) \end{bmatrix}. \tag{1}$$

The observation coordinates are then related to the RSO's relative ECI position and velocity as follows

$$\begin{aligned} \boldsymbol{p} &:= \rho \hat{\boldsymbol{p}}, \\ \dot{\boldsymbol{p}} &= \dot{\rho}\hat{\boldsymbol{p}} + \rho\left(\dot{\alpha}\hat{\alpha} + \dot{\delta}\hat{\delta}\right), \end{aligned} \tag{2}$$

where, as noted above, an estimate (or hypothesis) of the RSO's range ρ and range-rate $\dot{\rho}$ are required to fully determine it's orbital state.

2.2 The Constrained Admissible Region

Originally presented in [4], the Constrained Admissible Region (CAR) is a generalized approach to constrain the unobservable degrees of freedom $\boldsymbol{x}_u \in \mathbb{R}^u$ in a nonlinear observation model using a-priori feasibility constraints on the complete system state $\boldsymbol{x}^\top := [\boldsymbol{x}_d^\top, \boldsymbol{x}_u^\top] \in \mathbb{R}^n$. Specifically, under the nonlinear observation model

$$\boldsymbol{y} = \boldsymbol{o}(\boldsymbol{x}_d, \boldsymbol{k}), \tag{3}$$

we obtain observations $\boldsymbol{y} \in \mathbb{R}^m$ which depend only on the *observable* (or determined) component of the state $\boldsymbol{x}_d \in \mathbb{R}^m$ and the observation parameters $\boldsymbol{k} \in \mathbb{R}^\ell$. Reasonably assuming that $\boldsymbol{o}$ is bijective and continuously differentiable, the determined state component is uniquely determined by the inverse mapping

$$\boldsymbol{x}_d = \boldsymbol{o}^{-1}(\boldsymbol{y}, \boldsymbol{k}). \tag{4}$$

While the *unobservable* state component $\boldsymbol{x}_u$ cannot be determined directly via $\boldsymbol{o}$, one can constrain the potential values of $\boldsymbol{x}_u$ to a finite subset of $\mathbb{R}^u$ via feasibility constraints on the complete state vector $\boldsymbol{x}$. Specifically, a continuously differentiable nonlinear constraint of the form

$$\boldsymbol{g}(\boldsymbol{x}_u, \boldsymbol{x}_d, \boldsymbol{k}) \leq 0, \tag{5}$$

can be combined with (4) to yield the CAR constraint

$$\kappa(\boldsymbol{x}_u, \boldsymbol{z}) := \boldsymbol{g}(\boldsymbol{o}^{-1}(\boldsymbol{x}_u, \boldsymbol{y}, \boldsymbol{k}), \boldsymbol{k}), \tag{6}$$

where the vector $\boldsymbol{z}^\top := [\boldsymbol{y}^\top, \boldsymbol{k}^\top] \in \mathbb{R}^{m+\ell}$ combines the observation and observation parameters. Under the constraint (6) and observation $\boldsymbol{z}$, the CAR $\mathcal{D}(\boldsymbol{z})$ is then defined by

$$\mathcal{D}(\boldsymbol{z}) := \{\boldsymbol{x}_u : \kappa(\boldsymbol{x}_u, \boldsymbol{z}) \leq 0\},$$

In the absence of other sensor data, the region[3] $\mathcal{D}(\boldsymbol{z})$ may be used to define a uniform distribution for the unknown state component $\boldsymbol{x}_u$, and can be subsampled to return candidate hypotheses for $(\rho, \dot{\rho})$. Finally, combining these samples with the known state $\boldsymbol{x}_d$ yields feasible candidate hypotheses $\boldsymbol{x}$ which may be used to initialize a Multi-Hypothesis Filter or other associator.

While the above CAR notions are generalized to any nonlinear observation model, its application in existing literature has overwhelmingly focused on the short-angle EO observations considered in this work. Assuming the observations are pre-compressed to include angle-rate data, the state elements for this use case are defined as follows

$$\boldsymbol{x}_d^\top = \boldsymbol{y}^\top := \left[\alpha, \delta, \dot{\alpha}, \dot{\delta}\right], \quad \boldsymbol{x}_u^\top := [\rho, \dot{\rho}], \quad \boldsymbol{k}^\top := \left[\boldsymbol{q}^\top, \dot{\boldsymbol{q}}^\top\right]. \tag{7}$$

Note that unlike [9], (7) does not include observation time as an explicit observation parameter, as the effects of sample time uncertainty are not considered in this work (though they can be included in the above framework during compression). The two most prevalent CAR constraints in current literature are detailed in the following sections.

---

[2] Note that $\hat{\alpha}$ is not a true unit vector. The additional cosine term is traditionally included to simplify notation.

[3] Note that,depending on the underlying constraint surface, $\mathcal{D}(z)$ may contain multiple isolated connected components.

2.2.1 CAR Semi-Major-Axis Constraint

The first and arguably most prevalent CAR constraint is defined by bounding the RSO's feasible semi-major axis to the range $a \in [a^-, a^+]$ [4]. These bounds are then related to the RSO's position and velocity state via the Keplerian two-body specific energy equation

$$E := \frac{\|\dot{\boldsymbol{r}}\|^2}{2} - \frac{\mu}{\|\boldsymbol{r}\|}, \tag{8}$$

Using (2) to recoordinatize (8) into the EO observation frame and isolating the range and range-rate terms, we arrive at the polynomial form:

$$2E = \dot{\rho}^2 + w_1\dot{\rho} + F(\rho), \tag{9}$$

where

$$F(\rho) := w_2\rho^2 + w_3\rho + w_4 - \frac{2\mu}{\sqrt{\rho^2 + w_5\rho + w_0}}, \tag{10}$$

and

$$\begin{aligned} w_0 &:= \|\boldsymbol{q}\|^2, & w_1 &:= 2(\dot{\boldsymbol{q}} \cdot \hat{\boldsymbol{p}}), \\ w_2 &:= \dot{\alpha}^2 \cos^2(\delta) + \dot{\delta}^2, & w_3 &:= 2\dot{\alpha}(\dot{\boldsymbol{q}} \cdot \hat{\boldsymbol{\alpha}}) + 2\dot{\delta}\left(\dot{\boldsymbol{q}} \cdot \hat{\boldsymbol{\delta}}\right), \\ w_4 &:= \|\dot{\boldsymbol{q}}\|^2, & w_5 &:= 2(\boldsymbol{q} \cdot \hat{\boldsymbol{p}}). \end{aligned} \tag{11}$$

Next, recall that the two-body specific energy is also related to the orbit's semi-major axis $a \in \mathbb{R}_{\geq 0}$ via

$$E(a) := -\frac{\mu}{2a}, \tag{12}$$

where $E(a) < 0$ implies a bounded orbit. Combining this with (9) and re-arranging, we arrive at the constraint function

$$\kappa_a(\boldsymbol{x}_u, z, a) := \dot{\rho}^2 + w_1\dot{\rho} + F(\rho) + \frac{\mu}{a}. \tag{13}$$

Under these definitions, the CAR $\mathcal{D}_a(z, a^+)$ defined from $\kappa_a(\boldsymbol{x}_u, z, a^+) \leq 0$ specifies the feasible values of $\boldsymbol{x}_u = (\rho, \dot{\rho})$ for a bounded orbit with a semi-major axis $a \leq a^+$ under the observation z. Examples for the CAR resulting from (13) in GEO and LEO scenarios are shown in Fig. 1. Although the shape of the CAR's boundary varies substantially between these cases, we can quickly identify two important properties for CARs produced under this constraint:

1. $\mathcal{D}_a(z, a)$ has only *one* simply connected component.
2. $\mathcal{D}_a(z, a)$ is symmetric about a fixed $\dot{\rho}^*$ for a given observation.

These properties follow directly from (13) being a concave up quadratic in $\dot{\rho}$ with ρ and a only affecting the vertical offset. These properties make this CAR comparatively simple to model numerically, but also allow the resulting CAR to become quite large.

Thankfully, one can usually also infer a *lower* bound for the SMA of feasible orbits. In this case, the constraint (13) can also be used to enforce the lower bound $a \geq a^-$ in the form

$$-\kappa_a(\boldsymbol{x}_u, z, a^-) \leq 0.$$

Because both SMA constraints share the same quadratic form, CARs produced by each (and the combination of both) share the properties described above. This work will consider both upper and lower form SMA constraints, though the former will be prioritized in discussion due to their intrinsic similarities.

2.2.2 CAR Eccentricity Constraint

Like with the orbit's Semi-Major Axis, bounds on orbital eccentricity e can also be used to constrain the feasible range and range-rate space [1]. Specifically, an orbit's eccentricity can be related to its two-body specific energy E and orbital angular momentum $\boldsymbol{h}$ via the relation

$$-\mu^2(1 - e^2) = 2E\|\boldsymbol{h}\|^2. \tag{14}$$

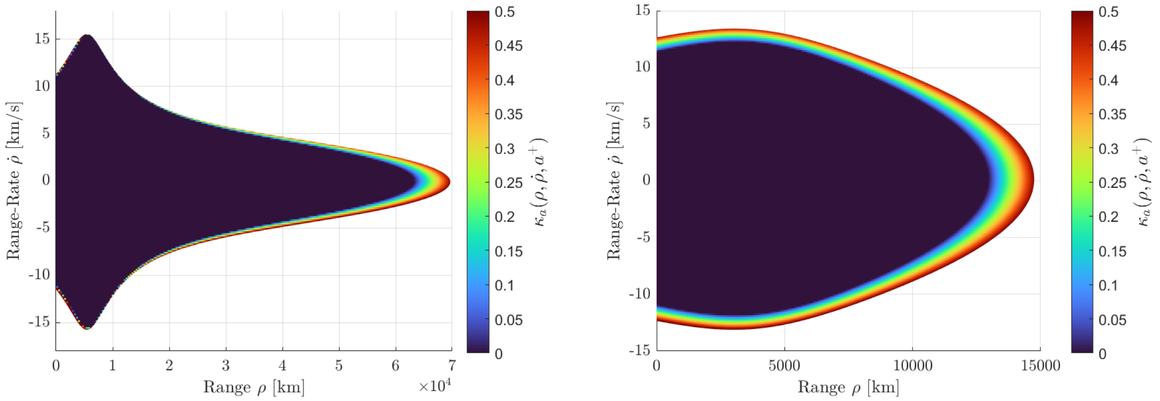

Fig. 1: Surface plots of the ($a^+ = \infty$) SMA constraint (13) for the (a) GEO and (b) LEO scenarios specified in table 1. The CAR itself is colored in black, while the curvature of the constraint surface is highlighted in color.

The angular momentum $\boldsymbol{h} := \boldsymbol{r} \times \dot{\boldsymbol{r}}$ can also be written in the EO observation frame as

$$\|\boldsymbol{h}\|^2 = c_0\dot{\rho}^2 + P(\rho)\dot{\rho} + U(\rho), \tag{15}$$

where the polynomial functions $P(\rho), U(\rho) \in \mathbb{R}$ are given by

$$P(\rho) := c_1\rho^2 + c2\rho + c_3, \tag{16a}$$

$$U(\rho) := c_4\rho^4 + c_5\rho^3 + c_6\rho^2 + c_7\rho + c_8, \tag{16b}$$

and the coefficients c_i and momentum terms $\boldsymbol{h}_i$ are defined as follows

$$\begin{aligned} c_0 &:= \|\boldsymbol{h}_1\|^2, & c_1 &:= 2\boldsymbol{h}_1 \cdot \boldsymbol{h}_2, & c_2 &:= 2\boldsymbol{h}_1 \cdot \boldsymbol{h}_3, \\ c_3 &:= 2\boldsymbol{h}_1 \cdot \boldsymbol{h}_4, & c_4 &:= \|\boldsymbol{h}_2\|^2, & c_5 &:= 2\boldsymbol{h}_2 \cdot \boldsymbol{h}_3, \\ c_6 &:= 2\boldsymbol{h}_2 \cdot \boldsymbol{h}_4 + \|\boldsymbol{h}_3\|^2, & c_7 &:= 2\boldsymbol{h}_3 \cdot \boldsymbol{h}_4, & c_8 &:= \|\boldsymbol{h}_4\|^2, \end{aligned} \tag{17}$$

$$\begin{aligned} \boldsymbol{h}_1 &:= \boldsymbol{q} \times \hat{\boldsymbol{p}}, & \boldsymbol{h}_2 &:= \hat{\boldsymbol{p}} \times \left(\dot{\alpha}\hat{\boldsymbol{\alpha}} + \dot{\delta}\hat{\boldsymbol{\delta}}\right), \\ \boldsymbol{h}_3 &:= \hat{\boldsymbol{p}} \times \dot{\boldsymbol{q}} + \boldsymbol{q} \times \left(\dot{\alpha}\hat{\boldsymbol{\alpha}} + \dot{\delta}\hat{\boldsymbol{\delta}}\right), & \boldsymbol{h}_4 &:= \boldsymbol{q} \times \dot{\boldsymbol{q}}. \end{aligned} \tag{18}$$

Finally, (14) can be expanded using (9) and (15) and re-arranged to yield the quartic eccentricity constraint

$$\begin{aligned} \kappa_e(\boldsymbol{x}_u, \boldsymbol{z}, e) &:= 2E\|\boldsymbol{h}\|^2 + \mu^2(1 - e^2) \\ &= a_4\dot{\rho}^4 + a_3\dot{\rho}^3 + a_2\dot{\rho}^2 + a_1\dot{\rho} + a_0, \end{aligned} \tag{19}$$

where the quartic polynomial coefficients a_i are given by

$$\begin{aligned} a_0 &= F(\rho)U(\rho) + \mu^2(1 - e^2), & a_1 &= F(\rho)P(\rho) + w_1U(\rho), \\ a_2 &= U(\rho) + c_0F(\rho) + w_1P(\rho), & a_3 &= P(\rho) + c_0w_1, \\ a_4 &= c_0. \end{aligned} \tag{20}$$

As with the SMA constraint (13), the eccentricity constrained region $\kappa_e(\boldsymbol{x}_u, \boldsymbol{z}, e^+) \le 0$ specifies the feasible values of $\boldsymbol{x}_u = (\rho, \dot{\rho})$ for a bounded orbit with eccentricity $e \in [0, e^+]$ under the observation $\boldsymbol{z}$. Unlike with the SMA constraint however, the CARs $\mathscr{D}_e(\boldsymbol{z}, e)$ resulting from (19) are neither symmetric nor simply connected, vary quite drastically between GEO and LEO observations (see Fig. 2 for examples), and are generally much smaller. Also unlike the SMA constraint, a lower bound e^- is not normally enforced, as circular orbits are nearly always considered admissable. This combination of factors makes the eccentricity constraint challenging to implement numerically, but also highly effective for drastically reducing the feasible hypothesis space.

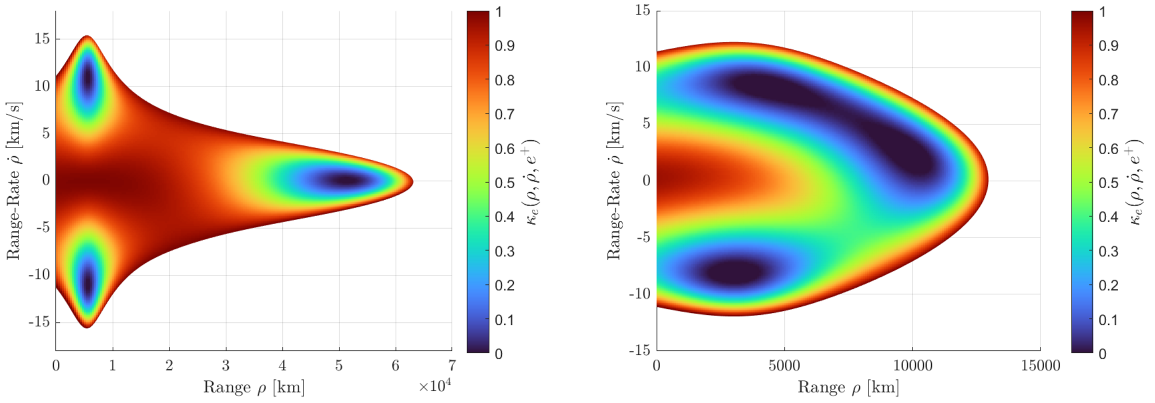

Fig. 2: Surface plots of the eccentricity constraint (19) for the (a) GEO ($e^+ = 0.08$) and (b) LEO ($e^+ = 0.2$) scenarios specified in table 1. The CAR is colored in black, while the curvature of the constraint surface is highlighted in color.

3. APPROACH

This section introduces the framework for incorporating observation parameter errors into a CAR using the distribution formed by its boundary contour. Specifically, it presents two approaches for estimating this distribution: (a) an existing 1st order approach presented by [9] which we refer to as *Differential Inflation*, and (b) an alternative approach employing an adaptation the Unscented Transform (UT). To determine the efficacy of these methods, a Monte-Carlo approach is presented for robustly estimating the true contour distribution. Finally, practical considerations and recommendations are discussed for implementing these algorithms.

3.1 Modeling CAR Parameter Error

While the original CAR definition presented in section 2.2 is highly effective at limiting the hypothesis space under appropriately chosen constraints, the assumption of *zero* observation error in its construction leads it to be inherently undersized. In reality, the source observation $\boldsymbol{y}$, observing sensor state parameters $\boldsymbol{k}$, and even the observation time t are all uncertain, all of which can substantially affect both the size and shape of the resulting CAR (e.g., consider the light blue regions in Fig. 2). To examine this effect in more detail, consider the Taylor series expansion of the CAR constraint (6) under the perturbations $\Delta\boldsymbol{x}_u$ and Δz:

$$\kappa(\boldsymbol{x}_u + \Delta\boldsymbol{x}_u, z + \Delta z) = \kappa(\boldsymbol{x}_u, z) + \frac{\partial\kappa}{\partial\boldsymbol{x}_u}\Delta\boldsymbol{x}_u + \frac{\partial\kappa}{\partial z}\Delta z + \cdots. \tag{21}$$

Like its defining constraint, the CAR $\mathcal{D}(z + \Delta z)$ is affected by this perturbation, contracting or inflating along its boundary $\partial\mathcal{D}$ depending on the specific perturbation Δz. For numerical purposes, it is both preferable and sufficient to model these perturbations in the CAR using its boundary, defined by

$$\partial\mathcal{D}(z + \Delta z) := \{\boldsymbol{s} : \kappa(\boldsymbol{s}, z + \Delta z) \equiv 0\}. \tag{22}$$

Through (21) and (22), a distribution of parameter errors $\delta\boldsymbol{Z}$ maps to a distribution of CAR boundary solutions. To capture this distribution numerically, let the boundary for each isolated (simply connected) CAR region[4] be parameterized as a periodic contour $\boldsymbol{s}(\tau)$ with $\tau \in [-1, 1]$. [5] Under this framework, the CAR boundary distribution can be described using the same parameterization: for each τ, the distribution of boundary contours $\delta\boldsymbol{s}$ can be described pointwise by the Gaussian $N(\boldsymbol{s}(\tau), P_s(\tau))$. The CAR boundary may also be *inflated* at each τ by projecting the covariance $P_s(\tau)$ along a direction $\hat{\boldsymbol{n}}(\tau)$ chosen 'outward' from the original boundary. For example, the traditional 3σ confidence interval can be captured by the contour:

$$\boldsymbol{s}_{3\sigma}(\tau) := \boldsymbol{s}(\tau) + 3\sigma_s(\tau)\hat{\boldsymbol{n}}(\tau), \tag{23}$$

---

[4] Distinct regions where the constraint surface $\kappa \leq 0$.

[5] The parameter τ is not necessary for the CAR boundary (22) definition, but adds the necessary vector space structure for modeling perturbations in the boundary contours (i.e. $\partial\mathcal{D}(z + \Delta z) - \partial\mathcal{D}(z)$ is not well defined). Approaches for selecting τ for the $u \leq 2$ case are discussed in later sections.

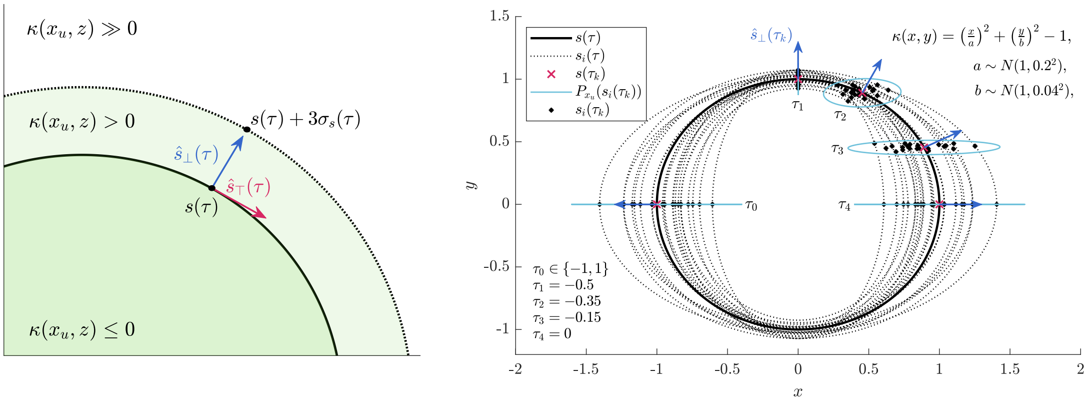

Fig. 3: Illustrations of modeling CAR parameter error inflation using boundary contours $\boldsymbol{s}(\tau)$. Panel (a) depicts the inflation of the CAR using a known step size and direction, while panel (b) depicts measuring this distribution on an example constraint using sample contours obtained via Monte-Carlo or unscented methods.

where the projected variance $\sigma_s^2(\tau) \in \mathbb{R}_{>0}$ is given by

$$\sigma_s^2(\tau) := \hat{\boldsymbol{n}}(\tau)^\top P_s(\tau)\hat{\boldsymbol{n}}(\tau), \tag{24}$$

An example of this inflation is illustrated in Fig. 3a, where the vector $\hat{\boldsymbol{n}}(\tau) = \boldsymbol{s}_\perp(\tau)$ orthogonal to the original boundary contour is chosen for the inflation direction: a standard, effective choice employed throughout this work.[6] The remainder of this section discusses potential (existing and novel) approaches for estimating the distribution $N(\boldsymbol{s}(\tau), P_{\boldsymbol{s}}(\tau))$.

3.2 CAR Differential Inflation

The first and most prevalent approach for estimating the CAR boundary distribution $N(\boldsymbol{s}(\tau), P_{\boldsymbol{s}}(\tau))$ is introduced in [9], and denoted here as the *Differential Inflation* (DI) approach. This method uses the gradient of the CAR constraint to linearly map perturbations in the observation parameters ($\Delta\boldsymbol{z}$) to perturbations in the unobservable space ($\Delta\boldsymbol{x}_u$). Specifically, consider the 1st order Taylor series (21) of the CAR constraint and recall that the CAR boundary is defined by the identity $\kappa(\boldsymbol{x}_u + \Delta\boldsymbol{x}_u, \boldsymbol{z} + \Delta\boldsymbol{z}) \equiv 0$. Combining these, (21) can be approximated to the local linear map

$$-\frac{\partial\kappa}{\partial\boldsymbol{x}_u}\Delta\boldsymbol{x}_u = \frac{\partial\kappa}{\partial\boldsymbol{z}}\Delta\boldsymbol{z}, \tag{25}$$

which relates the (first order) sensitivity of the observation parameters to the unobservable space through variations in the constraint function κ. However, as the linear map (25) is defined using the gradient of κ, it only relates variations orthogonal to the constraint boundary (i.e. along $\boldsymbol{s}_\perp(\tau) = {}^{\partial\kappa}/_{\partial\boldsymbol{x}_u}$) and thus is inherently rank deficient. The remaining null space (tangent to the boundary contour) can be spanned by the rows of any orthonormal matrix $T \in \mathbb{R}^{u-1\times u}$ satisfying $T\Delta\boldsymbol{x}_u = 0$. Combining these yields the following full rank version of (25)

$$-\begin{bmatrix} \frac{\partial\kappa}{\partial\boldsymbol{x}_u} \\ T \end{bmatrix}\Delta\boldsymbol{x}_u = \begin{bmatrix} \frac{\partial\kappa}{\partial\boldsymbol{z}} \\ 0 \end{bmatrix}\Delta\boldsymbol{z}. \tag{26}$$

Applying the linearity of Gaussian random variables, the perturbations $\Delta\boldsymbol{z}$ and $\Delta\boldsymbol{x}_u$ can be replaced by the error distributions $\delta\boldsymbol{Z} \sim N(0, P_z)$ and $\delta\boldsymbol{X}_u$ respectively. Thus, (26) can be written as

$$\delta\boldsymbol{X}_u = \underbrace{-\begin{bmatrix} \frac{\partial\kappa}{\partial\boldsymbol{x}_u} \\ T \end{bmatrix}^{-1}\begin{bmatrix} \frac{\partial\kappa}{\partial\boldsymbol{z}} \\ 0 \end{bmatrix}}_{M(\boldsymbol{x}_u, \boldsymbol{z}):=}\delta\boldsymbol{Z}, \tag{27}$$

---

[6] The alternate inflation direction $\boldsymbol{n}(\tau) = R\sqrt{D}R^\top\boldsymbol{s}_\perp(\tau)$ (where $RDR^\top$ diagonalizes the covariance $P_s(\tau)$) was also found to be effective, better tracking more drastic changes in the eccentricity CAR. However, it is not well supported analytically and so is omitted in this discussion.

which, along the CAR boundary contour $\boldsymbol{s}(\tau)$, reduces to the compact form

$$\sigma_s(\tau)\,\hat{\boldsymbol{s}}_\perp(\tau) = M(\boldsymbol{s}(\tau), z)\,\delta\boldsymbol{Z}. \tag{28}$$

Note that (28) directly yields the variance $\sigma_s^2(\tau)$ in lieu of a traditional covariance $P_s(\tau)$ (i.e. the step direction $\hat{\boldsymbol{n}}(\tau)$ in (24) is fixed). This is not necessarily a shortcoming, as the requirement for a full covariance matrix is introduced by the parameterization τ of the CAR boundary: a structure which the generic DI method in (27) does not inherently require. In practice, the DI method's greatest limitation is also a strength: the restriction to a linear approximation of κ makes DI extremely computationally efficient, but also omits any nonlinear effects or features (e.g. extremal points) which may substantially affect the shape of the inflated CAR. This trade-off is reflected in the results for the eccentricity constraint CAR presented later in paper.

3.3 The CAR Unscented Transform

While a linear approximation of a nonlinear function is frequently sufficient for modeling sensitivity to parameter error, highly nonlinear functions like the eccentricity constraint (19) can render this approach unreliable. A popular approach for effectively modeling the nonlinear transformation of a distribution is the Unscented Transform popularly applied in the design of Kalman filters [7]. Briefly, this method models the effect of a nonlinear function $f(\boldsymbol{x}) : \mathbb{R}^n \to \mathbb{R}^m$ on a multivariate Gaussian random variable $\boldsymbol{x} \sim N(\boldsymbol{\mu}_x, P_x)$ by using a set of representative *sigma points* which capture the major features of the distribution. A common selection for these points is

$$\begin{aligned} \boldsymbol{x}_0 &:= \boldsymbol{\mu}_x, \\ \boldsymbol{x}_i &:= \boldsymbol{\mu}_x + [\sqrt{n P_x}]_i, \qquad && i = 1, \ldots, n \\ \boldsymbol{x}_i &:= \boldsymbol{\mu}_x - [\sqrt{n P_x}]_{i-n}, \qquad && i = n+1, \ldots, 2n \end{aligned} \tag{29}$$

where the operator $[\sqrt{\cdot}]_i$ denotes the i^{th} row of the matrix square root. Specifically, the sigma points chosen in (29) capture the distribution's mean and the principle axes of its 1σ covariance ellipsoid. Mapping these sigma points through f, the sample mean and covariance can be obtained as follows

$$\boldsymbol{\mu}_y \approx \frac{1}{2n}\sum_{i=1}^{2n} f(\boldsymbol{x}_i), \qquad P_y \approx \frac{1}{2n}\sum_{i=1}^{2n} [f(\boldsymbol{x}_i) - \boldsymbol{\mu}_y]\cdot[f(\boldsymbol{x}_i) - \boldsymbol{\mu}_y]^\top. \tag{30}$$

Alternatively, the sample mean $\boldsymbol{\mu}_y$ may be fixed to the transformed mean sigma point $f(\boldsymbol{x}_0)$ if desired. This deceptively simple approach is highly efficient compared to traditional sample-based estimation methods (e.g. Monte-Carlo) and is at least third-order accurate for Gaussian inputs [7]. Notably however, the UT remains a local approximation and may still perform poorly if f is sufficiently nonlinear.

To leverage the appealing benefits of the UT for the CAR boundary estimation problem developed in this work, a mapping f must be found which is compatible with the above definitions. Given the constraint κ, the original definition (22) maps the observation parameter set z to a CAR boundary contour $\partial\mathcal{D}$. Correspondingly, a set of sigma points z_i selected from the parameter distribution would produce a set of *Sigma Contours* $\partial\mathcal{D}_i$ which collectively describe the key nonlinear transformations of the CAR boundary. Intuitively, it then follows that (30) can be applied point-wise along these Sigma Contours to estimate the boundary distribution. More specifically, given a common parameterization τ for the family of Sigma Contours $\boldsymbol{s}_i(\tau)$, the UT statistics (30) can be rewritten as

$$\boldsymbol{\mu}_s(\tau) := \boldsymbol{s}_0(\tau), \qquad P_s(\tau) := \frac{1}{2n}\sum_{i=1}^{2n} [\boldsymbol{s}_i(\tau) - \boldsymbol{\mu}_s(\tau)]\cdot[\boldsymbol{s}_i(\tau) - \boldsymbol{\mu}_s(\tau)]^\top, \tag{31}$$

where the estimated mean $\boldsymbol{\mu}_s(\tau)$ is fixed to the mean solution contour $\boldsymbol{s}_0(\tau)$. This concept is illustrated in Fig. 3b, which depicts a distribution of boundary contours for an ellipsoidal example CAR and the estimated 3σ covariance ellipsoids at the sample times τ_k. This example problem also illustrates two challenges inherent to this approach:

1. τ must be consistently defined (i.e. identify similar points) across all sigma contours $\boldsymbol{s}_i$.
2. The inflation direction $\hat{\boldsymbol{n}}(\tau)$ in (24) must be selected to estimate the 3σ CAR bound from $P_s(\tau)$.

Beginning with (1), recall from section 2.2, that the boundary contours $\partial\mathcal{D}$ are collections of points which possess no inherent ordering. To apply the UT, the parameterization τ must provide this ordering with each fixed τ identifying a consistent location on each contour. For the ellipsoidal example shown in Fig. 3b, the angle of the contour normal $\boldsymbol{s}_{\perp}(\tau)$ is a consistent feature across the contour family which can be used to associate similar points. For example, letting $\tau \in [-1,1]$ in Fig. 3b be clockwise increasing and periodic, the top, rightmost, and bottom edge of each ellipse may be assigned the the times τ = -0.5, 0, and 0.5 respectively, with the leftmost edge assigned the periodic reset time $\tau \in \{-1,1\}$. In cases where the contours are not uniquely point-wise identifiable (e.g. the CAR constraints depicted in Figs. 1 and 2), it is often sufficient to fit a set key feature points like those specified above using a smooth, monotonic increasing regression to consistently define τ along the remainder of the contour. In practice, a suitable parameterization will produce a dominant eigenvalue in the covariance $P_s(\tau)$ at almost all points (see 3b), orienting the CAR's local inflation direction.

Regarding challenge (2) above, note that even in the simplified example shown in Fig. 3b, the covariance ellipsoid $P_s(\tau)$ is not always dominantly skewed along the contour normal $\hat{\boldsymbol{s}}_{\perp}(\tau)$. In some cases, it may prove beneficial to select an inflation direction $\hat{\boldsymbol{n}}(\tau)$ reflecting the covariance estimate (e.g. along the dominant eigenvector) rather than $\hat{\boldsymbol{s}}_{\perp}(\tau)$ as assumed in this work. In practice, such adaptations can correct for error introduced by imperfect parameterizations; specifically those derived for contour families which drastically vary in shape. For simplicity, alternate inflation directions are not discussed in this work, though initial testing provided compelling evidence for their use.

3.4 Monte-Carlo Estimation of CAR Boundary

In order to compare the accuracy of the DI and UT approaches, the true distribution of the CAR boundary contours must also be determined. As in many applications, this is most reliably accomplished by directly and comprehensively sampling the distribution in order to directly measure its Probability Density Function (PDF). As with the unscented approach, a parameterization of the solution contours is ultimately required to estimate the resultant contour distribution as well as to compare them to the DI and UT solutions. Given N parameter samples $\boldsymbol{z}_i$ generated from the multivariate Gaussian distribution $\boldsymbol{z} + \delta\boldsymbol{Z}$, the estimated boundary contour distribution statistics are given by

$$\boldsymbol{\mu}_s(\tau) := \frac{1}{N}\sum_{i=1}^{N}\boldsymbol{s}_i(\tau), \qquad P_s(\tau) := \frac{1}{N-1}\sum_{i=1}^{N}[\boldsymbol{s}_i(\tau)-\boldsymbol{\mu}_s(\tau)]\cdot[\boldsymbol{s}_i(\tau)-\boldsymbol{\mu}_s(\tau)]^{\top}, \tag{32}$$

where, in contrast to (31), the sample mean is estimated to accurately capture the underlying distribution.

3.5 Estimating CAR Boundary Contours

In practice, each of the above algorithms requires one or more parameterized estimates of the CAR boundary contour $\boldsymbol{s}(\tau)$ in order to compute and apply the parameter-error inflation. When comparing the many techniques available in the literature for estimating surface contours, the reader should consider the following two requirements

1. The numeric form of the contours $\boldsymbol{s}(\tau)$ must admit support for interpolation along the entire boundary $\partial\mathcal{D}$.
2. The selected CAR contour solver must be robust to numerically ill-conditioned units and constraints.

Briefly, requirement (1) ensures that any samples of the original boundary contour $\boldsymbol{s}(\tau)$ or samples of the boundary contour distribution $\boldsymbol{s}_i(\tau)$ accurately reflect the CAR boundary $\partial\mathcal{D}$ and do not include error or numerical artifacts added by the interpolation scheme. Additionally, this requirement ensures that the *entirety* of the CAR boundary is captured in the numeric representation of $\boldsymbol{s}(\tau)$, ensuring that any regions of high sensitivity are captured by the chosen inflation algorithm. The second requirement ensures that the CAR boundary is accurately determined to an appropriate tolerance for the selected application units: a challenging prospect given the numerical ill-conditioning introduced by the scale of astronomical units. An practical example of this issue is the numerical artifacts observed in the SMA contours presented by [9] (identified as erroneous against MC simulations). To address these challenges, this work employs canonical unit scaling and numerical integration techniques in determining the CAR boundary, as these techniques (1) provide both relative and absolute tolerance guarantees along the computed solution mesh and (2) provide built-in interpolation schemes which also satisfy those fixed tolerances.

In addition to accurately computing the individual boundary contours, one must also account for CARs composed of multiple isolated regions. For example, the surface map shown in Fig. 2a reveals a CAR composed of three

Table 1: CAR Observation Parameters

	α [°]	δ [°]	$\dot{\alpha}$ [°/s]	$\dot{\delta}$ [°/s]	$\boldsymbol{q}$ [km]	$\dot{\boldsymbol{q}}$ [km/s]
GEO	270.292	-5.417	3.72E-3	1.77E-4	[−1381.427, 5122.805, 3527.9]	[−3.74E-1, −1.01E-1, 0]
LEO	309.091	17.322	-3.84E-2	0	[−1359.374, 5128.701, 3527.9]	[−3.74E-1, −9.91E-2, 0]

Table 2: CAR Observation Uncertainties

	Sample Period	$\sigma_\alpha, \sigma_\delta$ [arc s]	$\sigma_{\dot{\alpha}}, \sigma_{\dot{\delta}}$ [arc s/s]	σ_{q_i} [m]	$\sigma_{\dot{q}_i}$ [m/s]
GEO	1 min	5.7735	0.1179	1	1
LEO	1 sec	5.7735	7.0711	1	1

sub-regions, while that shown in Fig. 2b shows a CAR which, through inflation, will *merge* two of its three sub-regions. While large portions of these sub-regions may be removed when combined with other constraints, the regions which remain are equally likely under the uniform distribution assumed by the CAR approach. Thankfully, this complication is simple to overcome in practice: each of these regions may be inflated independently using any of the above approaches, and the overall CAR for that constraint is obtained as the union of the sub-regions.

4. NUMERICAL EVALUATION

This section compares the efficacy of the Differential, Unscented, and Monte-Carlo approaches for estimating the CAR region for representative examples scenarios in GEO and LEO respectively.

4.1 Experimental Design

To comprehensively evaluate the accuracy of the DI and UT CAR parameter error estimation schemes, this section examines the EO CAR inflation for two representative example scenarios: a compressed observation for a GEO object taken with a 1-min sample period, and a compressed observation for a LEO object taken with a 1-sec sample period. The specific observation, station parameters, and (compressed) observation uncertainties are shown in tables 1 and 2 respectively. These scenarios are identical to those presented in [9] (absent errors in sampling time), but are compressed from 3 angles-only measurement arcs to single angle/angle-rate measurements. For each scenario, the 3σ DI and UT CAR inflation boundaries are then compared against an $N = 10,000$ case Monte-Carlo simulation. For simplicity, comparisons of the eccentricity CARs only include the largest CAR sub-region (though the merger of two of these regions is discussed in the LEO example). The exact computational expense of each method is not strictly compared in this study, though the approximate run time of the DI, UT, and MC methods was observed to scale closely with the required number of contour solutions for each approach (i.e. 1, 21, and 10,000 respectively).

4.2 GEO Scenario

We first examine the inflations of the GEO scenario CARs depicted in Fig. 4 with the SMA bounds $a \in [4 \cdot 10^4, \infty)$ km and eccentricity bounds $e \in [0, 0.08]$. Looking first at Fig. 4a, we note that the DI, UT, and MC inflation regions reasonably agree for both the SMA and ECC constraints, widening the resulting CAR in manner consistent with the surface maps shown in Figs. 1a and 2a. Comparing the individual inflation boundaries more closely using the contour distributions shown in Fig. 4b, we see that the UT approach is, on average, around 10x closer to the MC boundary than the DI boundary. In practice, this accuracy gain is likely insignificant as the majority of the DI boundary differs from the MC boundary by less than 0.25%. The large apparent errors in the eccentricity CAR near $\tau = -0.5$ and 0.5 correspond to the top and bottom of the CAR in 4a and are artifacts of normalizing that region's near-zero inflation.

Combining these results, we may conclude that (in this scenario) the UT approach provides superior accuracy to the DI approach at the cost of additional computational expense. More critically, note that the combined CAR composed by intersecting each of these constraints nearly *doubles* in size under the inclusion of observation parameter error. This emphasizes the critical importance of including this error in practical applications of the CAR approach.

4.3 LEO Scenario

We next examine the inflations of the LEO scenario CARs depicted in Fig. 5 with the SMA bounds $a \in [7 \cdot 10^3, \infty)$ km and eccentricity bounds $e \in [0, 0.2]$. As with the GEO example, we first note from Fig. 5a that the SMA constraint

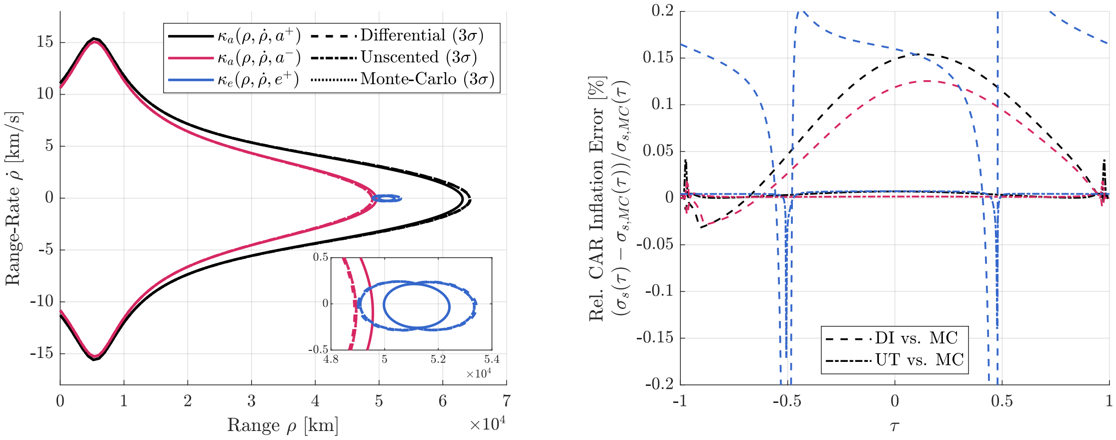

Fig. 4: Comparison of CAR 3σ inflation results in GEO scenario using different approaches. Panel (a) depicts the entire region, while (b) displays relative inflation error along the contour normal w.r.t. the Monte-Carlo results.

inflation boundary is largely identical between the DI and UT methods. Looking at Fig. 5b, the DI and UT approaches slightly under/overestimate the inflated region respectively, though the deviations of each are $< 0.25\%$ with respect to the MC estimate.

Unlike in the GEO scenario however, each approach's estimate of the eccentricity constraint boundary is far more nuanced. As shown by the colormap in Fig. 2b, the eccentricity CAR inflation shown in 5 captures the merging of two CAR sub-regions. As a result, the distributions modeled by the UT and MC approaches becomes heavily bimodal, with the solution contours split above and below the saddle point between the two sub-regions. This large variation in the Gaussianity of the contour distribution is reflected in the variation between the solutions, with high agreement in the (Gaussian) lower right lobe, and high variation in the (bimodal) upper left. Naturally, as all of the above approaches (including the MC) assume Gaussianity in the contour distributions, their relative accuracy in this example is harder to measure. From a practical viewpoint, the combination of the low gradients near saddle points and the bimodal contour distribution produced by merging CAR sub-regions should render the DI approach more conservative and the UT approach more comprehensive. This interpretation is reflected in the results shown in Fig. 5a with Fig. 2b, where the UT approach appears to better track the behavior of the constraint surface than the DI approach and the MC solution appears to drastically over estimate the inflated CAR.

Regardless of the selected inflation algorithm, this example also reveals a composite CAR that is *substantially* (3-10×) larger than the original CAR when considering reasonable parameter errors. This further emphasizes the importance of including parameter errors when implementing the CAR approach.

5. CONCLUSION

In this work, we examined alternative strategies for incorporating parameter uncertainty in the Constrained Admissible Region for short-arc EO observations. To this end, we first reproduced and expanded upon the Differential Inflation approach developed by [9] to include both prevalent CAR constraint bounds on orbital eccentricity and semi-major axis. This included the introduction of canonical scaling and numerical integration techniques in defining the constrain boundary, reduce intrinsic numerical ill-conditioning and addressing key shortcomings reported by [9]. To examine the higher order moments of the parameter-uncertainty inflated CAR, we also developed an adaptation the Unscented Transform (UT) to measure nonlinear distortions in the CAR boundary through careful parameterization of the solution contours. The accuracy and numerical efficiency of the above approaches were then compared against Monte-Carlo simulations in both the LEO and GEO orbital regimes. This comparison yields the following key conclusions:

1. The size and shape of the CAR is *massively* expanded by the inclusion of standard EO observation error in both the GEO and LEO regimes (3-10×),

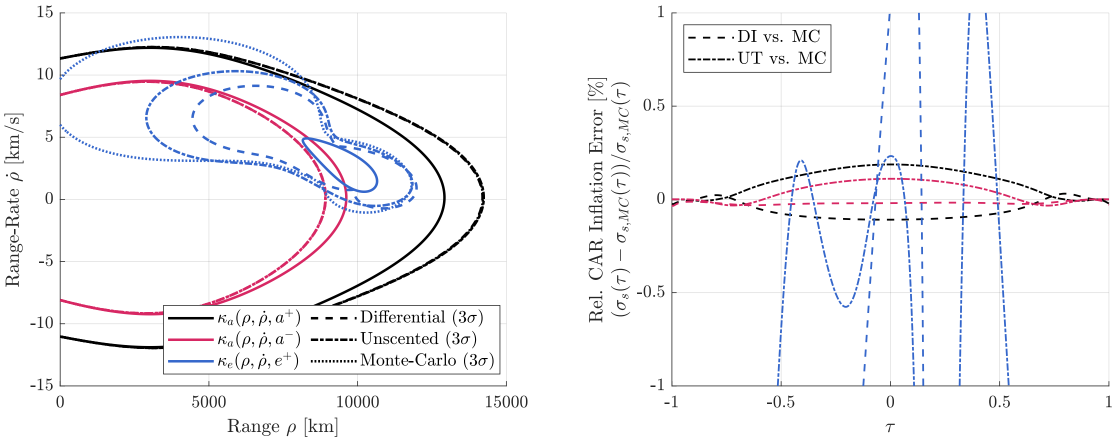

Fig. 5: Comparison of CAR 3σ inflation results in LEO scenario using different approaches. Panel (a) depicts the entire region, while (b) displays relative inflation error along the contour normal w.r.t. the Monte-Carlo results.

2. In most cases, both the DI and UT approaches accurately model the CARs inflation ($\leq 0.25\%$ relative error),
3. Both approaches become unreliable near saddle points in the constraint surfaces (where the CAR boundary is more bimodal than Gaussian. In this case, the UT method better estimates the CAR's (discontinuous) expansion.

These results highlight the importance of parameter error in admissible region generation and present multiple approaches to combine the respective robustness and computational efficiency benefits of the Probabilistic Admissible Region and CAR approaches in practical applications.

6. REFERENCES

[1] Kyle J. Demars and Moriba K. Jah. Probabilistic initial orbit determination using gaussian mixture models. *Journal of Guidance, Control, and Dynamics*, 36:1324–1335, 2013.

[2] Kyle J. Demars, Moriba K. Jah, and Paul W. Schumacher. Initial orbit determination using short-arc angle and angle rate data. *IEEE Transactions on Aerospace and Electronic Systems*, 48, 2012.

[3] Islam I. Hussein, Christopher W.T. Roscoe, Michael Mercurio, Matthew P. Wilkins, and Paul W. Schumacher. Probabilistic admissible region for multihypothesis filter initialization. *Journal of Guidance, Control, and Dynamics*, 41:710–724, 2018.

[4] Andrea Milani, Giovanni F Gronchi, Mattia De' Michieli Vitturi, Zoran KneÅ, and Evic Evic2. Orbit determination with very short arcs. i admissible regions. *Celestial Mechanics and Dynamical Astronomy*, 90:59–87, 2004.

[5] Jason Stauch, Jason Baldwin, Moriba Jah, Tom Kelecy, and Keric Hill. Mutual application of joint probabilistic data association, filtering, and smoothing techniques for robust multiple space object tracking. In *AIAA/AAS Astrodynamics Specialist Conference 2014*. American Institute of Aeronautics and Astronautics Inc., 2014.

[6] David A Vallado and Jacob D Griesbach. Simulating space surveillance networks. In *AAS/AIAA Astrodynamics Specialist Conference*. American Institute of Aeronautics and Astronautics Inc., 2011.

[7] Eric A Wan and Rudolph Van Der Merwe. The unscented kalman filter for nonlinear estimation. *Proceedings of the IEEE 2000 Adaptive Systems for Signal Processing, Communications, and Control Symposium*, 2000.

[8] R. M. Weisman, M. Majji, and K. T. Alfriend. Analytic characterization of measurement uncertainty and initial orbit determination on orbital element representations. *Celestial Mechanics and Dynamical Astronomy*, 118:165–195, 2 2014.

[9] Johnny L. Worthy and Marcus J. Holzinger. Incorporating uncertainty in admissible regions for uncorrelated detections. *Journal of Guidance, Control, and Dynamics*, 38:1673–1689, 2015.