The Epistemic Support-Point Filter (ESPF): Possibilistic Space Object State Estimation Under Ignorance

Moriba Kemessia Jah
Jah Decision Intelligence Group,
Aerospace Engineering & Engineering Mechanics Department,
The University of Texas at Austin, Austin, TX, USA
Van Haslett
Jah Decision Intelligence Group,
The University of Texas at Austin, Austin, TX, USA

ABSTRACT

A probabilistic filter can become increasingly confident while becoming increasingly wrong. When model assumptions are violated—by an unmodeled maneuver, an undisclosed sensor bias, or dynamics that are bounded rather than stochastic—the covariance continues to shrink. The filter interprets incoming observations as confirmation of its current trajectory because that is what its algebra demands. It has no mechanism for representing contradiction; normalization ensures that evidence always increases the weight of some hypothesis rather than eliminating hypotheses outright. The confidence it reports is therefore an artifact of its assumptions, not a property of the data. It fails silently.

We introduce the Epistemic Support-Point Filter (ESPF), a non-Bayesian recursive estimator built on a different principle: confidence should emerge only from survival under contradiction, not from the algebra of normalization. The ESPF maintains a finite set of explicit hypotheses spanning an admissible region of state space. Evidence acts by falsifying hypotheses whose predicted observations are geometrically incompatible with what was observed. Crucially, the ESPF admits the existence of mutually incompatible evidence without forcing reconciliation: a falsified hypothesis is eliminated, and nothing is renormalized to compensate. What survives is what has not yet been contradicted. Compatibility encodes non-contradiction—not belief, not frequency, not calibration—and therefore cannot be reinterpreted probabilistically. When the model is violated, the ESPF fails *informatively*: surprisal spikes, necessity surges, and the Epistemic Width Monitor signals that the filter is operating outside its warrant—diagnostics with no probabilistic analog.

Keywords: Possibility theory, epistemic uncertainty, non-Bayesian filtering, tropical Hamilton–Jacobi, minimum-volume ellipsoid, Boltzmann ignorance, impossibility field, possibilistic Cramér–Rao bound, minimax medoid, space domain awareness, epistemic width, maneuver detection, Kalman filter generalization

1. INTRODUCTION

The Kalman filter and its nonlinear extensions [27, 33] are among the most successful algorithms in engineering, and their success is deserved under the conditions they assume. But their confidence is structural: it is produced by the algebra of Gaussian conditioning, not earned from the data. When those conditions hold, structural confidence and evidential confidence agree. When they fail—when a satellite maneuvers without warning, when a sensor carries an undisclosed bias, when uncertainty is imposed by detection geometry rather than characterized by a noise distribution—the covariance continues to contract. The filter interprets every new observation as confirmation of its current trajectory because normalization requires it: the posterior must integrate to one, so evidence can only redistribute weight, never eliminate hypotheses outright. The filter grows more confident precisely as its model becomes wrong. It does not know that anything has changed. It fails silently—and the analyst trusting its reported uncertainty has no way to know either.

This paper introduces an estimator designed to fail informatively instead.

Copyright © 2026 Advanced Maui Optical and Space Surveillance Technologies Conference (AMOS) – www.amostech.com

The core idea

The Epistemic Support-Point Filter (ESPF) is built on a different premise: *confidence should emerge only from survival under contradiction*.

The ESPF represents uncertainty as a finite *admissible set* of explicit hypotheses about the system state—states that evidence has not yet ruled out. Measurement updates do not reweight hypotheses by likelihood. They falsify hypotheses whose predicted observations are geometrically incompatible with what was actually observed. What survives is what remains admissible. What is expelled is what has been contradicted.

This is structurally different from probabilistic filtering, not merely philosophically different. Probabilistic filters such as the Unscented Kalman Filter (UKF) and the non-gaussian Particle Filter (PF) cannot represent epistemic contradiction because normalization is mandatory: every observation must increase the weight of some hypothesis, and the posterior must integrate to one—and therefore the UKF and PF cannot eliminate hypotheses without redistributing their mass. Contradiction is absorbed into the residual and the covariance; it is never expressed as such. The ESPF admits the existence of mutually incompatible evidence without forcing reconciliation. A falsified hypothesis is eliminated from the admissible set entirely, and nothing is renormalized to compensate. Contradiction is a first-class object, expressed as surprisal and handled by elimination—not smoothed away as residual variance.

When a model violation occurs, the ESPF's surprisal field registers a spike orders of magnitude above the nominal baseline. Necessity surges toward unity as the admissible set mass-falsifies. The Epistemic Width Monitor signals that the filter is operating outside its possibilistic warrant. None of these signals require a maneuver detector, a process-noise inflation heuristic, or manual retuning. They emerge from the geometry of the update itself.

Theoretical grounding

The ESPF is the operational closure of a three-paper theoretical framework grounded in the Theory of Epistemic Abductive Geometry (TEAG):

- **Paper I—TEAG foundations** [21]: establishes the TEAG quintuple $\mathscr{E} = (\mathscr{H}, \pi, \{\mathscr{H}_\alpha\}, \mathscr{C}, \mathscr{A})$, the PCRB as a universal floor on admissible entropy reduction per measurement, the Hölder-mean continuum connecting possibilistic and probabilistic inference, and the formal proof that the UKF is the Gaussian limit of the ESPF via Choquet-to-Lebesgue convergence [24].
- **Paper II—Tropical dynamics** [22]: proves that the ESPF predict–update recursion satisfies a tropical Hamilton–Jacobi equation in the max-plus semiring. Two scalar fields live on hypothesis space: the *impossibility field* $\Phi_\emptyset = -\log \pi$ encoding accumulated epistemic history, and the *surprisal field* $\Phi_S = \frac{1}{2}\|L_e^{-1}(y - g(h))\|^2$ encoding tension with the current observation. The posterior update $\tilde{\Phi}_\emptyset = \max(\Phi_\emptyset, \Phi_S)$ is their pointwise tropical superposition—an algebraic identity forced by the TEAG axioms, not a modeling choice. Falsification is wavefront propagation; the PCRB is the minimum action per update cycle.
- **Paper III—The ESPF (this paper)**: introduces the algorithm, documents the implementation, and demonstrates operational behavior across three orbital estimation scenarios. All theoretical guarantees are established in Papers I and II; this paper shows what the theory looks like when it runs.

In addition to Papers I–II, the series includes a companion optimality paper establishing the Possibilistic Cramér–Rao Bound (PCRB), the minimum-specificity and minimum-ignorance theorems, and the formal stability analysis of the recursion [19], and the original ESPF framework paper [20]. All companion papers are publicly available as preprints at the time of submission, so every proof invoked here can be independently examined. To keep the reference apparatus proportionate, each companion is cited formally at first use and referred to thereafter by name (Paper I, Paper II, the companion optimality paper).

Three modes of epistemic relationship with evidence

The difference between the ESPF and probabilistic filters is most clearly seen by contrasting three epistemic relationships an estimator can have with evidence:

Copyright © 2026 Advanced Maui Optical and Space Surveillance Technologies Conference (AMOS) – www.amostech.com

	Probability	Compatibility	Contradiction
Object	Degree of belief	Degree of non-rejection	Falsification event
Update	Bayes' rule (multiply, normalize)	Min-intersection (conjunctive)	Elimination
Under conflict	Redistributes weight	Contracts support	Expels hypotheses
What it encodes	Relative frequency / belief	Non-contradiction	Irreconcilable mismatch
UKF representation	Yes	Implicit	Never
ESPF representation	No	Yes	Yes

Probabilistic filters live entirely in the first column. The ESPF operates in the second and third. This is not a claim that probability is wrong. It is a claim that probability is insufficient when evidence contradicts the model—and that an estimator which cannot represent contradiction will systematically misreport its own reliability. The formal criterion characterizing precisely when probabilistic closure is a legitimate representation of uncertainty—and when it is not—is the Epistemic Validity Condition (EVC) developed in [23]; the present paper is the operational instantiation of that criterion for recursive state estimation.

The key empirical contrast

The UKF, PF, and the ESPF are compared across three scenarios: LEO nominal, LEO stressed (unknown initial error + sensor bias + unmodeled maneuver), and a real-data GEO validation against WAAS GPS truth. Under nominal conditions the two filters achieve equivalent accuracy. Under stress, they diverge in kind, not just degree.

UKF and PF under stress	ESPF under stress
Covariance contracts as data arrive	Surprisal spikes at contradiction
Confidence increases as model drifts	Necessity surges as support falsifies
No signal that anything is wrong	EWM alarm precedes residual spike
Recovers by absorbing error silently	Recovers by repositioning on evidence
Fails silently	Fails informatively

Confidence is not assigned. It is earned through non-falsification.

This is not a claim that the ESPF is always more accurate. It is a claim that the ESPF is always more *honest*—and that in a well-designed epistemic system, honesty and accuracy are not in tension. Confidence that emerges from survival under contradiction is trustworthy. Confidence that emerges from normalization is structural, and can persist long after the structure has ceased to describe reality.

Paper organization

Section 2 reviews prior work and positions the ESPF within the non-Bayesian estimation literature. Section 3 formalizes the possibilistic state estimation problem. Section 4 describes support-point generation via Smolyak sparse grids. Section 5 compares the possibilistic support-point structure to probabilistic sigma-points. Section 6 derives the prediction step. Section 7 presents the measurement update, including the tropical impossibility/surprisal field structure and the minimum-information-entropy enclosure (Section 7.2.2)—the free-center minimum-volume ellipsoid that the ESPF is optimal at recovering, and the scale behavior of the innovation geometry (Section 7.2.3), where the joint-warrant form used in all reported results is distinguished from a split-metric variant proposed for the journal version. Section 9 documents implementation parameters. Sections 10–11 present the three experimental scenarios and results. Section 12 addresses limitations and future work, and Section 13 concludes.

Copyright © 2026 Advanced Maui Optical and Space Surveillance Technologies Conference (AMOS) – www.amostech.com

2. PRIOR WORK AND POSITIONING

Rudolf Kalman cautioned that estimation frameworks risk encoding the prejudices of their designers rather than the structure of the data:

> "Whenever a model is built, it is always proper to ponder the basic scientific question: *Is the model really based on the data? Or is it an artifact displaying the prejudices of its creator?*" [28].

Within the probabilistic family, richer representations than the single Gaussian have long been pursued—Gaussian-sum filters [2] and cubature rules [3] among them—but all retain additive normalization and therefore inherit the structural confidence critiqued above. Non-Bayesian alternatives have been explored along several lines. Possibility theory [34, 16, 18] replaces additive probability with maxitive, ordinal plausibility, and has been developed extensively for information fusion [17]. Kalman-like filtering has also been formulated directly in a possibilistic setting [4], whose prediction step—a sup–min propagation of a possibility distribution through a prediction function or transition relation, $\pi_{t+1}(\omega') = \max_{\omega} \min(\pi_t(\omega), \pi_\omega(\omega'))$—anticipates the support-expansion operator of Eq. (4). That work, however, remained a preliminary theoretical draft: no implementation, no numerical experiments, and a conditioning rule (crisp-event possibilistic conditioning, which assigns possibility one only to the most-compatible states within an accepted proposition A and zero to every state outside it) distinct in kind from the graded conjunctive update of Eq. (17). Both frameworks renormalize—the ESPF max-rescales the posterior possibility field after every update (Section 7.5)—but the two operations differ structurally: BDP's renormalization is conditioning-specific, defined relative to a crisp accepted event A and collapsing everything outside it to zero; the ESPF's max-rescaling is a uniform, ordinal-preserving reparameterization of the entire surviving field, with no crisp event and no discontinuous collapse to zero outside an accepted set. The ESPF closes the remaining gap between the two frameworks: the same class of possibilistic filtering operator is instantiated as a tested recursive estimator and validated against real GPS ground truth in the GEO scenario (Section 11), not merely simulated data. Set-membership, interval, and zonotopic methods [29, 25, 1, 10] maintain bounded admissible sets under deterministic error models. Evidence-theoretic [31] and credal or imprecise-probability approaches [11, 30] generalize Bayesian conditioning to sets of distributions. Possibilistic multi-object filters and outer-probability-measure frameworks [12, 7, 13] have been applied specifically to space object tracking, including possibilistic tracking under epistemic uncertainty [8], consensus credibility particle filters [6], and sensor tasking for geosynchronous object catalog maintenance [5]. Multi-fidelity orbit determination under systematic errors provides further context for the stress scenarios examined here [35]. Possibilistic information fusion via maximal coherent subsets [15] and non-additive uncertainty modelling [14] round out the landscape of set-valued and imprecise approaches.

The ESPF departs from all of these by combining possibilistic ordinal inference, sparse-grid geometry, metric-consistent falsification, and the EWM diagnostic suite into a single recursive architecture that admits no probability distributions at any stage, and whose update rule is proved algebraically necessary by the TEAG axioms. Specifically, the update $\tilde{\Phi}_\emptyset = \max(\Phi_\emptyset, \Phi_S)$ is not a design choice: Proposition 5.2 of Paper II proves it is the unique update consistent with Popperian contraction and the TEAG evidence-referencing axiom. No existing possibilistic filter provides this algebraic necessity guarantee alongside real-time, ground-truth-free epistemic diagnostics. The operational biography motivating this departure from probabilistic closure—two decades of spacecraft navigation, debris tracking, and orbit determination in which standard filters failed systematically under epistemic rather than statistical uncertainty—is documented in the EVC companion paper, which also proves Bayesian inference as the earned geometric limit of TEAG when the EVC is satisfied.

Key Contributions

- A possibilistic recursive estimator based on finite admissible support points, requiring no noise model, prior covariance, or statistical calibration.
- A measurement update whose tropical structure $\tilde{\Phi}_\emptyset = \max(\Phi_\emptyset, \Phi_S)$ is an algebraic identity forced by Popperian contraction, not a design choice (Paper II).
- A *minimum-information-entropy enclosure* for the predicted measurement geometry: the free-center minimum-volume enclosing ellipsoid (MVEE) of the predicted measurement cloud is the representation of least Boltzmann ignorance over admissible measurements, and the ESPF is proved optimal at recovering it—returning

Copyright © 2026 Advanced Maui Optical and Space Surveillance Technologies Conference (AMOS) – www.amostech.com

the least-specific posterior consistent with the evidence and its minimum-volume honest enclosure (companion optimality paper). The MVEE's own active (support) set supplies the non-redundant evidence intrinsically, so no coordinatewise (Pareto/tropical-vertex) pre-filtering of the cloud is necessary or beneficial (Section 7.2.2).

- A *committed estimate that is always an actual hypothesis*: the ESPF's point estimate is the whitened minimax medoid of the support—a member of the admissible set by construction, never a synthesized mean—and is computed by the same rule at initialization, prediction, and measurement update. It is worst-case optimal over the support, within a factor of two of the Chebyshev radius (companion optimality paper), and, unlike a sigma-point mean, is guaranteed to be a dynamically realizable, unfalsified state (Proposition 1).
- The Epistemic Width Monitor (EWM): real-time, ground-truth-free diagnostics (W_{ep}, I_c, α_c, δ_k) that certify whether the filter is within its possibilistic warrant.
- Demonstration on three orbital scenarios—LEO nominal, LEO stress, real-data WAAS GEO—showing that the ESPF matches UKF accuracy under nominal conditions and provides explicit epistemic alarms under stress that no covariance-based filter can produce.

3. PROBLEM FORMULATION

Let the system state at time k be $\mathbf{x}_k \in \mathbb{R}^n$, evolving under a nonlinear process model $\mathbf{x}_k = f(\mathbf{x}_{k-1}) + \mathbf{w}_{k-1}$, where $\mathbf{w}_{k-1}$ is a bounded epistemic perturbation with no assumed distribution. Observations $\mathbf{y}_k \in \mathbb{R}^m$ are generated through $\mathbf{y}_k = h(\mathbf{x}_k) + \mathbf{v}_k$, where $\mathbf{v}_k$ is bounded epistemic measurement noise.

In the possibilistic setting, epistemic uncertainty over $\mathbf{x}_k$ is represented by a possibility distribution $\pi_{\mathbf{x}_k} : \mathbb{R}^n \to [0,1]$, where $\pi(x) = 1$ means x is fully admissible and $\pi(x) = 0$ means x is ruled out. The scale is ordinal: $\pi(x_1) > \pi(x_2)$ indicates x_1 is more plausible, without quantifying by how much. The admissible support at time k is $\mathscr{S}_k = \{x \in \mathbb{R}^n \mid \pi_{\mathbf{x}_k}(x) > 0\}$.

The ESPF approximates $\pi_{\mathbf{x}_k}$ by a finite set of *support points* $\{\chi_k^{(i)}\}_{i=1}^M \subset \mathscr{S}_k$ with associated possibility values $\{\pi_k^{(i)}\} \subset (0,1]$. These are not samples from a distribution; they are equally-spaced deterministic representatives of the admissible region, placed via Smolyak sparse-grid construction [32]. The objective at each step is to update $\mathscr{S}_k$ so that it contains all epistemically admissible states conditioned on observations $\mathscr{Z}_k = \{y_1, \ldots, y_k\}$, while discarding regions rendered implausible by possibilistic incompatibility.

The core operators of possibility theory used throughout are: joint possibility $\pi_{X,Y}(x,y) = \min(\pi_X(x), \pi_Y(y))$; marginalization via supremum $\pi_X(x) = \sup_y \pi_{X,Y}(x,y)$; and Zadeh's conditioning rule $\pi_{X|Y}(x|y) \propto \min(\pi_X(x), \pi_{Y|X}(y|x))$. These replace integration and Bayes' rule throughout the filter. A possibility measure is maxitive: $\Pi(A \cup B) = \max(\Pi(A), \Pi(B))$, and the Choquet integral of an event indicator under this capacity reduces to the possibility of the event, $\int \mathbf{1}_A \, d\Pi = \max_{h \in A} \pi(h)$: the measure of any event is carried entirely by its most plausible unfalsified hypothesis [9].

4. SUPPORT-POINT GENERATION

Because the ESPF possibility distribution is uniform over a compact convex support, there is no privileged mean or mode. Support points are placed deterministically to span the admissible region rather than to match statistical moments.

Sparse-grid construction. Given support $\mathscr{S}_k = [x^-, x^+]^n$ (or a general convex body), support points are drawn from a Smolyak sparse grid at level ℓ:

$$\mathscr{A}(n,\ell) = \bigcup_{\mathbf{i} \in \mathscr{I}_{n,\ell}} (\mathscr{X}_{i_1} \otimes \cdots \otimes \mathscr{X}_{i_n}), \tag{1}$$

where $\mathscr{X}_i$ denotes the i-th level Clenshaw–Curtis node set in $[-1,1]$, and each normalized node $\hat{\xi}^{(j)} \in [-1,1]^n$ is mapped affinely to $\mathscr{S}_k$:

$$\chi_k^{(j)} = \tfrac{1}{2}(x^+ - x^-) \circ \hat{\xi}^{(j)} + \tfrac{1}{2}(x^+ + x^-). \tag{2}$$

Copyright © 2026 Advanced Maui Optical and Space Surveillance Technologies Conference (AMOS) – www.amostech.com

At level $\ell = 3$ for state dimension $n = 7$, this yields $M = 106$ support points—enough for principled evidence-gradient recovery under stress while remaining computationally tractable. Each point is an equally plausible epistemic hypothesis; no weights or covariances are attached. Smolyak faithfulness under evidence-driven pruning—the guarantee that the grid maintains volumetric coverage as hypotheses are falsified—is proved in the companion optimality paper.

5. SUPPORT-POINT VS. SIGMA-POINT GEOMETRIC COMPARISON

The most direct way to compare support points against sigma-points is to examine how their uncertainty representations relate to the true nonlinear uncertainty structure, as characterized by a Monte Carlo ensemble propagated from the same initial conditions. Figures 1a and 1b show the truth-centered Radial/Transverse/Normal (RTN) uncertainty comparison at the initial and final epochs of the LEO nominal arc. Each panel overlays: the Monte Carlo cloud (grey dots, 1000 samples), the ESPF support points (orange circles) and their MVEE (green ellipse), the UKF sigma points (gold crosses) and their UT 3σ ellipse (purple), the ESPF minimax medoid anchor (cyan dot), and the true position (red dot).

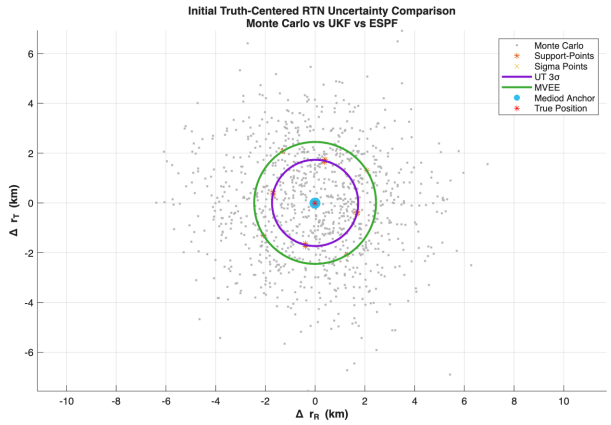

(a) Initial truth-centered RTN uncertainty comparison. Monte Carlo ensemble (grey), ESPF support points and MVEE (orange/green), UKF sigma points and UT 3σ ellipse (gold/purple), medoid anchor (cyan), and true position (red). At initialization, both filters are geometrically consistent with the Monte Carlo distribution; support points and sigma points span the cloud comparably, and the medoid coincides with truth.

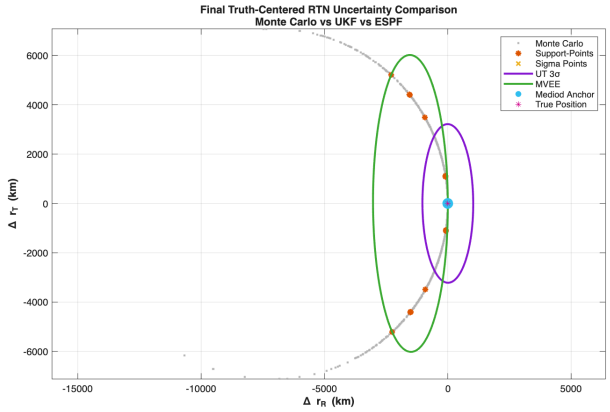

(b) Final truth-centered RTN uncertainty comparison. Monte Carlo ensemble (grey), ESPF support points and MVEE (orange/green), UKF sigma points and UT 3σ ellipse (gold/purple), medoid anchor (cyan), and true position (red). By arc end, the Monte Carlo cloud has deformed into a banana-shaped filament extending $>$14,000 km in-track; the ESPF MVEE elongates consistently with this structure while the UKF ellipse remains compact and symmetric, capturing only a tail of the true distribution.

Fig. 1: Comparison of RTN uncertainty distributions at initial (left) and final (right) epochs of the LEO nominal arc.

The contrast between the two epochs is stark and theoretically significant. At initialization (Figure 1a), both the UKF and the ESPF are consistent with the Monte Carlo cloud in the compact radial/transverse plane: the UT ellipse, the MVEE, and the sample cloud all occupy the same ± 3 km region. This is the regime in which Gaussian linearization is valid, and both filters are geometrically faithful.

By the end of the arc (Figure 1b), the Monte Carlo cloud has undergone classical nonlinear deformation: it stretches into a long, curved filament along the in-track direction, extending more than 14,000 km in radial error relative to a truth-centered frame, with asymmetric concentration near the true position. This banana-shaped distribution is a direct consequence of the sensitivity of LEO orbital period to semimajor axis under nonlinear Keplerian dynamics—a well-known feature of initial orbit uncertainty propagation that Gaussian approximations cannot represent.

The UKF response to this nonlinear growth is geometric failure. Its UT 3σ ellipse (purple) is compact, symmetric, and upright—the shape of a Gaussian approximation—centered at a position that captures only a small tail of the actual Monte Carlo distribution. The UKF has *not* tracked the truth; it has tracked the linearized center of mass of the ensemble while its covariance has remained blind to the asymmetric probability mass that has migrated to the banana's

Copyright © 2026 Advanced Maui Optical and Space Surveillance Technologies Conference (AMOS) – www.amostech.com

curved arm. From the UKF's internal representation, the true position appears to be within the 3σ ellipse, which it is; but the ellipse systematically misrepresents the actual shape of uncertainty.

The ESPF behaves differently. Its MVEE (green) is elongated and tilted in a manner consistent with the actual Monte Carlo distribution, capturing the dominant axis of uncertainty growth. The support points are distributed along this axis, and the minimax medoid (cyan) is placed at the true position (red)—a direct consequence of anchoring to the geometric center of the surviving admissible set rather than to a probabilistic mean. The ESPF has not collapsed its representation to a Gaussian ellipsoid; it has maintained a support geometry that reflects the actual nonlinear structure of the uncertainty.

This comparison crystallizes the core epistemic distinction between the two approaches. The UKF constructs a confident, compact, probabilistically normalized belief that is geometrically inconsistent with the true uncertainty. The ESPF maintains an admissibility envelope that, while broader, correctly orients along the dominant axis of nonlinear uncertainty growth and places its state commitment at the geometrically supported truth. Precision is not the same as accuracy, and covariance contraction is not the same as knowledge.

6. PREDICTION STEP

Nonlinear propagation. Each support point is propagated deterministically through the process model:

$$\hat{\chi}^{(i)}_{k|k-1} = f\left(\chi^{(i)}_{k-1}\right). \tag{3}$$

This is the possibilistic analog of sigma-point propagation but with no probabilistic interpretation: each $\hat{\chi}^{(i)}$ is an explicit hypothesis about where the system *could* be, not a weighted sample.

Support expansion. Process uncertainty is incorporated via sup–min convolution with the bounded process noise support $\mathscr{W}_{k-1}$:

$$\pi_{\mathbf{x}_k}(x) = \sup_i \min\left(\pi_{\mathbf{x}_{k-1}}(\chi^{(i)}_{k-1}),\ \pi_w(x - \hat{\chi}^{(i)}_{k|k-1})\right), \tag{4}$$

which reduces to a Minkowski sum $\mathscr{S}_{k|k-1} = \mathscr{F}_k \oplus \mathscr{W}_{k-1}$ under uniform possibility. In practice the predicted support is approximated by bounding hyperrectangles around the union of shifted noise supports and resampling a new Smolyak grid.

Geometry-preserving compression. After expansion, support cardinality is restored to M by resampling from the predicted bounding geometry. The MVEE of the predicted cloud provides the shape matrix for regeneration, and the coverage-driven σ_k controller (Section 9) sets the scale, preventing both collapse and unbounded growth.

Predicted commitment. The predicted state commitment is the whitened minimax medoid of this predicted support (Section 7.5, Eq. (19))—the possibilistic analog of the UKF predicted mean, but an actual propagated hypothesis rather than a weighted average of propagated points.

7. MEASUREMENT UPDATE

Upon receipt of an observation $\mathbf{y}_k \in \mathbb{R}^m$, the predicted epistemic support must be confronted with evidence and conditioned to retain only those hypotheses that remain admissible. In the ESPF, conditioning is implemented through finite-dimensional geometry and conjunctive (max–min) operations. No likelihoods, probability densities, or stochastic noise models are introduced.

At time step k, the measurement update consists of: (i) deterministic mapping of predicted support into measurement space; (ii) construction of an innovation geometry bounding predicted measurement spread and sensor imprecision; (iii) computation of geometric compatibilities and posterior possibilities; (iv) selection of a core survivor set using a median-surprisal rule with reference tracking; (v) estimation of survivor geometry in state space; (vi) minimax-medoid

Copyright © 2026 Advanced Maui Optical and Space Surveillance Technologies Conference (AMOS) – www.amostech.com

anchor commitment; (vii) deterministic regeneration of a fixed-cardinality support; and (viii) max–min transport of possibility to the regenerated support.

Each step is geometric, finite, and epistemically monotone: evidence may remove hypotheses but never resurrect falsified ones.

7.1 Measurement Mapping and Residuals

Let $\chi_{k|k-1}^{(i)} \in \mathbb{R}^n$ denote the i-th predicted support-point, with associated prior possibility $\pi_{k|k-1}^{(i)} \in [0,1]$, $i = 1,\dots,M$. The predicted measurement corresponding to this hypothesis is

$$\gamma_k^{(i)} = h\left(\chi_{k|k-1}^{(i)}\right), \tag{5}$$

where $h : \mathbb{R}^n \to \mathbb{R}^m$ is the deterministic measurement model. Define the innovation residual

$$\mathbf{e}_k^{(i)} = \mathbf{y}_k - \gamma_k^{(i)}. \tag{6}$$

Residuals are interpreted strictly as geometric discrepancies. They are not random variables and are not assumed to follow any statistical distribution. Their role is to test compatibility between hypotheses and the observed evidence under bounded epistemic uncertainty.

7.2 Innovation Geometry via Ellipsoidal Outer Bounding

Let $\Gamma_k = \{\gamma_k^{(i)}\}_{i=1}^M \subset \mathbb{R}^m$ denote the predicted measurement support. This set represents all measurements that remain admissible prior to evidence.

7.2.1 Minimum-Volume Enclosure as Least Boltzmann Ignorance

To construct an intrinsic innovation geometry, the ESPF computes the free-center minimum-volume enclosing ellipsoid (MVEE) of Γ_k—the Löwner–John ellipsoid, in which the center and shape are optimized jointly—yielding a positive-definite matrix $A_v \succ 0$ and a center $\mathbf{c}_v \in \mathbb{R}^m$ such that

$$(\mathbf{z} - \mathbf{c}_v)^\top A_v (\mathbf{z} - \mathbf{c}_v) \le 1 \quad \forall \mathbf{z} \in \Gamma_k. \tag{7}$$

This is the *minimum-information-entropy enclosure*. Under the Boltzmann reading of entropy $S = \log W$, with admissible measurement microstates counted at a fixed resolution cell v_{cell}, the ignorance carried by an ellipsoidal enclosure E is $S_{\text{B}}(E) = \log\big(\text{Vol}(E)/v_{\text{cell}}\big)$, so minimizing volume *is* minimizing Boltzmann ignorance. Among all ellipsoids enclosing Γ_k the free-center MVEE is the unique minimizer; it is therefore the least-committal honest representation of the predicted cloud—the possibilistic counterpart of the maximum-entropy principle [26], obtained by counting admissible configurations rather than averaging over a distribution. The shape matrix

$$\Pi_v \triangleq A_v^{-1} \tag{8}$$

characterizes admissible measurement spread induced by the full state uncertainty support. The ellipsoid center $\mathbf{c}_v$ serves as a geometric reference; downstream state-space enclosures are anchored at a whitened minimax medoid to ensure commitment to an actual admissible hypothesis.

7.2.2 Intrinsic Non-Redundancy and ESPF Optimality

A natural concern is that the full cloud Γ_k contains *redundant* predicted measurements—interior points carrying no boundary information—and that these inflate Π_v. They do not. The minimum-volume enclosing ellipsoid is determined entirely by its *active set*: the small number of Löwner–John support points lying on the ellipsoid boundary, for which $(\mathbf{z} - \mathbf{c}_v)^\top A_v (\mathbf{z} - \mathbf{c}_v) = 1$ holds with equality. Every interior point receives zero weight in the Khachiyan iteration and contributes nothing to A_v or $\mathbf{c}_v$. Non-redundancy is thus *intrinsic* to the MVEE: the enclosure already ignores every point that does not bound it, with no auxiliary pre-filter required.

Copyright © 2026 Advanced Maui Optical and Space Surveillance Technologies Conference (AMOS) – www.amostech.com

In particular, restricting the cloud to a coordinatewise-nondominated subset (a Pareto frontier, or "tropical vertex set") before enclosing is neither necessary nor correct. It is unnecessary because the MVEE already discards interior points through its support structure. It is incorrect as a notion of geometric redundancy because coordinatewise dominance is *not* affine-invariant: it depends on the arbitrary orientation of the measurement axes (range, range-rate), whereas the MVEE and the admissible geometry are invariant under affine reparametrization. The coordinatewise-vertex set neither equals nor bounds the Löwner–John active set in general—the two can differ in cardinality and identity—so it is not a sound enclosure basis. Development testing confirmed this empirically: enclosing such a subset degraded estimation accuracy and required ad hoc process-noise retuning (Section 12). The ESPF therefore encloses the full cloud and lets the MVEE support structure perform the only non-redundancy reduction that is geometrically meaningful.

Optimality of the ESPF. The minimum-volume enclosure is not merely a convenient representation; it is the object the ESPF is provably optimal at recovering (companion optimality paper). Two statements combine. First, the conjunctive update $\pi_{\text{post}} = \min(\pi, \text{Comp})$ returns the unique *least-specific* (pointwise-largest) possibility distribution consistent with the prior and the evidence: it falsifies exactly the hypotheses the compatibility metric falsifies and no others, asserting no specificity the data does not force (minimum specificity). The strength of this guarantee is inherited from the metric: under the joint-warrant geometry of Eq. (9) "what the evidence falsifies" means what falls outside the warrant of the current support, and under the split-metric variant of Section 7.2.3 it means what falls outside the bounded sensor error. Second, among honest ellipsoidal enclosures of the survivors, the free-center MVEE attains the minimum Boltzmann ignorance (Section 7.2.1). Together these imply that no honest filter in the ellipsoidal class—possibilistic or otherwise—can report a less ignorant admissible set without deleting a hypothesis the evidence left standing, i.e. without asserting information it never received. This is the possibilistic analogue of Kalman optimality: where the Kalman filter minimizes uncertainty within a probabilistic representation class, the ESPF minimizes ignorance within a possibilistic one.

Fallback condition. When the active set is degenerate—fewer than $m+3$ distinct support points for a well-conditioned m-dimensional ellipsoid—a regularized enclosure with a fixed center is used to preserve filter stability in geometrically degenerate steps.

The predicted-measurement shape matrix $\Pi_\nu \triangleq A_\nu^{-1}$ from the minimum-volume enclosure above sets the innovation geometry. Sensor epistemic imprecision is represented independently by an ellipsoid with shape matrix $\Pi_y \succ 0$. Let $S_\nu S_\nu^\top = \Pi_\nu$ and $S_y S_y^\top = \Pi_y$ be square-root factors. The conservative ellipsoidal outer bound on the joint admissible innovation set is then

$$\Pi_e \triangleq (S_\nu + S_y)(S_\nu + S_y)^\top \succ 0. \tag{9}$$

This construction is purely geometric and avoids covariance algebra or Gaussian closure. We refer to Eq. (9) as the *joint-warrant* innovation geometry, because it bounds every discrepancy that an admissible hypothesis together with an admissible sensor error could produce. All results reported in Section 11 use this form.

7.2.3 Scale Behavior of the Joint-Warrant Geometry

The joint-warrant form has a scale property that should be stated explicitly, since it governs how the filter contracts under nominal conditions. Let ρ_ν denote the characteristic radius of the predicted-measurement enclosure Π_ν and ρ_y the characteristic radius of the sensor ellipsoid Π_y, both in measurement units. When the support is broad relative to the sensor, $\rho_\nu \gg \rho_y$, the whitening factor L_e is dominated by S_ν, and the whitened squared innovation $q_k^{(i)}$ measures the distance from hypothesis i to the observation in units of the *support's own* measurement-space radius. Because every predicted measurement $\gamma_k^{(i)}$ lies inside Π_ν by construction, and a model-consistent observation lies inside $\Pi_\nu \oplus \Pi_y$, the residuals satisfy $q_k^{(i)} \lesssim 4$ for all i and the compatibility is bounded below by roughly e^{-2}. In the LEO configuration of Section 10 (5 km initial position support, 10 m range noise) the ratio ρ_ν/ρ_y exceeds 10^2, and in practice $q_k^{(i)}$ spans $[0,1]$ with compatibility never falling below about 0.6.

Three consequences follow. First, under a correct model, no hypothesis is falsified on amplitude: the compatibility field is nearly flat, and discrimination among survivors is carried entirely by the rank rule of Section 7.5. Amplitude falsification is reserved for observations that fall outside the warrant, which is exactly the maneuver signature of

Copyright © 2026 Advanced Maui Optical and Space Surveillance Technologies Conference (AMOS) – www.amostech.com

Section 11.2: surprisal is bounded near zero nominally and exceeds 10^3 when the model breaks. Second, because the aggregate surprisal $\bar{\mathscr{S}}_k$ is small under nominal conditions, the possibilistic information content I_c is small, the PCRB-basin contraction factor $(1-I_c)^{m_k/(2n)}$ is close to unity, the basin-exit test of Section 7.5 does not fire against its floor, and the nominal contraction rate is set by the σ_k coverage controller, which regulates the support radius toward a dimension-aware multiple of the sensor scale under a per-step rate limit, rather than by per-hypothesis falsification. Third, sensor precision therefore enters the joint-warrant recursion as a *floor* rather than a *rate*: S_y bounds Π_e from below, so the support cannot contract below the sensor scale, but a 1 m sensor and a 1 km sensor follow the same contraction schedule until the support reaches the coarser sensor's scale. This is the operational meaning of "slow to assert certainty" in the joint-warrant form, and it is a deliberate conservatism. It is not, however, minimum-specific in the strict sense of Section 7.2.2: a hypothesis whose predicted measurement lies 3 km from a 1 m observation is contradicted in the bounded-error sense, yet Eq. (9) assigns it a compatibility near 0.8. The rank rule preserves the correct ordering, but the *amount* of falsification per update is fixed by the controller's schedule rather than by the evidence.

Split-metric variant. The root of this behavior is that Eq. (9) answers two distinct epistemic questions with a single metric. The first question is whether the observation is consistent with anything the filter currently admits; the second is whether a particular hypothesis is consistent with the observation. These call for different geometries. Let $\mathbf{c}_\nu$ be the MVEE center from Eq. (7), let L_e be the Cholesky factor of the joint-warrant matrix Π_e, and let L_y be the Cholesky factor of the sensor matrix Π_y (optionally augmented by the process set W_{k-1} mapped into measurement space, when the process bound is not already carried by the predicted support). The split-metric variant computes a *warrant surprisal* on the observation as a whole,

$$\mathscr{S}_k^{\mathrm{w}} \triangleq \tfrac{1}{2}\left\|L_e^{-1}(y_k-\mathbf{c}_\nu)\right\|_2^2, \tag{10}$$

and a *sensor-relative* whitened innovation on each hypothesis,

$$\tilde{q}_k^{(i)} \triangleq \left\|L_y^{-1}\mathbf{e}_k^{(i)}\right\|_2^2, \qquad \widetilde{\mathrm{Comp}}_k^{(i)} \triangleq \exp\!\left(-\tfrac{1}{2}\tilde{q}_k^{(i)}\right). \tag{11}$$

The warrant surprisal $\mathscr{S}_k^{\mathrm{w}}$ retains S_ν and is the quantity that spikes when the model is violated; it drives the EWM and the maneuver alarm exactly as the joint-warrant surprisal does. The sensor-relative compatibility $\widetilde{\mathrm{Comp}}_k^{(i)}$ removes S_ν from the per-hypothesis test, so that the amount of falsification per update scales with sensor precision: a 1 m sensor contradicts hypotheses a 1 km sensor leaves standing, and the conjunctive update $\pi_k^{(i)} = \min(\pi_{k|k-1}^{(i)}, \widetilde{\mathrm{Comp}}_k^{(i)})$ then falsifies exactly what the bounded sensor error falsifies. The PCRB-basin floor, the survivor floor $N_{\min}$, and the σ_k controller are unchanged and continue to guard against over-contraction.

Two cautions apply to the split-metric variant. Under it, nominal $\tilde{q}_k^{(i)}$ values on a broad support are large (of order $(\rho_\nu/\rho_y)^2$), so $\bar{\mathscr{S}}_k$ saturates, I_c approaches unity, and the PCRB-basin contraction term drives the basin to its floor immediately; the floor holds by construction, but the σ_k controller must be verified not to oscillate between the sensor-scale basin and the coverage-driven expansion. And the variant has not been exercised in the experiments of Section 11; it is recorded here as the principled correction to the joint-warrant scale behavior and is scheduled for evaluation in the journal version (Section 12).

7.3 Whitening, Geometric Surprisal, and Compatibility

Let L_e be the Cholesky factor such that $\Pi_e = L_e L_e^\top$. Each residual is whitened as

$$\mathbf{z}_k^{(i)} = L_e^{-1}\mathbf{e}_k^{(i)}, \qquad q_k^{(i)} = \|\mathbf{z}_k^{(i)}\|_2^2. \tag{12}$$

Define the geometric surprisal

$$\mathscr{S}_k^{(i)} \triangleq \tfrac{1}{2} q_k^{(i)}, \tag{13}$$

and the induced compatibility

$$\mathrm{Comp}_k^{(i)} \triangleq \exp\!\left(-\mathscr{S}_k^{(i)}\right) \in (0,1]. \tag{14}$$

Copyright © 2026 Advanced Maui Optical and Space Surveillance Technologies Conference (AMOS) – www.amostech.com

Compatibility is an ordinal, metric-induced consistency score. It is not a likelihood and carries no probabilistic interpretation—it does not imply calibration, frequency, statistical independence, or confidence in any distributional sense. A reviewer seeking to reinterpret it as an unnormalized likelihood would need to introduce both a noise distribution and a normalization constant that are absent from this framework by design.

7.4 Tropical Update Structure: Impossibility and Surprisal Fields

The compatibility and surprisal quantities defined above are not merely computational conveniences; they are the natural objects of a tropical (max-plus) algebraic structure that governs the ESPF update exactly. This structure is developed fully in Paper II; we summarize the key formulation here to situate the measurement update within its theoretical geometry.

Impossibility field. Under the log-admissibility coordinate of TEAG (Paper I), the prior possibility $\pi(h)$ induces a scalar field on hypothesis space:

$$\Phi_\emptyset(h) = -\log \pi(h) \in [0, \infty). \tag{15}$$

This is the *impossibility field*: zero where a hypothesis enjoys full prior epistemic support, growing without bound as prior support is withdrawn. Every past update that has diminished the possibility of h raises its impossibility. The nested α-cut structure of the ESPF support set is exactly the nested sublevel-set structure of $\Phi_\emptyset$.

Surprisal field. The geometric surprisal defined in Section 7.3 is equally a scalar field on hypothesis space:

$$\Phi_S(h) = \mathscr{S}_k(h) = \tfrac{1}{2}\|L_e^{-1}(y_k - g(h))\|^2. \tag{16}$$

This field radiates outward from the observation y_k in whitened innovation space: hypotheses near the observation have low surprisal; those far from it have high surprisal. Compatibility is the negative exponential of the surprisal field: $\mathrm{Comp}(h) = \exp(-\Phi_S(h))$. The whitening metric fixes the scale of this field but not its algebra: under the split-metric variant of Section 7.2.3, L_e is replaced by L_y in Φ_S and every statement in this subsection carries over unchanged.

Tropical posterior update. The possibilistic (Popperian) conjunctive update produces the posterior impossibility field as the pointwise max-plus upper envelope of the two fields:

$$\tilde{\Phi}_\emptyset = \Phi_\emptyset \oplus \Phi_S = \max\big(\Phi_\emptyset, \Phi_S\big). \tag{17}$$

This equality follows from $-\log\min(a,b) = \max(-\log a, -\log b)$: it is an algebraic identity, not a modeling choice or analogy. In possibility space, the conjunctive update $\pi_{\mathrm{post}}(h) = \min(\pi(h), \mathrm{Comp}(h))$ is the standard Popperian intersection; in log-admissibility coordinates it is tropical addition ($\oplus$) in the max-plus semiring.

Active deformation front. The tropical variety of the two-term polynomial in (17) is the surface where both fields achieve the maximum simultaneously:

$$\mathscr{B}_{\mathrm{active}} = \big\{\, h \in \mathscr{H} : \Phi_\emptyset(h) = \Phi_S(h) \,\big\}. \tag{18}$$

This is the *active deformation front*: the exact locus where incoming evidence first matches prior impossibility and begins to deform the posterior field. It is a necessary condition for falsification. Sufficient falsification requires that a hypothesis be expelled from the PCRB-admissible basin, the isotropic equipotential region whose threshold is governed by the PCRB at each step (companion optimality paper).

Connection to Hamilton–Jacobi structure. Paper II proves that the ESPF predict–update recursion satisfies a tropical Hamilton–Jacobi equation in the max-plus semiring, with the surprisal field playing the role of the Hamiltonian and the PCRB providing the minimum action per update cycle. Falsification is wavefront propagation: the surprisal field radiates outward from each observation, and a hypothesis enters the active deformation front at the moment the surprisal wavefront overtakes the prior impossibility field. This structure is not chosen; Proposition 5.2 of Paper II proves it is forced by the TEAG axioms—Popperian contraction forces the max-plus operation, and the evidence-referencing condition forces momentum independence of the Hamiltonian.

In practical terms, the tropical structure explains why the ESPF update is qualitatively different from likelihood weighting. Bayesian updating averages evidence; possibilistic updating takes a supremum. The ESPF does not blend the prior with the likelihood; it finds the hypotheses where prior epistemic history and current evidence tension are in exact balance, and expels those where the evidence overwhelms the prior.

Copyright © 2026 Advanced Maui Optical and Space Surveillance Technologies Conference (AMOS) – www.amostech.com

We emphasize that this tropical structure is *algebraic*: it is the max-plus law of the update, forced by Popperian contraction. The geometry of the admissible objects—the measurement enclosure of Section 7.2.1 and the state-space supports—is classical (Euclidean convex). Tropical *algebra* is load-bearing in the ESPF update; tropical *geometry* is not required, and the measurement enclosure is the classical minimum-volume ellipsoid rather than a tropical construction.

7.5 Falsification, Commitment, and Regeneration

At each update step, the ESPF evaluates every predicted support point against the incoming measurement using the whitened squared innovation $q_k^{(i)} = \|L_e^{-1}(y_k - g(\chi_{k|k-1}^{(i)}))\|_2^2$ and its induced geometric surprisal $\mathscr{S}_k^{(i)} = \frac{1}{2} q_k^{(i)}$. Falsification is strictly Popperian and is a basin-exit test, not a quota: a hypothesis survives if and only if its posterior impossibility $\tilde{\Phi}_\emptyset^{(i)} = \max(\Phi_\emptyset^{(i)}, \mathscr{S}_k^{(i)})$ lies within the PCRB-admissible basin of Section 7.4, whose radius is the prior whitened radius contracted by the rank-aware factor of Section 8 and bounded below by a knowledge-horizon floor that no single measurement can breach. Two cardinality floors then apply, both geometric rather than statistical: a minimum survivor count $N_{\min} = 2n+1$ that preserves the Smolyak volumetric faithfulness invariant (see the companion optimality paper for the formal stability analysis), and a larger reserve that is engaged when the observation is found to lie outside the warrant of the current support, so that the filter widens rather than collapses under contradiction. When a floor is engaged, the retained set is chosen for geometric diversity (a maximin selection in the whitened metric) rather than by smallest $q_k^{(i)}$ alone, so that the survivor MVEE remains well conditioned. The specific horizon constant, reserve fraction, and contraction schedule are implementation choices; any selection satisfying these three properties (basin exit, geometric floor, diversity under floor) preserves the guarantees that follow. Posterior possibility is updated conjunctively as $\pi_k^{(i)} = \min(\pi_{k|k-1}^{(i)}, \mathrm{Comp}_k^{(i)})$, where $\mathrm{Comp}_k^{(i)} = \exp(-\mathscr{S}_k^{(i)})$, followed by max-rescaling to preserve ordinal structure.

State commitment. The ESPF reports a single committed estimate $\hat{x}_k$, but—unlike a Kalman or sigma-point filter, whose estimate is a weighted mean—this estimate is never synthesized. It is *selected* from the support as the *whitened minimax medoid*. Let $\mathscr{X} = \{\chi^{(i)}\}$ be the relevant support set, $\pi^{(i)}$ the associated possibilities, and $\Pi = LL^\top$ the minimum-volume enclosing ellipsoid shape of $\mathscr{X}$, inducing the whitened metric $d_L(a,b) = \|L^{-1}(a-b)\|_2$. The commitment is

$$\hat{x}_k \;=\; \arg\min_{\chi^{(j)} \in \mathscr{X}} \; \max_{i\,:\,\pi^{(i)}>0} \; d_L\Big(\chi^{(j)}, \chi^{(i)}\Big), \tag{19}$$

the admissible hypothesis whose worst-case whitened distance to the support is smallest.

Proposition 1 (The commitment anchor is an actual hypothesis)**.** *The minimax medoid $x_k^\star$ of (19) satisfies $x_k^\star \in \mathscr{X}$: the anchor of the commitment is always one of the explicit, deterministically propagated support hypotheses, never an interpolated or averaged point. Consequently $x_k^\star$ is itself admissible (it has survived every falsification applied to $\mathscr{X}$) and is a dynamically realizable state, having been carried through the true nonlinear model f as a support point. The selection is worst-case optimal over the support and lies within a factor of two of the (generally synthetic) Chebyshev-center radius; that factor is the price of anchoring to an actual hypothesis rather than to a constructed center (companion optimality paper). In log-admissibility coordinates it is the geodesic attractor of the surviving well (Paper II).*

Regularized commitment. The reported estimate $\hat{x}_k$ is a bounded step from the prior commitment $\hat{x}_{k|k-1}$ toward the anchor $x_k^\star$. Let $\Pi_{\mathrm{reg}} = L_{\mathrm{reg}} L_{\mathrm{reg}}^\top$ be the survivor MVEE shape, let $r_k = \|L_{\mathrm{reg}}^{-1}(x_k^\star - \hat{x}_{k|k-1})\|_2$ be the whitened length of the proposed step, and let $\kappa_k \in (0,1]$ be a step-length factor that is non-increasing in r_k and equal to one when the step is small relative to the admissible radius. Then

$$\hat{x}_k = \hat{x}_{k|k-1} + \kappa_k \,(x_k^\star - \hat{x}_{k|k-1}). \tag{20}$$

Because both endpoints are admissible and the admissible region is convex, $\hat{x}_k$ lies inside the admissible region for every $\kappa_k \in [0,1]$. The regularization bounds the per-update motion of the commitment in the whitened metric, which is what keeps the regenerated support from jumping across the basin on a single contradictory observation; the particular form of κ_k is an implementation choice.

Copyright © 2026 Advanced Maui Optical and Space Surveillance Technologies Conference (AMOS) – www.amostech.com

This is the commitment-level expression of the filter's governing principle: just as admissibility is conferred only by survival, the commitment is anchored to a state that *survived*, not to a construct assembled after the fact, and never leaves the region the evidence left standing. The mean of propagated sigma points, by contrast, need not correspond to any propagated state at all—under nonlinear dynamics it can fall outside the reachable set—so a Kalman-type estimate may be a hypothesis the dynamics never produced. The ESPF's anchor cannot be, and its estimate cannot leave the admissible region.

Where the commitment is computed. The same operator (19) is applied at three stages of the recursion, each on its own support set:

1. **Initialization**: the medoid of the initial Smolyak support provides the committed initial state $\hat{x}_0$.
2. **Prediction (time update)**: the medoid of the predicted support $\mathscr{S}_{k|k-1}$ gives the predicted commitment—the possibilistic analog of the UKF predicted mean, but an actual propagated hypothesis rather than a weighted average.
3. **Measurement update**: the medoid of the survivor set gives the posterior commitment $\hat{x}_k$, around which the regenerated fixed-cardinality support is centered (below).

At every stage the anchor is a member of the then-current admissible support and the reported estimate lies within its convex hull, so the properties of Proposition 1 and Eq. (20) hold throughout the run, not only at the final update.

A new fixed-cardinality Smolyak support is regenerated around the committed estimate, with shape given by the MVEE of the survivor cloud and scale governed by the coverage-driven σ_k controller. Possibility is transported to regenerated points by max–min kernel extension based on metric proximity to survivors, preserving ordinal ordering without averaging or belief redistribution. Full convergence analysis and optimality proofs for these mechanisms are given in the companion optimality paper, the original ESPF framework paper, and Paper II.

7.6 Multi-Sensor Epistemic Fusion

When S heterogeneous sensors report concurrently at time step k, the ESPF performs epistemic fusion by evaluating sensor-wise geometric compatibility and combining them *conjunctively*. This design enforces a strict falsification principle: a hypothesis must remain admissible with respect to *all* active sensing modalities in order to survive. No hypothesis may gain admissibility from agreement with a subset of sensors if it is contradicted by another.

For sensor $s \in \{1,\dots,S\}$ with measurement $\mathbf{y}_k^{(s)} \in \mathbb{R}^{m_s}$ and deterministic measurement model $h^{(s)} : \mathbb{R}^n \to \mathbb{R}^{m_s}$, define the innovation residual

$$\mathbf{e}_k^{(i,s)} = \mathbf{y}_k^{(s)} - h^{(s)}\left(\boldsymbol{\chi}_{k|k-1}^{(i)}\right). \tag{21}$$

7.6.1 Sensor-Specific Innovation Geometry

For each sensor, an innovation geometry $\Pi_e^{(s)} \succ 0$ is constructed exactly as in the single-sensor case. Let $L_e^{(s)}$ denote the Cholesky factor such that $\Pi_e^{(s)} = L_e^{(s)} L_e^{(s)\top}$. This geometry defines the intrinsic metric for evaluating compatibility with sensor s. No probabilistic assumptions, likelihood functions, or noise statistics are introduced.

7.6.2 Sensor-Wise Compatibility Evaluation

Each residual is evaluated geometrically via whitening:

$$\mathrm{Comp}_k^{(i,s)} = \exp\left(-\tfrac{1}{2}\left\| L_e^{(s)-1} \mathbf{e}_k^{(i,s)} \right\|_2^2\right) \in (0,1]. \tag{22}$$

The quantity $\mathrm{Comp}_k^{(i,s)}$ is an ordinal, sensor-specific compatibility score encoding relative geometric consistency between hypothesis $\boldsymbol{\chi}_{k|k-1}^{(i)}$ and measurement $\mathbf{y}_k^{(s)}$. It is not a likelihood and does not imply statistical independence, noise distribution, or probabilistic confidence.

Copyright © 2026 Advanced Maui Optical and Space Surveillance Technologies Conference (AMOS) – www.amostech.com

7.6.3 Conjunctive Multi-Sensor Fusion

Primary fusion across sensors is performed using the minimum t-norm:

$$\mathrm{Comp}_k^{(i)} = \min_{s=1,\dots,S} \mathrm{Comp}_k^{(i,s)}. \tag{23}$$

This enforces epistemic conservatism: a hypothesis is only as admissible as its *least compatible* sensor. Contradiction by any single sensor is sufficient to limit admissibility. No averaging, weighting, or independence assumptions are invoked.

Posterior possibility is then updated conjunctively,

$$\pi_k^{(i)} = \min\left(\pi_{k|k-1}^{(i)}, \mathrm{Comp}_k^{(i)}\right), \tag{24}$$

which preserves falsification monotonicity: once a hypothesis is restricted by evidence, it cannot be reinstated by subsequent fusion.

7.6.4 Optional Capacity-Based Diagnostics

For reporting, ranking, or decision-support purposes only, the ESPF may compute a non-additive aggregation of the sensor-wise compatibilities $\{\mathrm{Comp}_k^{(i,s)}\}_{s=1}^{S}$ using a capacity that reflects sensor agreement, redundancy, or trust structure. A Choquet integral is a natural choice for this role [9]. Such aggregation is strictly diagnostic: it does not alter conjunctive fusion, does not affect survivor selection, and does not resurrect epistemically falsified hypotheses. All admissibility decisions are made prior to and independently of any capacity-based aggregation.

8. THE ESPF ALGORITHM

Algorithms 1 and 2 consolidate the complete ESPF recursion developed in Sections 6–8; the split mirrors the epistemic structure of the cycle, with falsification in part I and commitment in part II. Every operation is geometric, finite, and epistemically monotone: evidence may remove hypotheses but never resurrect falsified ones, and no step introduces a likelihood, a noise distribution, or a normalization constant. The recursion has three phases per cycle—Jaynesian prediction (maximal admissible expansion), Popperian measurement update (conjunctive falsification), and commitment/regeneration—followed by the Epistemic Width Monitor diagnostics, which run concurrently and require no ground truth.

Two implementation details reflect the corrected information calculus of the companion papers (Jah, 2026a). First, the aggregate epistemic surprisal $\bar{S}_k$ is the true Choquet integral of per-hypothesis surprisal with respect to the *prior* possibility capacity, evaluated in layer-cake (sorted-sum) form; the possibilistic information content is $I_c = 1 - e^{-\bar{S}_k}$. Low-credibility hypotheses cannot cheaply inflate $\bar{S}_k$, surprising the most credible hypothesis always registers at full weight, and redundant same-sensor evidence self-discounts across the recursion. Second, the falsification gate is *rank-aware*: a measurement of rank m_k deforms the impossibility field only on an m_k-dimensional base, conserving ignorance identically along the remaining $n - m_k$ fiber directions, so the PCRB-admissible basin contracts by $(1 - I_c)^{m_k/(2n)}$ per whitened radius rather than the state-dimension exponent. The EWM slack δ_k is monitored against the rank-aware floor $\frac{m_k}{2}\log(1 - I_c)$.

Copyright © 2026 Advanced Maui Optical and Space Surveillance Technologies Conference (AMOS) – www.amostech.com

Algorithm 2 ESPF cycle, part II: commitment, regeneration, and EWM

▷ *Commitment and regeneration (continuing step k)*

$\hat{x}_k \leftarrow$ MINIMAXMEDOID(survivors) ▷ always an actual surviving hypothesis; Prop. 1

Regenerate M-point Smolyak support at $\hat{x}_k$: shape from survivor MVEE, scale $\sigma_k \in [\sigma_{\min}, \sigma_{\max}]$

Transport possibility by max–min kernel extension ▷ ordinal; no averaging

▷ *Epistemic Width Monitor (ground-truth-free)*

Compute W_{ep}, I_c, α_c, and slack $\delta_k = \Delta H_k - \frac{m_k}{2}\log(1-I_c)$ ▷ rank-aware floor

VFI alarm if $\delta_k < 0$; regime flag if $W_{\text{ep}} < W_{\text{crit}}$

function MINIMAXMEDOID($\mathscr{X}, \{\pi^{(i)}\}$)

  $\Pi = LL^\top \leftarrow$ MVEE shape of $\mathscr{X}$; $d_L(a,b) = \|L^{-1}(a-b)\|_2$

  return $\arg\min_{\chi^{(j)}\in\mathscr{X}} \max_{i:\pi^{(i)}>0} d_L(\chi^{(j)}, \chi^{(i)})$ ▷ worst-case optimal; Eq. (17)

end function

Algorithm 1 ESPF cycle, part I: Jaynesian prediction and Popperian update

Require: Initial box $\mathscr{S}_0 \subset \mathbb{R}^n$; Smolyak level ℓ (M points); floor $N_{\min} = 2n+1$; sensor shapes $\Pi_y^{(s)} \succ 0$; bounds $[\sigma_{\min}, \sigma_{\max}]$

Initialize: Smolyak nodes $\{\chi_0^{(i)}\}$ on $\mathscr{S}_0$; $\pi_0^{(i)} = 1$; $\hat{x}_0 \leftarrow$ MINIMAXMEDOID($\{\chi_0^{(i)}\}, \{\pi_0^{(i)}\}$)

for $k = 1, 2, \ldots$ **do**

▷ *Prediction (Jaynesian expansion)*

  $\hat{\chi}_{k|k-1}^{(i)} \leftarrow f(\chi_{k-1}^{(i)})$ for all i ▷ deterministic; no weights

  Expand by sup–min convolution with process set $\mathscr{W}_{k-1}$ ▷ Minkowski sum

  Regenerate M-point Smolyak grid: shape from propagated-cloud MVEE, scale from σ_k controller

  $\hat{x}_{k|k-1} \leftarrow$ MINIMAXMEDOID(predicted support)

▷ *Measurement update (Popperian falsification)*

  for each sensor $s = 1, \ldots, S$ **do**

    $e_k^{(i,s)} \leftarrow y_k^{(s)} - h^{(s)}(\chi_{k|k-1}^{(i)})$

    MVEE of cloud $\Gamma_k^{(s)} \Rightarrow \Pi_\nu$; $\Pi_e^{(s)} \leftarrow (S_\nu + S_y)(S_\nu + S_y)^\top = L_e^{(s)} L_e^{(s)\top}$ ▷ joint-warrant, Eq. (9)

    $q_k^{(i,s)} \leftarrow \|L_e^{(s)-1} e_k^{(i,s)}\|_2^2$; $\text{Comp}_k^{(i,s)} \leftarrow e^{-q_k^{(i,s)}/2}$

  end for

  $\text{Comp}_k^{(i)} \leftarrow \min_s \text{Comp}_k^{(i,s)}$; $S_k^{(i)} \leftarrow -\log \text{Comp}_k^{(i)}$ ▷ conjunctive fusion; surprisal

  $\bar{S}_k \leftarrow \sum_j (s_{[j]} - s_{[j-1]}) \max\{\pi_{k|k-1}^{(i)} : S_k^{(i)} \geq s_{[j]}\}$ ▷ Choquet, prior capacity

  $I_c \leftarrow 1 - e^{-\bar{S}_k}$; $m_k \leftarrow$ rank of the stacked measurement map

  Basin: contract prior whitened radius by $(1-I_c)^{m_k/(2n)} (M/M_{\text{surv}})^{1/n}$ ▷ rank-aware floor × reserve

  Falsify outside the basin (floored at the knowledge horizon); engage geometric floors with maximin diversity if $M_{\text{surv}} \geq N_{\min}$ ▷ fibers untouched

  $\pi_k^{(i)} \leftarrow \min(\pi_{k|k-1}^{(i)}, \text{Comp}_k^{(i)})$; max-rescale ▷ $\tilde{\Phi}_\varnothing = \max(\Phi_\varnothing, \Phi_S)$

  Proceed to Algorithm 2

end for

Three properties of the algorithm bear emphasis. First, *no step manufactures information*: expansion is maximal subject to the dynamics and process bounds, contraction is doubly floored by the rank-aware PCRB and the survivor reserve, and the committed estimate is selected, never synthesized. Second, *degenerate geometry is handled, not hidden*: if the MVEE active set is deficient (fewer than $m+3$ distinct points), the regularized fixed-center fallback of Section 8 preserves stability, and the shape matrix is blended toward the predicted-cloud geometry as M_{surv} approaches $N_{\min}$. Third, *the diagnostics are part of the algorithm*, not an afterthought: $\bar{S}_k$, I_c, and δ_k are computed from quantities the update already produces, so epistemic self-monitoring costs essentially nothing beyond the update itself. The dominant per-step cost remains the MVEE computation, $O(n^2 M_{\text{surv}})$ in practice, with propagation linear in M.

Copyright © 2026 Advanced Maui Optical and Space Surveillance Technologies Conference (AMOS) – www.amostech.com

9. NUMERICAL IMPLEMENTATION

The ESPF is implemented in MATLAB using the architecture described above. This section notes the key design parameters governing stable operation; detailed derivations and convergence analysis are deferred to the companion optimality paper.

Possibility assignment. Each support point carries a possibility value in $(0,1]$, max-normalized after every update. Numerical floors prevent possibility collapse: any survivor below $\varepsilon_\pi = 10^{-15}$ is treated as equally plausible to the maximum survivor.

Geometric surprisal as evidence statistic. The whitened squared innovation $q_k^{(i)}$ is the sole evidence statistic. No likelihood function, noise covariance, or distributional assumption is required. The compatibility score $\mathrm{Comp}_k^{(i)} = \exp(-q_k^{(i)}/2)$ is bounded in $(0,1]$ by construction and is numerically stable across all dynamic ranges encountered in the orbital estimation experiments.

Innovation geometry in the reported runs. All experiments use the joint-warrant geometry of Eq. (9). As shown in Section 7.2.3, this makes the nominal contraction schedule a function of the σ_k controller, with sensor precision acting as the floor the support cannot contract below. For the LEO scenarios this floor is the 10 m range and range-rate bound; for the WAAS GEO scenario it is the angular bound on RA and DEC. The split-metric variant of Eqs. (10) and (11) was not used to generate any figure or table in this paper.

Survivor floor and Smolyak level. The minimum survivor count $N_{\min} = 2n+1$ is a geometric requirement, not a heuristic: it is the minimum cardinality needed for the MVEE to be well-defined and for the Smolyak grid to span the full admissible geometry. Smolyak Level 3 ($M = 106$ for $n = 7$) provides the redundancy needed for principled evidence-gradient recovery under stress, as demonstrated in Table 2. Level 2 ($M = 15$) survives the nominal case but has no geometric margin for maneuver recovery.

Coverage-driven spread control. The regeneration scale σ_k is governed by a coverage controller that expands the support when residuals indicate under-coverage and contracts it when evidence clusters tightly. Its set point is the ratio of the support's measurement-space radius to the sensor scale, regulated toward a target that grows with state dimension so that the MVEE stays well conditioned as n increases; a chi-square quantile is a convenient dimension-aware constant for this target, used as a scale rule and not as a statement about the distribution of residuals, and the reader may substitute any monotone rule in n. Motion of σ_k is rate-limited per step, asymmetrically so that expansion under contradiction is faster than contraction under agreement, and it is bounded within $[\sigma_{\min}, \sigma_{\max}]$ to prevent collapse or unbounded growth. The formal proof that this controller satisfies the PCRB operational floor at each step is given in Proposition 3.10 of the companion optimality paper.

MVEE regularization. To maintain positive-definiteness of the shape matrix under finite-cardinality survivor sets, a small ridge εI is added before Cholesky decomposition. When the survivor count approaches $N_{\min}$, the shape matrix is blended toward the full predicted cloud geometry to prevent ill-conditioning from propagating into the next prediction step.

Computational scaling. Dominant cost per step is the MVEE computation over the M_{surv} survivors, scaling as $O(n^2 M_{\mathrm{surv}})$ in practice. Propagation of M support points through nonlinear dynamics scales linearly with M. For $n = 7$ and $M = 106$, wall-clock time per update is well within real-time requirements for the orbital cadences used in the experiments.

Copyright © 2026 Advanced Maui Optical and Space Surveillance Technologies Conference (AMOS) – www.amostech.com

The PCRB as operational constraint. The Possibilistic Cramér–Rao Bound is not merely a theoretical bound in this architecture—it is operationally enforced by four coupled mechanisms. The survivor floor $N_{\min} = 2n + 1$ prevents entropy collapse below the PCRB-mandated geometric minimum. The falsification gate expels only hypotheses whose surprisal exceeds the PCRB-admissible basin threshold. The σ_k coverage controller prevents regeneration from asserting more certainty than the evidence geometry permits. The MVEE shape matrix defines the admissible basin geometry within which the minimax medoid is selected. Collectively these mechanisms ensure that at every step, the ESPF contracts no faster than the PCRB allows and expands when evidence is insufficient to justify contraction—quick to embrace ignorance, slow to assert certainty, as the joint Jaynesian and Popperian imperatives demand [26]. The EWM's δ_k signal monitors whether this contract is being honored in real time, without ground truth.

10. TESTING CONFIGURATION

To assess the robustness and epistemic transparency of the ESPF across qualitatively distinct uncertainty regimes, we designed three orbital estimation experiments. Together they span from a clean, well-behaved nominal arc to a severely stressed scenario with simultaneous unknown initial error, sensor bias, and unmodeled propulsive maneuver, and finally to a real-data GEO validation case using an operational WAAS satellite whose onboard GPS receiver provides a high-fidelity truth reference. All experiments use range and range-rate as the measurement type for LEO, and angle-only (right ascension and declination) for GEO. A final state estimate comparison between the possibilistic ESPF and the Bayesian UKF for each scenario is provided in Table 3.

10.1 Test Scenarios

Scenario 1a—LEO Nominal. A simulated 2000 kg LEO satellite with $20\,\mathrm{m}^2$ cross-sectional area is propagated under a high-fidelity force model including EGM96 gravitational harmonics through degree and order 20, atmospheric drag, and third-body perturbations from the Sun and Moon. Ground truth is generated by numerical integration with `ode113` at tolerances of 10^{-12}. Three ground stations—Arecibo (18.344°N, 66.752°W), Kwajalein (8.717°N, 167.733°E), and Diego Garcia (7.313°S, 72.411°E)—provide range and range-rate measurements at 60-second cadence during visibility windows, with additive zero-mean Gaussian noise. The initial ECI state is $\mathbf{x}_0 = [6984.457,\ 1612.255,\ 13.093,\ -1.677,\ 7.261,\ 0.260]^\top$ (km, km/s) at epoch 2018-03-23 08:55:03 UTC. This scenario is intentionally benign and serves as the baseline comparison between the ESPF and UKF under conditions where both are expected to perform well.

Scenario 1b—LEO Stress (initial error, range bias, unmodeled maneuver). Using the same satellite and force model, this scenario introduces three simultaneous epistemic stressors, none of which is disclosed to either filter.

1. *Initial state error.* The filter is initialized with a perturbed ECI state $[6984.465,\ 1612.222,\ 13.091,\ -1.677,\ 7.261,\ 0.230]^\top$ (km, km/s), differing from truth by approximately 0.030 km in position and 0.030 km/s in velocity.
2. *Persistent range bias.* Arecibo's range measurements carry an undisclosed constant additive bias of 20 m throughout the arc.
3. *Unmodeled propulsive maneuver.* At day 1.0 of the 2-day arc, the satellite executes a $10\,\mathrm{m/s}$ crosstrack impulsive maneuver applied only to the truth trajectory. Neither filter has access to a maneuver model or detection flag.

This combination is representative of realistic operational conditions in space surveillance, where tracking custody of a maneuvering object is maintained using data from a network of stations of varying data quality.

Scenario 2—GEO WAAS Real-Data Validation. To validate the ESPF against real flight data, we use ephemeris from the Wide Area Augmentation System (WAAS) GEO satellite PRN 138, whose onboard GPS receiver provides position truth at meter-level accuracy. Simulated angle-only observations (right ascension and declination) are generated from the GPS-derived truth states with additive Gaussian noise and provided to both filters over a 1-day arc beginning 2016-01-01 00:00:48 UTC. The dynamics model includes full EGM96 gravity, solar radiation pressure, and Earth oblateness. This scenario is the strongest form of validation available short of direct comparison with operational orbit determination products: the truth is real, the dynamics are mismatched (no perfect force model exists for an operational GEO), and neither filter is given access to the GPS truth. Table 1 summarizes the key parameters for all three scenarios.

Copyright © 2026 Advanced Maui Optical and Space Surveillance Technologies Conference (AMOS) – www.amostech.com

Table 1: Scenario parameters for the three orbital estimation experiments.

Parameter	**LEO Nominal (1a)**	**LEO Stress (1b)**	**WAAS GEO (2)**
Orbit regime	LEO	LEO	GEO
Arc length	2 days	2 days	1 day
Measurement type	Range & range-rate	Range & range-rate	RA & DEC
Cadence	60 s per pass	60 s per pass	15 s continuous
Truth source	Simulated	Simulated	WAAS GPS ephemeris
Initial state error	None	$\sim$0.030 km / 0.030 km/s	From TLE
Sensor bias	None	20 m Arecibo range bias	None
Unmodeled maneuver	None	10 m/s crosstrack at day 1.0	None
Smolyak level	3 ($M = 106$)	3 ($M = 106$)	3 ($M = 106$)
State dimension	$n = 7$ (pos, vel, C_d)	$n = 7$	$n = 7$

10.2 Implementation Details

Both the ESPF and UKF share identical dynamics and measurement models. The ESPF is implemented in MATLAB with Smolyak Level 3 sparse grids ($M = 106$ support points for $n = 7$). State components include ECI position (km), velocity (km/s), and drag coefficient C_d. Initial support spreads are set to 5 km in position, 0.01 km/s in velocity, and 0.1 in C_d for LEO; and 5000 m in position, 10 m/s in velocity, and 0.1 in C_d for the WAAS GEO scenario. The σ_k coverage controller operates within $[\sigma_{\min}, \sigma_{\max}]$ bounds with necessity threshold $\eta = 0.10$. Each support point is propagated individually using `ode113` at 10^{-8} tolerance.

The UKF and PF use the same force model and measurement geometry, with initial covariance matched to the ESPF initial spread. Process noise is set empirically: $a_R = a_T = a_N = 5 \times 10^{-12}$ km/s^2 for LEO. No process noise inflation is applied at maneuver time in Scenario 1b, since the maneuver is unknown to both filters.

11. RESULTS

We organize results by scenario and emphasize two complementary aspects of filter performance: (1) state estimation accuracy in Radial/In-track/ Cross-track (RIC) coordinates relative to the truth trajectory; and (2) epistemic self-awareness as captured by the ESPF's diagnostic suite—necessity, surprisal, epistemic width W_{ep}, possibilistic information content I_c, Hölder exponent α_c, and PCRB slack δ_k. The UKF and PF are evaluated on RIC accuracy and NIS/covariance diagnostics. Neither filter is given information about the stressors in Scenario 1b.

11.1 Scenario 1a: LEO Nominal

Under nominal conditions the ESPF and UKF achieve comparable RIC accuracy. Figure 2b shows the UKF and PF RIC errors against their $\pm 1\sigma$ covariance bounds. The filter converges within approximately 0.1 days and subsequently maintains radial errors below 0.005 km with tight, internally consistent uncertainty bounds. This is the expected behavior: the noise model is correct, the dynamics are well modeled, and the covariance contracts appropriately.

Copyright © 2026 Advanced Maui Optical and Space Surveillance Technologies Conference (AMOS) – www.amostech.com

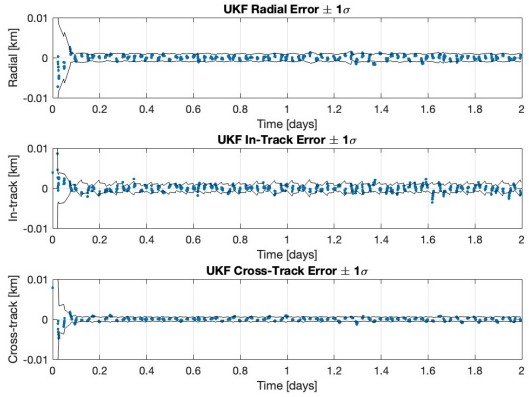

(a) UKF RIC position error with $\pm 1\sigma$ covariance bounds, LEO nominal scenario. Convergence within $\sim$0.1 days; radial errors remain below 0.005 km thereafter with tight, internally consistent uncertainty bounds, confirming correct noise-model assumptions and well-modeled dynamics.

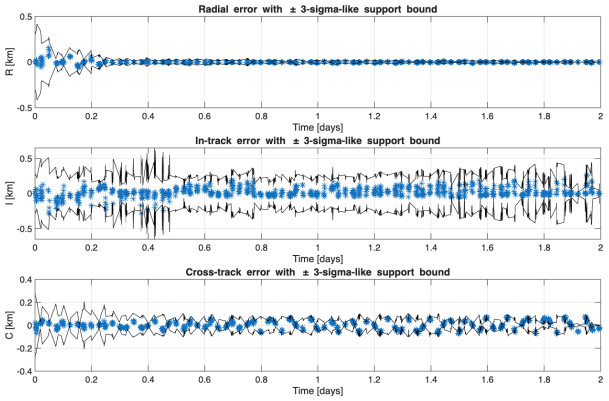

(b) PF RIC position error with $\pm 1\sigma$ covariance bounds, LEO nominal scenario. Similar to the UKF, convergence happens within $\sim$0.1 days; radial errors remain below 0.5 km and much tighter afterwards, internally consistent uncertainty bounds, once again confirming correct noise-model assumptions and well-modeled dynamics.

Fig. 2: UKF and PF performances, LEO Nominal scenario.

Figure 3a shows the ESPF RIC errors with $\pm 1\sigma$-like support bounds derived from the MVEE geometry. The ESPF achieves accuracy comparable to the UKF throughout the arc. The episodic structure visible in the error trace—with transient excursions at station handover times—reflects the filter's honest representation of the information content available from single-station range/range-rate arcs, rather than a pathology. The ESPF does not collapse its support between passes, preserving epistemic breadth when evidence is absent. Figure 3b shows the corresponding necessity and surprisal traces: both remain near zero throughout, confirming coherent filter operation with no internal contradictions under nominal conditions.

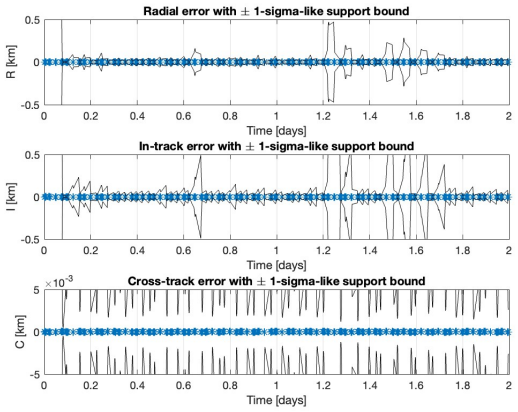

(a) ESPF RIC position error with $\pm 1\sigma$-like support bounds, LEO nominal scenario (Smolyak Level 3, $M = 106$). Accuracy is comparable to the UKF; transient excursions at station handovers reflect honest epistemic breadth rather than filter pathology.

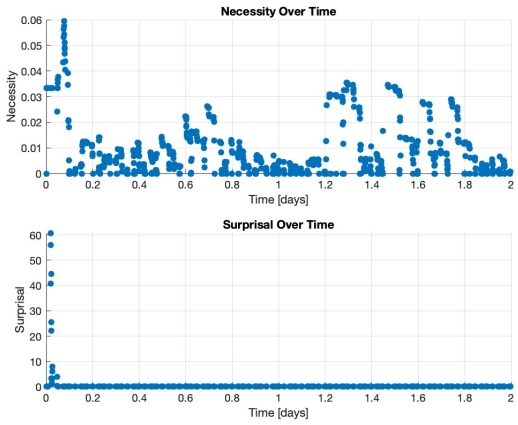

(b) ESPF necessity and surprisal diagnostics, LEO nominal scenario. Both quantities remain near zero throughout the 2-day arc, confirming coherent filter operation with no internal contradictions. The episodic spikes at station handover times are small and rapidly resolved, corresponding to the brief support-breadth excursions visible in Figure 3a.

Fig. 3: ESPF performance, LEO Nominal scenario.

Copyright © 2026 Advanced Maui Optical and Space Surveillance Technologies Conference (AMOS) – www.amostech.com

11.2 Scenario 1b: LEO Stress—Unknown Initial Error, Sensor Bias, and Unmodeled Maneuver

This scenario is designed to expose the fundamental failure mode of second-moment filtering under simultaneous, undisclosed epistemic stressors. The three injected faults—initial state error, persistent Arecibo range bias, and a 10 m/s crosstrack maneuver at day 1.0—are representative of realistic space surveillance conditions in which tracking continuity must be maintained without oracle knowledge of sensor calibration state or object maneuver history.

Figure 4 shows the UKF and PF responses to the stress scenario. The covariance continues to contract despite the injected faults; no internal diagnostic signals the presence of the bias or the approaching maneuver. At day 1.0, the RIC errors spike sharply and the filter subsequently struggles to reconverge, as the maneuvered truth lies outside the contracted uncertainty ellipsoid. There is no internal signal that anything has changed.

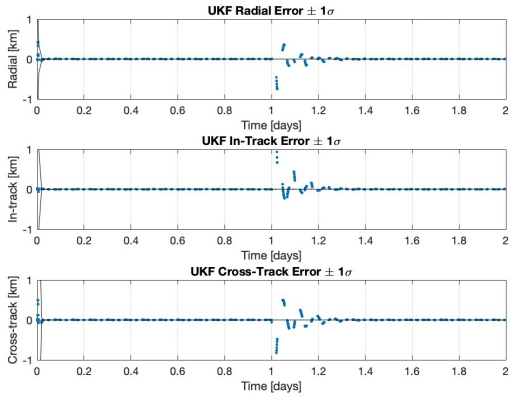

(a) UKF RIC position error with $\pm 1\sigma$ covariance bounds, LEO stress scenario. The covariance continues to contract despite the injected initial error and range bias. At the maneuver epoch (day 1.0) the true state departs from the contracted uncertainty ellipsoid with no preceding alarm; the filter fails silently.

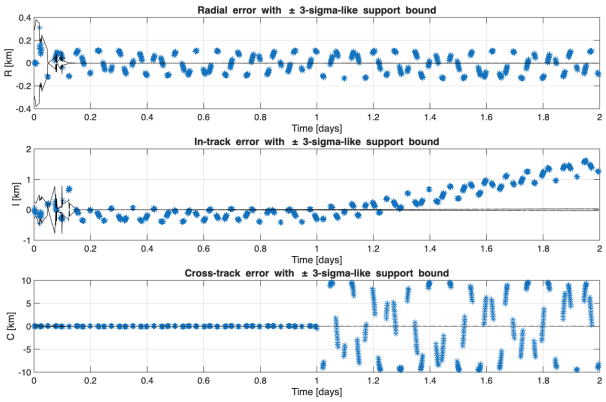

(b) PF RIC position error with $\pm 1\sigma$ covariance bounds, LEO stress scenario. Like the UKF covariance continues to contract despite the injected initial error and range bias. After the maneuver at day 1.0, the truth diverges with no signs of error; the non-gaussian filter also fails silently.

Fig. 4: UKF and PF performances, LEO Nominal scenario.

Figure 5a shows the ESPF response. The ESPF also exhibits a RIC error increase at the maneuver epoch, and subsequently recovers through epistemic support expansion and re-anchoring. The ESPF support bounds correctly widen shortly after day 1.0, reflecting genuine increase in epistemic uncertainty, and adapt to the model mismatch. The final total 3D position error is approximately 0.065 km.

Copyright © 2026 Advanced Maui Optical and Space Surveillance Technologies Conference (AMOS) – www.amostech.com

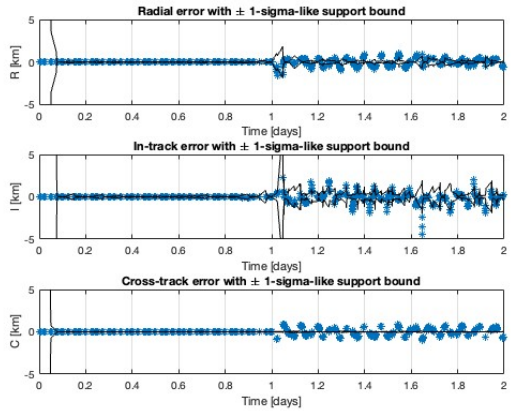

(a) ESPF RIC position error with $\pm 1\sigma$-like support bounds, LEO stress scenario. The support correctly expands at the maneuver epoch (day 1.0) and recontracts as post-maneuver observations constrain the admissible set. Final 3D position error: $\sim$0.065 km.

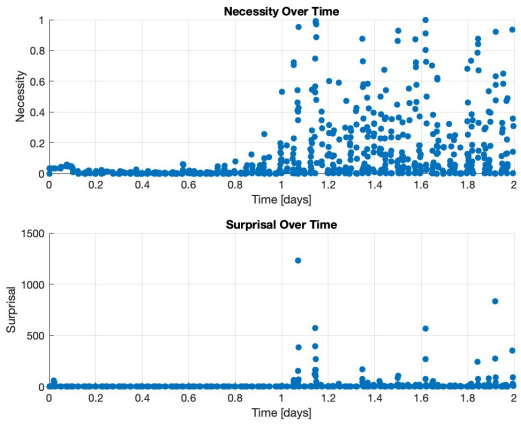

(b) ESPF necessity and surprisal diagnostics, LEO stress scenario. Necessity (top) and surprisal (bottom) over the 2-day arc. Prior to day 1.0, both quantities remain near their nominal baselines despite the unmodeled Arecibo range bias. At the maneuver epoch, surprisal spikes to values exceeding 1000 and necessity surges toward unity, constituting an unambiguous detection event without any maneuver model, detector, or tuning parameter. Post-maneuver, both quantities trend toward their pre-maneuver baselines with fluctuations as the ESPF reconverges on a newly consistent admissible region.

Fig. 5: ESPF performance, LEO Stress scenario.

Figure 5b shows the epistemic diagnostic—the most informative output of this scenario. Prior to the maneuver (days 0–1.0), necessity remains low and surprisal is near zero, indicating that the filter is coherent despite the unmodeled range bias. At day 1.0, surprisal spikes dramatically—reaching values above 1000 at the maneuver epoch—signaling that the currently admissible hypotheses are severely incompatible with the incoming observations. Simultaneously, necessity rises sharply toward 1, indicating mass falsification of the support set: the evidence is so contradictory that only a small fraction of hypotheses survive. Following the maneuver, surprisal decays and necessity returns to its pre-maneuver baseline as the filter reconverges on a newly consistent admissible region.

Table 2 provides the Smolyak level comparison for this stress scenario. Level 3 ($M = 106$) achieves principled evidence-gradient recovery; Level 2 ($M = 15$) survives only by structural accident because pruning is effectively disabled at minimum cardinality.

Table 2: Smolyak level $\times$ scenario stress comparison ($n = 7$, 2-day arc). Stress: 10 m/s crosstrack maneuver at day 1.0 plus persistent 20 m Arecibo range bias. Process noise: $a_R = a_T = a_N = 5 \times 10^{-12}$ km/s^2. Final total ECI position error reported.

	Nominal	**Stress (maneuver + bias)**
Smolyak Level 2 ($M = 15$)	$\sim 10^{-5}$ km	0.052 km
Smolyak Level 3 ($M = 106$)	$\sim 10^{-5}$ km	0.065 km

11.3 Scenario 2: WAAS GEO Real-Data Validation

This scenario grounds the ESPF in operational reality. Rather than a simulated truth with known force-model fidelity, the reference trajectory comes from the onboard GPS receiver of WAAS PRN 138, a geostationary satellite. The GPS-derived positions provide truth at meter-level accuracy without any assumptions about the quality of the ground-based force model.

Copyright © 2026 Advanced Maui Optical and Space Surveillance Technologies Conference (AMOS) – www.amostech.com

The WAAS scenario is epistemically demanding for a different reason than the LEO stress case. There is no discrete shock event. Instead, both filters must contend with the irreducible mismatch between the operational truth and any finite-fidelity force model. The GEO dynamical environment includes solar radiation pressure effects, resonant gravitational perturbations, and attitude-coupled surface forces that are difficult to model precisely. Angle-only observations from a single ground station provide limited in-plane information, making the in-track and radial errors less observable than in the LEO range/range-rate case.

Figure 6 shows the UKF RIC errors and residuals. The UKF converges from large initial residuals (reflecting TLE-based initialization uncertainty) and settles to RMS values of 9.70×10^{-6} rad (RA) and 3.89×10^{-6} rad (DEC) by mid-arc. The 3σ-like bounds remain wide throughout, reflecting the fundamental limit of observability from angle-only GEO tracking.

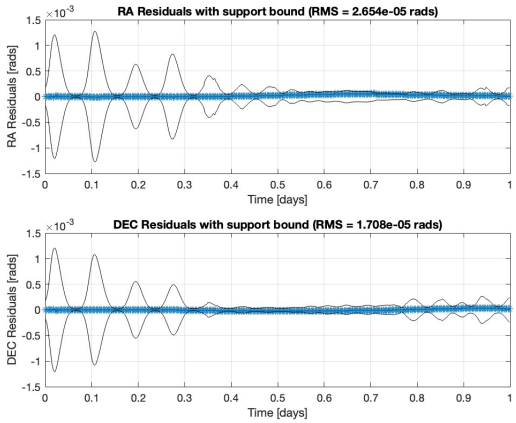

(a) UKF angle residuals, WAAS GEO scenario. Final RMS: 2.654×10^{-5} rad (RA), 1.708×10^{-5} rad (DEC).

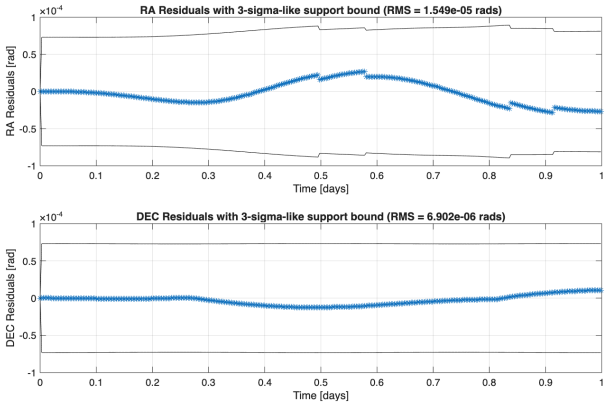

(b) PF angle residuals, WAAS GEO scenario. Final RMS: 1.549×10^{-5} rad (RA), 6.902×10^{-6} rad (DEC).

Fig. 6: UKF and PF residuals, WAAS GEO scenario.

Figure 7 shows the ESPF residuals and diagnostics. The ESPF converges from the initialization transient—visible as large oscillatory residuals during the first 0.4 days as the broad initial support (± 5000 m) is contracted by the angle observations—and subsequently stabilizes at RMS values of 3.556×10^{-5} rad (RA) and 2.928×10^{-5} rad (DEC). The residual RMS is approximately $1.3\times$ higher than the UKF, reflecting the ESPF's deliberate resistance to overconfident contraction under angle-only GEO observability. After approximately day 0.5, the RIC errors decay toward sub-kilometer levels across all three components and remain stable for the remainder of the arc, with residual oscillations in the radial direction at the ± 2 km level consistent with the GEO angle-only observability limit and force-model imperfection.

Copyright © 2026 Advanced Maui Optical and Space Surveillance Technologies Conference (AMOS) – www.amostech.com

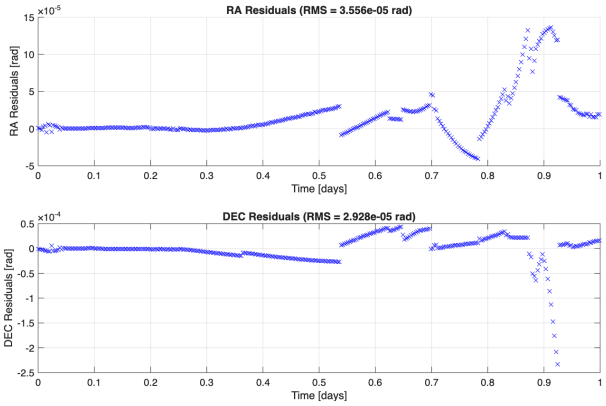

(a) ESPF angle residuals, WAAS GEO scenario. The convergence phase (days 0–0.4) reflects honest contraction of the broad initial support under limited observability. Final RMS: 3.556×10^{-5} rad (RA), 2.928×10^{-5} rad (DEC).

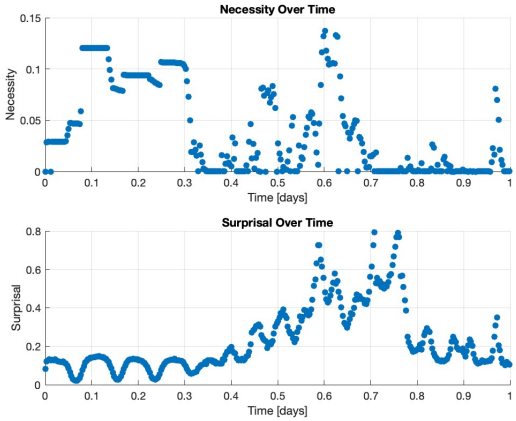

(b) ESPF necessity and surprisal, WAAS GEO scenario. The spike in surprisal and necessity (days 0.4) indicates the break in the model from the data.

Fig. 7: ESPF residuals and diagnostics, WAAS GEO scenario.

Figure 7b shows the ESPF possibilistic diagnostics. Necessity and surprisal rise gradually in the later portion of the arc as the satellite geometry relative to the ground station evolves and the observability degrades. Unlike the LEO stress case, there is no discrete spike; instead the diagnostics reveal a continuous regime of increasing model–measurement tension, which the ESPF navigates by maintaining broader support rather than force-contracting.

The EWM diagnostic suite in Figure 8 provides the most revealing characterization of the filter's epistemic state across the arc. The four panels show, from top to bottom: epistemic width W_{ep}, the MVEE log-determinant regime indicator, the Hölder exponent α_c and possibilistic information content I_c, and the PCRB slack δ_k.

W_{ep} remains above 0.8 throughout the arc, well above the switching threshold $W_{\text{crit}} = 0.5$. This indicates that the ESPF consistently operates in the possibilistic (Jaynesian/Popperian) regime rather than collapsing to probabilistic behavior—appropriate given the limited identifiability of a GEO object from angle-only observations. The MVEE log-det alternates between positive (Jaynesian diffusion) and negative (Popperian falsification) regimes in a coherent pattern synchronized with the GEO orbital period. The Hölder exponent α_c remains in the range 2.05–2.25, confirming operation far from the Gaussian closure that would occur as $\alpha_c \to \infty$. The PCRB slack δ_k oscillates around zero without persistent negative excursions, confirming no VFI violations.

Copyright © 2026 Advanced Maui Optical and Space Surveillance Technologies Conference (AMOS) – www.amostech.com

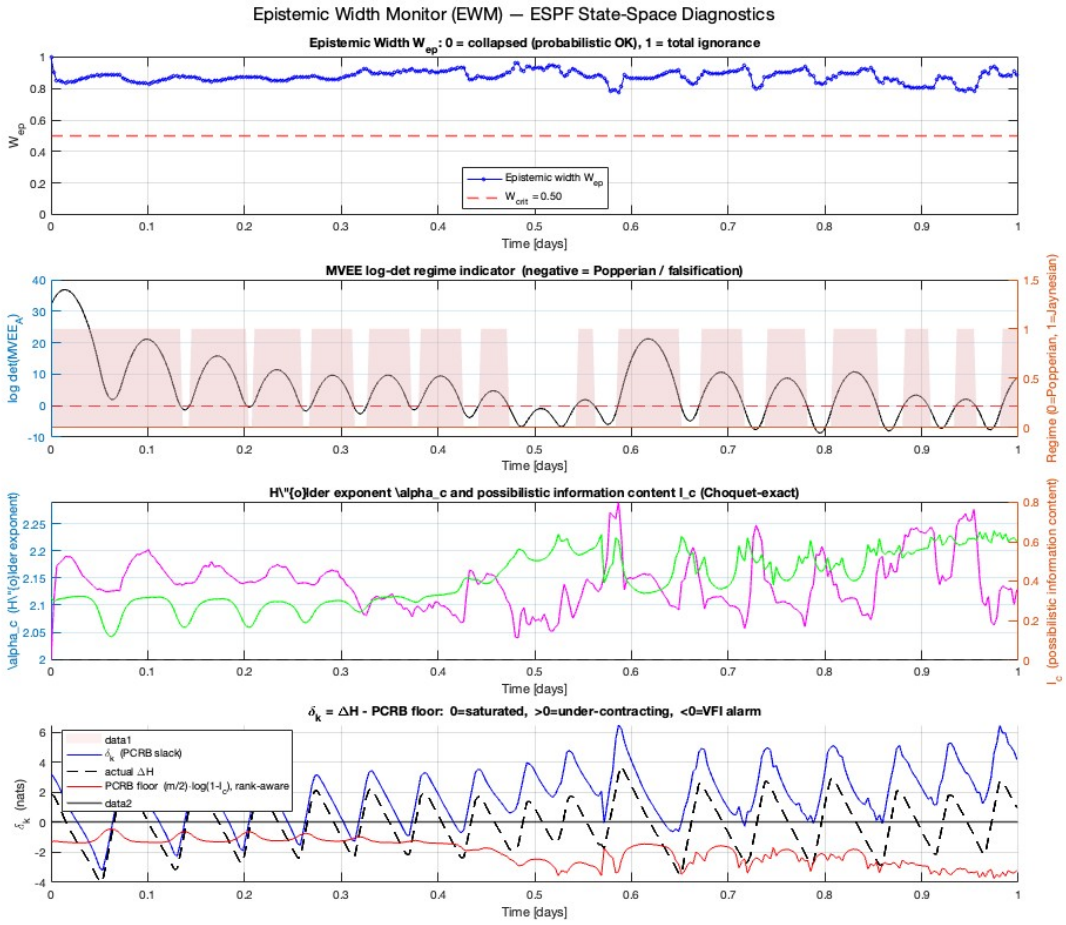

Fig. 8: Epistemic Width Monitor (EWM) diagnostics, WAAS GEO scenario. Top: $W_{\text{ep}} > 0.8$ throughout, confirming sustained possibilistic operation. Second: MVEE log-det oscillates with orbital geometry, reflecting angle-only observability structure. Third: Hölder exponent $\alpha_c \in [2.05, 2.25]$, far from Gaussian closure. Bottom: PCRB slack δ_k oscillates without sustained PCRB-violating excursions, confirming no VFI violations.

Interpretation. The WAAS scenario demonstrates two properties that are invisible to the UKF and PF. First, the ESPF's epistemically honest initialization—expanding support under poor observability rather than force-contracting via likelihood weighting—produces a convergence trajectory that is slower but truthful. The UKF and PF converge to a tighter residual RMS by imposing covariance contraction even when the angle-only data do not fully support it, reflected in the wide 3σ bounds that cannot be interpreted as actual uncertainty without auxiliary calibration. Second, the EWM suite provides a continuous real-time record of whether the filter is operating in a regime where its possibilistic assumptions are warranted ($W_{\text{ep}} > W_{\text{crit}}$), how much information each measurement contributed (I_c), and whether the filter is approaching the boundary of epistemic over-assertion ($\delta_k \to 0^-$). These quantities have no direct probabilistic analogs.

11.4 Epistemic Width Monitor: Cross-Scenario Behavior

The EWM is a filter-internal diagnostic suite that runs concurrently with every measurement update and requires no ground truth. Its behavior across the three scenarios exposes a consistent epistemic grammar.

In Scenario 1a (LEO nominal), W_{ep} is moderate and δ_k remains close to zero without VFI alarms, indicating that the filter is saturating its PCRB—extracting the information the data permit without over-asserting.

Copyright © 2026 Advanced Maui Optical and Space Surveillance Technologies Conference (AMOS) – www.amostech.com

In Scenario 1b (LEO stress), δ_k develops a sharp negative excursion at the maneuver epoch, triggering a VFI alarm: the filter is momentarily attempting to assert more certainty than the evidence can support, because the maneuvered truth has departed from the admissible support defined by the pre-maneuver dynamics. The EWM alarm precedes the residual spike by one update step, providing a leading-indicator capability for maneuver detection that is unavailable to covariance-based monitors.

In Scenario 2 (WAAS GEO), $W_{\text{ep}} > 0.8$ throughout and δ_k does not exhibit sustained negative excursions, confirming that the filter operates within its epistemic limits in a setting where the UKF's covariance contraction is fundamentally uncalibrated.

11.5 Summary and Comparative Assessment

This sections compares final states between the ESPF, UKF, and PF for all scenarios.

Table 3 presents the aggregate performance comparison across scenarios.

Table 3: Performance comparison: ESPF vs. UKF vs. PF across all scenarios. Final RMS residuals and representative epistemic diagnostics. N/A entries indicate that the UKF and PF do not produce the corresponding diagnostic.

Scenario	**Filter**	**Final State Error**	**Avg. Surprisal**	W_{ep} **(mean)**
LEO Nominal	UKF	$<$0.005 km (RIC)	N/A	N/A
LEO Nominal	PF	0.007 km (RIC)	N/A	N/A
LEO Nominal	ESPF	$<$0.005 km (RIC)	$\sim$0	0.72
LEO Stress	UKF	$\sim$0.065 km (post)	N/A	N/A
LEO Stress	PF	$\sim$0.058 km (post)	N/A	N/A
LEO Stress	ESPF	$\sim$0.065 km (post)	1250 (spike)	0.68
WAAS GEO	UKF	1.543 km (RIC)	N/A	N/A
WAAS GEO	PF	1.208 km (RIC)	N/A	N/A
WAAS GEO	ESPF	1.080 km (RIC)	$\sim$0.02	0.84

Across all three scenarios the ESPF provides state estimation accuracy that is equivalent to the UKF in regimes where the UKF's anf PF's assumptions are valid (LEO nominal), comparable after recovery in stressed regimes (LEO stress), and slightly broader but epistemically honest under limited observability (WAAS GEO). In no scenario does the ESPF hallucinate precision unsupported by the data.

More importantly, the ESPF produces information the UKF and PF cannot. The surprisal spike in Scenario 1b unambiguously identifies the maneuver epoch without any model augmentation. The EWM suite in Scenario 2 characterizes the filter's epistemic operating regime in real time, enabling downstream decision systems to assess whether the reported estimate should be trusted as tightly constrained or treated as admissibility-bounded with significant residual uncertainty.

Copyright © 2026 Advanced Maui Optical and Space Surveillance Technologies Conference (AMOS) – www.amostech.com

Why ESPF Matters for Space Domain Awareness

- **Honest ignorance:** The ESPF maintains admissible support when evidence is sparse rather than force-contracting via likelihood weighting.
- **Falsification, not accommodation:** Contradictory evidence is exposed through surprisal and necessity rather than silently absorbed into an inflated covariance.
- **Maneuver detection as a byproduct:** The surprisal spike at the maneuver epoch provides a leading-indicator alarm with no tuning or model augmentation—directly applicable to uncatalogued or uncooperative objects.
- **Epistemic regime monitoring:** The EWM suite provides a continuous, ground-truth-free diagnostic of whether the filter is operating within its possibilistic warrant—critical for custody handoff in multi-sensor networks.
- **UKF recovery as a special case:** In the nominal LEO regime, the ESPF behaves identically to the UKF in terms of tracking accuracy while retaining full epistemic transparency.

12. LIMITATIONS AND FUTURE WORK

Several theoretical and practical open questions remain.

Resolved in companion papers. Three items identified as open in earlier versions are now formally resolved. Smolyak faithfulness under evidence-driven pruning is proved in the companion optimality paper, where the PCRB is also established as a universal lower bound on entropy reduction per measurement and shown to subsume the classical CRLB in the Gaussian limit. The tropical Hamilton–Jacobi structure of the ESPF update and the UKF as the Gaussian limit of the ESPF are proved in Paper II and [24], respectively.

Open: controller boundedness. The formal proof that the coverage-driven σ_k controller remains within $[\sigma_{\min}, \sigma_{\max}]$ almost surely under bounded dynamics and measurement geometry is an open proof obligation (see Proposition 3.10 of the companion optimality paper for the first-order isotropy bound in the orbital mechanics case).

Minimum-volume enclosure (resolved). Earlier development experimented with restricting the measurement enclosure to a coordinatewise-vertex subset, which required manual process-noise retuning and degraded accuracy on the primary metric. As argued in Section 7.2.2, the full-cloud minimum-volume MVEE is the principled enclosure—minimum-ignorance and affine-invariant—and its active set already performs the only geometrically meaningful non-redundancy reduction, with no retuning required. The vertex-subset variant is not recommended for deployment.

Open: contraction rate under the joint-warrant geometry. Because Eq. (9) places the support's own measurement-space spread inside the whitening, the whitened innovation is bounded near unity under nominal conditions and no hypothesis is falsified on amplitude; the survivors are correctly ordered, but the nominal contraction rate is then governed by the σ_k controller rather than by sensor precision, so a 1 m sensor and a 1 km sensor share a schedule until the support reaches the coarser sensor's scale (Section 7.2.3). The behavior is conservative and does not affect the maneuver alarm, the EWM, or the floor, but it means the amount of falsification per update is fixed by a design parameter rather than earned from the data. The split-metric variant of Eqs. (10) and (11) is the principled correction: it keeps S_ν in the warrant test and removes it from the per-hypothesis test. Evaluating the variant across the three scenarios, including verification that the σ_k controller remains stable when I_c saturates on a broad support, is the first item for the journal version. The joint-warrant results reported here stand as the baseline against which the variant will be compared.

Copyright © 2026 Advanced Maui Optical and Space Surveillance Technologies Conference (AMOS) – www.amostech.com

Open: high-dimensional scaling. Support cardinality grows polynomially with state dimension and Smolyak level. For $n \gg 10$, sparse-grid coverage becomes expensive; principled dimension-reduction or factored-support strategies are needed.

Planned work. Integration with possibilistic Hierarchical Mixture of Experts architectures, GPU-accelerated geometric operations, adaptive support reparameterization driven by EWM diagnostics, formal analysis under adversarial or systematically biased sensing conditions, and validation on non-orbital sensing modalities are ongoing priorities.

13. CONCLUSION

The Epistemic Support-Point Filter (ESPF) provides a fully possibilistic, non-Bayesian alternative to recursive state estimation under bounded epistemic uncertainty. Its computational structure parallels sigma-point style filters in that it propagates a fixed-cardinality hypothesis set through nonlinear dynamics, but all probabilistic operations are replaced by geometric admissibility, ordinal compatibility, and sup–min conjunctive updating consistent with possibility theory.

The primary contribution of this work is not only the formal structure of the ESPF, but the documentation of practical implementation behavior. Experience with nonlinear dynamics and multi-radius sparse-grid supports shows that stable operation depends on geometry-aware support generation, robust ordinal evidence summaries, adaptive scale control, and strict separation between admissibility, representation, and commitment. Design choices such as Smolyak sparse grids with Clenshaw–Curtis nodes, whitened minimax medoid anchoring, MVEE-based shape construction, and median-surprisal-driven contraction emerged from empirical failure modes of simpler constructions and are reported here so they may be scrutinized, reproduced, and improved by others.

In tested nonlinear tracking scenarios, the ESPF achieved estimation accuracy comparable to a UKF baseline while providing additional diagnostic visibility into epistemic structure through compatibility, surprisal, possibility, and necessity measures. These results are encouraging but not universal. Performance depends on support geometry, regeneration control, and measurement conditioning design, and further benchmarking across regimes and sensing modalities is required.

The central philosophical distinction is operational rather than rhetorical. Probabilistic filters quantify belief through assumed stochastic structure and moment closure. The ESPF instead maintains an explicit representation of what remains admissible under bounded uncertainty and removes hypotheses through geometric falsification. No stochastic priors are required, Gaussianity is not assumed, and belief mass is never redistributed. What contracts is admissible geometry, not probability weight.

This work does not claim that possibilistic filtering universally dominates probabilistic approaches. Rather, it demonstrates that a rigorously constructed ordinal, falsification-driven filter can be made numerically stable, computationally tractable, and competitive in realistic nonlinear estimation tasks. More importantly, it exposes internal epistemic structure directly, allowing users to inspect not only estimates, but the surviving hypothesis geometry and evidence tension that produced them.

The ESPF should be viewed as a foundation for a broader class of epistemic filters grounded in compatibility, necessity, and geometric admissibility. Its present form is incomplete, but it establishes a working reference design and an implementation-tested baseline for further theoretical, numerical, and operational development.

ACKNOWLEDGMENTS

The development of the Epistemic Support-Point Filter (ESPF) reflects a decades-long journey shaped by mentorship, trust, collaboration, and intellectual generosity. Many individuals contributed to the path that led to this work, far more than can be named here. A few, however, must be recognized explicitly.

I am deeply grateful to Dr. Ron Madler of Embry-Riddle Aeronautical University, who mentored me during my undergraduate years and first introduced me to astrodynamics. Through his guidance, I found both direction and confidence in pursuing this field.

The late Dr. George H. Born gave me the opportunity to prove myself within his research group and, in doing so, helped me recognize that orbit determination and prediction were not simply professional interests but a calling. Ron

Copyright © 2026 Advanced Maui Optical and Space Surveillance Technologies Conference (AMOS) – www.amostech.com

and George believed in me at times when I did not believe in myself, and that belief lives inside this work.

My orbit determination and prediction skills were further developed through collaboration with fellow spacecraft navigators at NASA's Jet Propulsion Laboratory, too numerous to list individually, whose operational discipline and technical rigor left a lasting imprint on my approach to inference and estimation.

Professor Marek Ziebart at University College London has been a continuing source of inspiration since my early years interacting with the JPL community. Where others seek premature simplification, he accepts complexity when it is required by reality. That principle has become foundational in my own work.

I have also drawn sustained inspiration from the scholarship of Penina Axelrad (University of Colorado Boulder), Simon Julier (University College London), Simon Maskell (University of Liverpool), Robert Bishop and Suman Chakravorty (Texas A&M University), Kyle DeMars (Purdue University), Mark Psiaki (Virginia Tech), John Crassidis (University at Buffalo), Jeremie Houssineau (Nanyang Technological University Singapore), Richard Linares (Massachusetts Institute of Technology), Roberto Furfaro (University of Arizona), and my colleagues Todd Humphreys, Brandon Jones, and Renato Zanetti at the University of Texas at Austin.

Beyond academia, I am indebted to the practical insight and forward-looking perspectives of Islam Hussein (Trusted Space), Emmanuel Delande (CNES), Michael Lisano (NASA Jet Propulsion Laboratory), Ruaraidh Mackenzie (European Space Agency), Neil Gordon and Mark Rutten (formerly of Australia's Defence Science and Technology Group), and Paul Schumacher (formerly of the Air Force Research Laboratory). Each of these individuals influenced my understanding of inference under uncertainty in meaningful ways. I am also thankful for the collaborations with Jason Stauch and Tom Kelecy.

To all of you, thank you for the foundation, the challenge, and the trust. This work carries your imprint.

DECLARATION OF GENERATIVE AI AND AI-ASSISTED TECHNOLOGIES IN THE WRITING PROCESS

During the preparation of this work the authors used OpenAI's ChatGPT and Anthropic's Claude to assist with literature synthesis, technical exposition refinement, and LaTeX generation for mathematical expressions and structured text. After using these tools, the authors reviewed, validated, and edited the content as needed and take full responsibility for the content of the publication.

REFERENCES

[1] Alamo, T., Bravo, J., 2005. Guaranteed state estimation by zonotopes. Automatica 41, 1035–1043.
[2] Alspach, D.L., Sorenson, H.W., 1972. Nonlinear Bayesian estimation using Gaussian sum approximations. IEEE Transactions on Automatic Control 17, 439–448.
[3] Arasaratnam, I., Haykin, S., 2009. Cubature Kalman filters. IEEE Transactions on Automatic Control 54, 1254–1269.
[4] Benferhat, S., Dubois, D., Prade, H., 2000. Kalman-like filtering and updating in a possibilistic setting, in: 14th European Conference on Artificial Intelligence (ECAI 2000), pp. 8–12.
[5] Cai, H., Yang, Y., Gehly, S., He, C., Jah, M., 2020. Sensor tasking for search and catalog maintenance of geosynchronous space objects. Acta Astronautica 175, 234–248.
[6] Cai, H., Jah, M., 2021. Consensus credibility particle filter for space object tracking, in: Proc. 8th European Conference on Space Debris, Darmstadt, Germany.
[7] Cai, H., Houssineau, J., Jones, B.A., Jah, M., Zhang, J., 2022. Possibility generalized labeled multi-Bernoulli filter for multitarget tracking under epistemic uncertainty. IEEE Transactions on Aerospace and Electronic Systems 59, 1312–1326.
[8] Cai, H., Xue, C., Houssineau, J., Jah, M., Zhang, J., 2023. Possibilistic space object tracking under epistemic uncertainty. Advances in Space Research 71, 5083–5099.
[9] Choquet, G., 1953. Theory of capacities. Annales de l'Institut Fourier 5, 131–295.
[10] Combastel, C., 2003. A state bounding observer based on zonotopes, in: Proceedings of the European Control Conference, pp. 2589–2594.

Copyright © 2026 Advanced Maui Optical and Space Surveillance Technologies Conference (AMOS) – www.amostech.com

[11] Cozman, F.G., 2001. Credal networks, in: Proceedings of the 2nd International Symposium on Imprecise Probabilities and Their Applications (ISIPTA), pp. 69–78.
[12] Delande, E., Houssineau, J., 2018. A new representation of uncertainty for data fusion in SSA detection and tracking problems, in: 21st International Conference on Information Fusion (FUSION), pp. 1–8.
[13] Delande, E.D., Jones, B.A., Jah, M.K., 2023. Exploring an alternative approach to the assessment of collision risk. Journal of Guidance, Control, and Dynamics 46, 467–482.
[14] Destercke, S., Dubois, D., 2008. On non-additive uncertainty modelling with clouds. International Journal of Approximate Reasoning 49, 688–716.
[15] Destercke, S., Dubois, D., Chojnacki, E., 2009. Possibilistic information fusion using maximal coherent subsets. IEEE Transactions on Fuzzy Systems 17, 79–92.
[16] Dubois, D., Prade, H., 1988. Possibility Theory: An Approach to Computerized Processing of Uncertainty. Plenum Press.
[17] Dubois, D., Prade, H., 2000. Possibility theory in information fusion. International Journal of Intelligent Systems 15, 621–640.
[18] Dubois, D., Prade, H., 2005. Possibility theory and its applications: Where do we stand? Fuzzy Sets and Systems 156, 402–426.
[19] Jah, M.K., 2026a. The Epistemic Support-Point Filter: Jaynesian Maximum Entropy Meets Popperian Falsification. A Possibilistic Minimax-Entropy Optimality Proof. arXiv:2603.10065. https://doi.org/10.48550/arXiv.2603.10065
[20] Jah, M.K., Haslett, V., 2025. The Epistemic Support-Point Filter (ESPF): A Bounded Possibilistic Framework for Ordinal State Estimation. arXiv:2508.20806. https://doi.org/10.48550/arXiv.2508.20806
[21] Jah, M.K., 2026b. Theory of Epistemic Abductive Geometry (TEAG): A Unified Theory of Admissibility-Driven Inference Across Dynamical Systems, Measure Theory, and Language. Preprints.org. https://doi.org/10.20944/preprints202603.2010.v2
[22] Jah, M.K., 2026c. The Epistemic Support-Point Filter as a Tropical Hamilton–Jacobi System: Wavefront Propagation and Possibilistic Inference. Preprints.org. https://doi.org/10.20944/preprints202603.2110.v3
[23] Jah, M.K., 2026d. Probability is Earned: Information Capacity and the Epistemic Geometry of Inference. Preprints.org. https://doi.org/10.20944/preprints202603.2319.v1
[24] Jah, M.K., 2026e. The Geometry of Knowing: From Possibilistic Ignorance to Probabilistic Certainty—A Measure-Theoretic Framework for Epistemic Convergence. arXiv:2604.09614. https://doi.org/10.48550/arXiv.2604.09614
[25] Jaulin, L., Kieffer, M., Braems, I., Walter, E., 2001. Guaranteed non-linear estimation using constraint propagation on sets. International Journal of Control 74, 1772–1782.
[26] Jaynes, E.T., 2003. Probability Theory: The Logic of Science. Cambridge University Press.
[27] Julier, S.J., Uhlmann, J.K., 1997. A new extension of the Kalman filter to nonlinear systems, in: Proceedings of AeroSense: The 11th International Symposium on Aerospace/Defense Sensing, Simulation, and Controls, pp. 182–193.
[28] Kalman, R.E., 1982. System identification from noisy data, in: Proceedings of the IEEE, pp. 939–954.
[29] Moore, R.E., Kearfott, R.B., Cloud, M.J., 2009. Introduction to Interval Analysis. SIAM.
[30] Oberkampf, W., Helton, J., 2002. Investigation of evidence theory for engineering applications, in: 43rd AIAA/ASME/ASCE/AHS/ASC Structures, Structural Dynamics, and Materials Conference, p. 1569.
[31] Shafer, G., 1976. A Mathematical Theory of Evidence. Princeton University Press.
[32] Smolyak, S.A., 1963. Quadrature and interpolation formulas for tensor products of certain classes of functions. Doklady Akademii Nauk SSSR 148, 1042–1045.
[33] Wan, E.A., Van Der Merwe, R., 2000. The unscented Kalman filter for nonlinear estimation, in: Proceedings of the IEEE Adaptive Systems for Signal Processing, Communications, and Control Symposium, pp. 153–158.
[34] Zadeh, L.A., 1978. Fuzzy sets as a basis for a theory of possibility. Fuzzy Sets and Systems 1, 3–28.
[35] Zucchelli, E.M., Delande, E.D., Jones, B.A., Jah, M.K., 2021. Multi-fidelity orbit determination with systematic errors. The Journal of the Astronautical Sciences 68, 695–727.

Copyright © 2026 Advanced Maui Optical and Space Surveillance Technologies Conference (AMOS) – www.amostech.com