Polar Clustering and Stabilized Cluster Likelihoods for Extended Object Maritime Radar Tracking

Lucas Paes Moreira, Michailas Romanovas
Institute of Communications and Navigation
German Aerospace Center (DLR)
Neustrelitz, Germany
{lucas.paesmoreira, michailas.romanovas}@dlr.de

***Abstract*—Accurate maritime situational awareness is essential for safety, security, and environmental protection. While the Automatic Identification System (AIS) is widely used, it is vulnerable to jamming, spoofing, and intentional manipulation. These limitations can be mitigated through independent validation using marine X-band radars, where high-resolution coastal sensors detect extended targets with distance-dependent sparsity and raw data in polar coordinates.**

This paper evaluates a polar-domain clustering algorithm together with its impact on a family of extended-object multi-target tracking (MTT) frameworks of increasing complexity. A Random Matrix Model (RMM) based estimator incorporating radar intensity information is useed for a combined centroid and extension tracking. Beyond the front end, we detail the tracker mathematics required for robust operation on high-resolution clusters. The modifications include a stabilized cluster-likelihood formulation that avoids the clutter-power penalty of point-based extended-object models, a fully log-domain hypothesis-generation backbone shared by the IPDA, JIPDA, and PHD filters, and a measurement-driven birth (MDB) model that removes the one-frame latency of the classical Gaussian-mixture PHD.

Performance is validated on real-world X-band radar data within selected Regions of Interest (ROI) and at different distances from the sensor. Results show that clustering in the polar domain is more robust against the distance variability of detections and benefits the overall performance of the tracking algorithms.

***Index Terms*—polar clustering, extended objects, multiple target tracking, radar tracking, maritime surveillance**

I. Introduction

Maritime monitoring is essential for the safety, security, and efficiency of global trade, particularly in congested areas such as ports and harbors. The growing interest in autonomous maritime operations further increases the demand for reliable, real-time data, typically achieved through multi-sensor integration for navigation and route planning. Automatic Identification System (AIS) data, although widely used, is vulnerable to manipulation and interference caused by Global Navigation Satellite System (GNSS) jamming and spoofing [1]. These limitations highlight the need for complementary, independent sensing systems. The network of X-band coastal marine radars provides an established complementary, and in many locations, already existing alternative, enabling the detection of non-cooperative targets and ensuring robust monitoring even in GNSS-denied environments.

Recent advances in Multiple Target Tracking (MTT) have significantly improved maritime surveillance, particularly through Extended Object Tracking (EOT) methods. Unlike point-target models, EOT exploits the fact that modern high-resolution sensors produce multiple detections per object, enabling a reliable estimation of both the object's kinematic state and its spatial extension. The Random Matrix Model (RMM) [2], [3] is a well-established approach, representing target extent as an ellipse while capturing uncertainties in both kinematic and geometric parameters within a probabilistic framework and using well-established and well-understood mathematics for its recursive estimation. While random-finite-set (RFS) frameworks process the multiple measurements per target directly [4], clustering provides a deterministic, cost-effective way to group spatially consistent detections, treating each cluster as a single measurement and avoiding combinatorial partitioning. Within the RMM, the cluster centroid provides the target position, while the scatter matrix describes its extent and orientation, modeled by an Inverse Wishart (IW) distribution [3]. Fig. 1 illustrates a large vessel detected by an X-band radar, with its estimated position and extension ellipse representation.

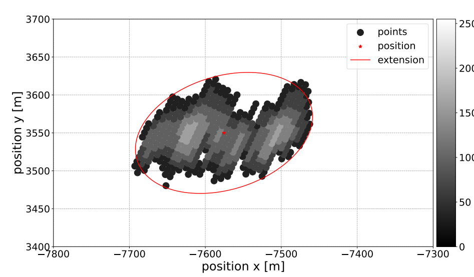

Fig. 1: Example for X-band radar detections of a large vessel with the intensity information and the resulting ellipse approximation. The sensor is located at the origin.

Marine X-band radar data are typically acquired in polar coordinates (range and azimuth) but are often transformed into Cartesian coordinates for further processing including MTT algorithms [5], [6]. This transformation introduces non-uniformities in measurement density due to the radar's angular resolution. Here the distant targets yield sparse point clouds, while closer ones produce denser measurements. Such variability complicates clustering over large observation ranges, and may also introduce distortions and inconsistencies.

This work makes two complementary contributions. First, we propose clustering directly in the polar domain, preserving the radar's native geometry for more consistent clustering across ranges. Second, and equally important for practice, we make explicit the tracker-side mathematics that high-resolution clusters demand but that textbook implementations rarely emphasize. We introduce a stabilized cluster likelihood, a log-domain hypothesis-generation backbone, and adopt measurement-driven birth (MDB) model of Ristic [7] for object proposing step. The resulting clusters are evaluated within several MTT frameworks of increasing complexity, namely the Sequential Probability Ratio Test (SPRT), (Joint) Integrated Probabilistic Data Association ((J)IPDA), and a cluster-based Gaussian-Inverse-Wishart Probability Hypothesis Density (CB-GIW-PHD) filter.

II. Methods

First let us consider the clustering problem. Raw radar data are inherently formed in polar domain (range r and azimuth ϕ). We denote the detection set at time k by $\mathcal{Z}_k = \{\mathbf{z}_k^i\}_{i=1}^{N_k}$, where N_k is the total number of detections, with $\mathbf{z}_k^i = [r_k^i, \phi_k^i]^T$. These point clouds may originate from true targets, Aids to Navigation (AtoN), static infrastructure, or clutter, and are grouped by a clustering algorithm. Consequently, each resulting cluster is processed as an independent source. Here the conventional approach maps detections to Cartesian coordinates before processing [8].

Several clustering algorithms are well established and can be integrated in MTT systems. Reference [9] investigated around ten different algorithms for automotive radars based on suitability and efficiency, where density-based DBSCAN shows better performance compared to others, which also shows reasonable results in X-band marine radar application [10]. In the Cartesian domain, however, the cluster density is range-dependent, turning more difficult and less efficient to tune a single set of hyper-parameter to cover a large spam on long-range sensors. Thus, we suggest to cluster the radar detections directly in the original polar domain, which is natural for rotational marine radars whose range and azimuth resolutions are approximately constant. Since the azimuth lives on S^1, all distances must respect its 2π-periodicity. Following [11], a polar point is mapped onto the lateral surface of a cylinder with a base radius R:

$$\mathbf{z}(r, \phi) \to \mathbf{z}(x, y, z) = \left[R\cos\phi, R\sin\phi, r\right]^T. \tag{1}$$

This acts as a scaling factor that weights azimuth differences relative to range. Unwrapping the surface onto a rectangular plane $(X', Y') = (R\phi, r)$ (Fig. 2) yields, for two points, the geodesic:

$$d(\mathbf{z}_k^i, \mathbf{z}_k^j) = \left[(r_k^i - r_k^j)^2 + R^2(\phi_k^i - \phi_k^j)^2\right]^{\frac{1}{2}}. \tag{2}$$

The two distinct distances $d(A, B_1)$ and $d(A, B_2)$ in Fig. 2 reflect the circularity of the azimuth and must be resolved by a signed angular difference. Setting $R = \sigma_r/\sigma_\phi$, with σ_r, σ_ϕ the range and azimuth resolutions, both restores metric (meter) units to (2) and removes R as a free hyper-parameter.

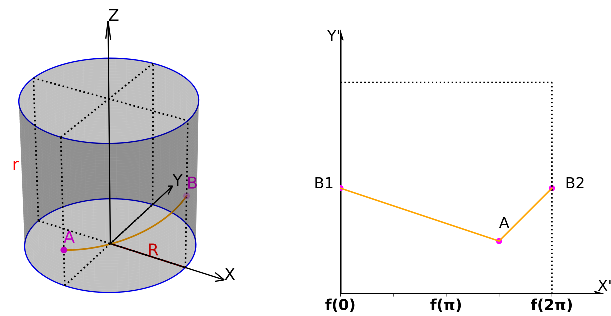

Fig. 2: Reconstruction of polar coordinates on the rectangular plane (X', Y') using cylindrical coordinates (adapted from [11]).

Our data are recorded in the ASTERIX CAT240 protocol [12], where each reflection is a 2D polar cell delimited by $(r_{\text{start}}, r_{\text{end}})$ and $(\phi_{\text{start}}, \phi_{\text{end}})$, so the resolutions are inferred per cell as $\sigma_r = r_{\text{end}} - r_{\text{start}}$ and $\sigma_\phi = \phi_{\text{end}} - \phi_{\text{start}}$. Note that these values do not necessarily reflect the physical sensor resolution and may vary in time, therefore are assessed per scan. The azimuth difference is evaluated with a four-quadrant $\arctan2$ to honor the S^1 topology.

Secondly, all presented trackers share the RMM [2] as the base kinematic and extend estimator, where the kinematic state $\mathbf{x}_k$ is Gaussian and the extent $\mathbf{X}_k$ follows an IW distribution, enabling joint estimation of 2D position, 2D velocity, and elliptical shape along with the associated uncertainties. In order to take into account the radar intensity information l_k^i, the intensity-weighted cluster statistics is used to form the centroid and scatter matrix measurements:

$$\bar{\mathbf{z}}_k^w = \frac{\sum_i l_k^i \mathbf{z}_k^i}{\sum_i l_k^i}, \quad \bar{\mathbf{Z}}_k^w = \frac{n_k}{\sum_i l_k^i} \sum_i l_k^i (\mathbf{z}_k^i - \bar{\mathbf{z}}_k^w)(\cdot)^T \tag{3}$$

where n_k is the number of detections in the cluster and $(\cdot)^T$ represents the transpose of the preceding term. This preserves the interpretation of as the measurement count while biasing the estimate toward strong reflections.

And finally, for hypothesis generation and data association we consider formulations of increasing complexity. The SPRT baseline [13] uses hard 1:1 assignments and updates existence from the probability of detection (P_D), the probability of false alarm (P_{FA}), and the gating probability (P_G) alone with its simplicity benefiting mainly low-clutter and scenes with well-separated objects. JPDA [14], [15] resolves ambiguous associations by evaluating all feasible joint events χ and forming marginal weights $\beta_j^t = \sum_\chi p(\chi|Z^k)\, \hat{\omega}_{jt}(\chi)$, $\beta_0^t = 1 - \sum_j \beta_j^t$. The JIPDA extension [16] additionally propagates a probabilistic existence variable. We also include a multi-instance IPDA [17] and a cluster-based PHD tracker [18] whose posterior is a mixture of GIW components fed with pre-clustered measurements. Here the clustering allows avoiding the combinatorial partitioning of classical EOT-PHD filters. In all the trackers the target motion uses a 2D constant-velocity model.

Every framework except SPRT consumes a spatial likelihood. The key departure from textbook EOT is how this likelihood is formed for clusters. A point-based partition likelihood treats the n_k members of a cluster as individual returns, so the clutter hypothesis contributes $\lambda_c^{n_k}$. For high-resolution radar this term over-penalizes and eventually deletes large clusters as n_k grows. Treating each cluster Z_j as a single observation with centroid $\bar{\mathbf{z}}_j$ and scatter $\bar{\mathbf{Z}}_j$ removes the exponent on λ_c. A naive single-observation Bayes factor $p(\bar{\mathbf{z}}_j|H_1)\, p(\bar{\mathbf{Z}}_j|H_1)/\lambda_c$ avoids the $\lambda_c^{n_k}$ penalty but introduces a dimensionality mismatch as the shape term follows a Wishart density that becomes extremely peaky as n_k grows. Thus the log-likelihood gets dominated by the shape magnitude and target hypotheses appear infinitesimal next to missed detections.

We therefore characterize clutter with a non-informative extent model $p(\bar{\mathbf{Z}}_j|H_0)$ and define the stabilized cluster likelihood that is passed to every hypothesis generator with:

$$f_t(Z_j \mid Z^{k-1}) = \frac{p(\bar{\mathbf{z}}_j|H_1)\, p(\bar{\mathbf{Z}}_j|H_1)}{p(\bar{\mathbf{Z}}_j|H_0)}, \tag{4}$$

where the centroid term is Gaussian and the shape terms are Wishart densities with $n_k - 1$ degrees of freedom:

$$p(\bar{\mathbf{Z}}_j|H_i) = \mathcal{W}_d(\bar{\mathbf{Z}}_j; n_k - 1, \mathbf{V}_i), \tag{5}$$

with $\mathbf{V}_1 = \rho\mathbf{X}_{k|k-1} + \mathbf{R}_k$ being the predicted target extent and $\mathbf{V}_0 = \sigma_c^2\mathbf{I}$ corresponds to a rather broad clutter model. Here with equal degrees of freedom the data-dependent factor $|\bar{\mathbf{Z}}_j|^{(n_k-d-2)/2}$, the $2^{(n_k-1)d/2}$ factor, and the multivariate gamma $\Gamma_d\big(\frac{n_k-1}{2}\big)$ all cancel between numerator and denominator, leaving the closed form as:

$$\frac{p(\bar{\mathbf{Z}}_j|H_1)}{p(\bar{\mathbf{Z}}_j|H_0)} = \left(\frac{|\mathbf{V}_0|}{|\mathbf{V}_1|}\right)^{\frac{n_k-1}{2}} \exp\big(-\tfrac{1}{2}\,\mathrm{tr}\big[(\mathbf{V}_1^{-1} - \mathbf{V}_0^{-1})\bar{\mathbf{Z}}_j\big]\big)\,. \tag{6}$$

The divergent n_k-dependent normalizers are therefore gone. Now n_k enters only through the bounded determinant-ratio exponent, keeping log-likelihoods comparable in magnitude to the existence/clutter terms.

This construction is internally consistent with the standard filtering recursions. Adopting the cluster measurement model in which a target generates a cluster with likelihood $g(Z_j|\mathbf{x}) = p(\bar{\mathbf{z}}_j|H_1)\, p(\bar{\mathbf{Z}}_j|H_1)$ and a clutter cluster has intensity $\kappa(Z_j) = \lambda_c\, p(\bar{\mathbf{Z}}_j|H_0)$, the PHD detection-case weight is:

$$w_D^{(t,j)} = \frac{P_D^t\, w_{k|k-1}^t\, g_t(Z_j)}{\lambda_c\, p(\bar{\mathbf{Z}}_j|H_0) + \sum_\ell P_D^\ell\, w_{k|k-1}^\ell\, g_\ell(Z_j)}. \tag{7}$$

Dividing numerator and denominator by $p(\bar{\mathbf{Z}}_j|H_0)$ leaves the normalized weights unchanged and recovers exactly the standard form with $g \to f$,

$$\begin{aligned} w_{ND}^t &= (1 - P_D^t P_G)\, w_{k|k-1}^t, \\ w_D^{(t,j)} &= \frac{P_D^t\, w_{k|k-1}^t\, f_t(Z_j)}{\lambda_c + \sum_\ell P_D^\ell\, w_{k|k-1}^\ell\, f_\ell(Z_j)}, \end{aligned} \tag{8}$$

Hence (4) is the exact spatial likelihood of a consistent cluster model, in which λ_c appears once (not raised to n_k) and the clutter shape $p(\cdot|H_0)$ is the per-cluster re-scaling. The same f_t is supplied to the (J)IPDA hypothesis generators, where classical point-target equations are recovered by dropping the Wishart factors. Finally, we decouple estimation step from the gating step. Now we have the innovation covariance for the state update to scale with cluster size,

$$\mathbf{S}_{k|k-1} = \mathbf{H}\mathbf{P}_{k|k-1}\mathbf{H}^T + \frac{\rho\mathbf{X}_{k|k-1} + \mathbf{R}_k}{n_k}, \tag{9}$$

while the validation gate uses the full extent, $\mathbf{S}_{k,\text{gate}} = \mathbf{H}\mathbf{P}_{k|k-1}\mathbf{H}^T + \rho\mathbf{X}_{k|k-1} + \mathbf{R}_k$, where $\mathbf{H}\mathbf{P}_{k|k-1}\mathbf{H}^T$ is the kinematic (measurement) uncertainty, $\mathbf{X}_{k|k-1}$ is the predicted object extension (from the Inverse Wishart distribution), $\mathbf{R}_k$ is the sensor noise covariance and ρ is the scaling parameter from [2].

What relates practical implementation, direct evaluation of the association weights underflows whenever likelihoods are small (Gaussian/GIW tails), the clutter density is large, or many measurements populate a gate. All hypothesis generators therefore share a single log-domain backbone built on the log-sum-exp identity:

$$\mathrm{LSE}(x_1, \ldots, x_n) = M + \log \sum_{i=1}^{n} e^{\,x_i - M}, \tag{10}$$

with $M = \max_i x_i$. Two implementation details matter in practice and are easy to miss: gated-out entries are encoded as $-\infty$ and vanish under (10) without producing `NaN` provided at least one finite argument is present, and the clutter log-term $\log \lambda_c$ provides exactly this finite anchor, so the argument set is never empty or all $-\infty$.

GM-PHD. For each predicted component t and cluster Z_j, we form the log-numerator and the PHD pseudo-likelihood denominator, respectively:

$$\log L_{t,j} = \log\big(P_D^t\, w_{k|k-1}^t\big) + \log f_t(Z_j), \tag{11}$$

$$\log D_j = \mathrm{LSE}\big(\log \lambda_c,\ \{\log L_{\ell,j}\}_{\ell \in \mathcal{A}_j}\big), \tag{12}$$

with $\mathcal{A}_j$ the set of components gating measurement j, so that $w_D^{(t,j)} = \exp(\log L_{t,j} - \log D_j)$ reproduces (8) without ever leaving the log-domain.

(J)IPDA. Defining the per-measurement log evidence ratio $\ell_j = \log f_j - \log P_G - \log \lambda_c$, the IPDA normalizer is assembled by a nested (10):

$$\log(1-\delta_k) = \mathrm{LSE}\Big(\log(1-P_D P_G),\ \log(P_D P_G) + \mathrm{LSE}_j\ \ell_j\Big) \tag{13}$$

yielding $\beta_0 = \exp[\log(1 - P_D P_G) - \log(1 - \delta_k)]$ and $\beta_j = \exp[\log(P_D P_G) + \ell_j - \log(1 - \delta_k)]$. The existence update is evaluated in the stable form $P_E^+ = P_E \,/\, [\, P_E + (1 - P_E)\, e^{-\log(1-\delta_k)}\,]$, which avoids the cancellation in $1 - \delta_k P_E$. For JIPDA the log-probability of a feasible event χ becomes the track-wise sum:

$$\mathcal{L}_t(\chi) = \begin{cases} \log(P_D^t P_E^t) + \log f_t(z_{j[t]}) - \log \lambda_c, & \delta_t(\chi) = 1 \\ \log\big(\max(1 - P_D^t P_G P_E^t,\ \epsilon)\big), & \delta_t(\chi) = 0 \end{cases} \tag{14}$$

and the events are normalized by $\log c = \mathrm{LSE}_\chi\left(\sum_t \mathcal{L}_t(\chi)\right)$ before marginalization. The ϵ-floor on the missed-detection factor prevents $-\infty$ for near-perfect sensors. Across all three filters the spatial input is the same stabilized factor f_t of (4) with the only difference in the combinatorial structure of the normalizer.

The classical GM-PHD [19] injects birth components in the prediction step at fixed locations, so a new target is first updated only at the follow-up scan, i.e. with a one-frame latency, and the tracker is effectively blind outside the predefined birth region. The PDA family (including SPRT as a degenerate case), by contrast, initiates tracks from unassociated measurements within the same frame. To match this responsiveness of PDA family in PHD trackers, we adopt the MDB model of Ristic et al. [7], which augments the state with a label $\beta \in \{0, 1\}$ (persistent/newborn) and shows that newborns are always detected, $P_D(\mathbf{x}) = 1$ for $\beta = 1$. The predicted intensity then separates as $D_{k|k-1}(\cdot, 1) = \gamma_{k|k-1}$ for newborns and the persistent recursion for $\beta = 0$, while the corresponding update denominators acquire the extra term $\langle g_k(z|\cdot), \gamma_{k|k-1}\rangle$.

Interesting is that the implementation of Ristic MDB does not need modifications the PHD update. For every cluster we propose one birth component via the inverse measurement model, assign it weight $w_\gamma = w_b/m_k$ so that the total birth mass per scan equals a constant w_b (the expected number of births), append it to the persistent set with $P_D = 1$, and run the unmodified generator over the augmented set. Setting $P_D = 1$ has two exact consequences. First, the missed-detection weight $w_{ND} = (1 - P_D P_G)w$ collapses to 0 for newborns, so no spurious missed-detection copy is propagated. Second, the newborn contribution enters the shared denominator (12) precisely as the $\langle g, \gamma\rangle$ term of [7]. The newborn is thus associated with its generating measurement and reported in the same frame, eliminating the latency while letting the standard PHD arithmetic perform the bookkeeping. The price is a larger pool of tentative tracks to be processed. Note that MDB makes the birth process depend on the current scan and is therefore not strictly Markov. However, we may accept it as a viable and flexible practice since it allows us to recover e.g. lost tracks anywhere in the field of view rather than only at predefined entry regions and does not need fundamental changes in the underlying PHD mathematics.

Because the PHD update splits each of T predicted components into $1 + M$ children (one missed-detection and up to M detection cases), the posterior intensity must be reduced at every scan. Following the reference architecture, pruning step shall precede the merging. Components below a weight threshold are first removed and weights are deliberately not re-normalized after pruning or capping, so the intensity continues to integrate to the expected target number.

Merging then reduces the surviving GIW components by greedy, single-pass moment matching [20]. The highest-weight component is taken as a pivot, and components i are absorbed when their equivalent innovation $\mathbf{S}_{\text{equiv}} = \mathbf{P}_j + \mathbf{P}_i + \rho(\mathbf{X}_j + \mathbf{X}_i)$ (full kinematic covariance plus the positional extent block) satisfies $(\mathbf{m}_i - \mathbf{m}_j)^T \mathbf{S}_{\text{equiv}}^{-1}(\mathbf{m}_i - \mathbf{m}_j) \leq \gamma$. The merged track sums the weights and applies weighted moment reduction independently to the Gaussian and the IW components. The single pass (merged tracks are not re-inserted) intentionally avoids the coalescence of closely spaced targets.

Differently from the PDA family, the identity handling in PHD tracker requires a dedicated step, since the original algorithm does not natively carry labels. The detection children inherit the tag of their parent component, so after the merging step several spatially separated components may share a tag, representing competing, mutually exclusive hypotheses for one target. These hypothesis are not moment-merged, but instead, consistent with [19], are grouped by tag and only the highest-weight branch retains the tag while the rest are pruned, which we found to outperform re-labeling the residual branches on the ID-switch metric. To prevent the freshly proposed tracks from cannibalizing the established ones (as proposed tracks almost perfectly fit the measurements), the merger pulls the surviving tag from the highest-weight persistent component and only falls back to a newborn when no persistent component is present in the merge set. A final computational cap retains the $J_{\max}$ heaviest components. This ceiling is a CPU/memory bound, distinct from and much larger than the plausible target count, and does not alter the cardinality estimate. Stable identities are essential in maritime use, where AIS fusion and consistency checks rely on persistent track IDs [21].

Detection and clutter are assumed constant within the ROI, with $P_D = 0.8$ and $\lambda_c = 10^{-8}$. The SPRT tracker uses $P_{FA} = 0.01$, and the PHD filter a survival probability $P_S = 0.95$. For the PDA family, tracks are initiated from unassociated clusters with existence probability 0.6, whereas in PHD every measurement creates a track [7] with small combined birth weight, distributed over all born targets, e.g. 0.1. Positions are initialized from measurements with zero velocity and inflated covariance. The initial extent is computed from the steady-state approximation:

$$\mathbf{X}_{\infty,k} = \frac{1}{\rho}\left(\frac{\bar{\mathbf{Z}}_k}{n_k} - \mathbf{R}_k\right). \tag{15}$$

The expression above is not guaranteed to yield a positive-definite matrix (required by the RMM update), thus the difference matrix $\bar{\mathbf{Z}}_k/n_k - \mathbf{R}_k$ may be repaired per axis via an eigenvalue floor, here with a minimum of 0.01 m in both major axis, which is negligible for actual target (vessel) dimensions and avoid degeneration of the extension estimation. Data association uses greedy $M{:}N$ assignment for JIPDA, 1:N for IPDA, and 1:1 for SPRT. The PDA family prunes on existence probability, kinematic uncertainty, missed detections, and extent uncertainty, whereas PHD relies solely on the weight-based pruning and merging.

III. Results

The evaluation used a real-world marine radar dataset recorded by a static station in the Baltic Sea, detecting reflections up to 30 km. Two scenarios were selected: a near-field case (up to 5 km) and a far-field case (beyond 20 km).

The set was manually labeled over one hour, annotating object pixels for cooperative vessels (AIS-confirmed), non-cooperative vessels, and AtoN devices such as buoys and these labels serve as ground truth. Two subsets, one per ROI, enable the comparison across the scenarios. The near-field subset has AtoN objects, frequently occluded by large vessels, and one non-cooperative target not transmitting AIS messages. The far-field subset has mainly cooperative vessels without AtoN.

DBSCAN is sensitive to ε and to the minimum number of points ($min\_samples$), so a grid search was performed over these parameters to find the optimal values for each domain and ROI. Considering the range resolution $dr \approx 6.9$ meters, the grid search considered values in the interval $10 \leqslant \varepsilon \leqslant 70$ and $3 \leqslant min\_samples \leqslant 15$. The outputs were scored by the Adjusted Rand Index (ARI) [22] and the difference between labeled and estimated cluster counts, with the optimum parameter set chosen for high ARI and low difference count. Fig. 3 shows the these metrics for each ROI per ε values. As the results show no significant differences for $min\_samples > 7$, only results for $min\_samples \in \{3, 5, 7\}$ are shown.

In the Cartesian domain, the ARI score shows high dependency on ε and has no single set of parameters with the highest scores. In polar domain, conversely, ε values above 45.0 for any values of $min\_samples$ show the highest ARI score of 0.995.

Trackers were evaluated with GOSPA [23] and T-GOSPA [24], the latter providing the ID-switch component not directly available from GOSPA. Both used $p = 2$, $c = 10$ (and $\gamma = 5$ for T-GOSPA). GOSPA components were computed per frame, while T-GOSPA returns scalar summaries for scenarios. Fig. 4 shows the per-frame GOSPA distance for the near- and far-field scenarios for all trackers, while the Table I lists the T-GOSPA values for near and far field scenarios.

TABLE I: T-GOSPA

ROI	Tracker	Distance	Localization	Missed	False	ID Switch
Near-field	SPRT	162.5	**61.6**	**70.0**	123.3	50.2
	P-IPDA	172.3	73.6	77.1	123.9	54.5
	P-JIPDA	172.3	73.6	77.1	123.9	54.5
	CB-GIW-PHD	**136.4**	86.6	76.2	**68.2**	**25.5**
Far-field	SPRT	**266.8**	**143.5**	112.0	194.0	19.4
	P-IPDA	270.3	144.3	113.6	197.5	18.7
	P-JIPDA	270.3	144.3	113.6	197.5	18.7
	CB-GIW-PHD	268.7	168.1	**97.2**	**185.6**	**5.0**

The overall GOSPA distance (Fig. 4) shows a better performance for all trackers in the near-field scenario, as well as a relatively stable performance over time. In the near-field the clusters are well separated with high detection density, so the localization and missed tracks (false negatives) errors are relatively low for all trackers, compared to the far-field counterpart (Table I). The higher values for false positives are an artifact of frequent object occlusion, usually AtoN occluded by large vessels, combined with the use of labels as reference. During an occlusion no measurement and hence no label (ground truth) is available, yet the tracks persist, as intended, for a limited time, producing false positives. Metrics, therefore, penalize the desired behavior of maintaining occluded tracks. On the other hand, using AIS positions as reference would distort the metrics for non-cooperative vessels and AtoN.

In the far-field ROI, targets generate less and more sparse detections. Additionally, due to the long distance, target clusters are often misdetected. When the track is persistent, its state vector is not updated from measurements, increasing their uncertainties, consequently also increasing the localization and false positive errors. However, as these misdetections are shorter in time compared to the occlusion of AtoN in the first scenario, the observed ID switch penalty is less prominent in the far-field.

In general all trackers present similar performance for each metric, with the PHD tracker performing slightly better especially for ID preservation. This is consistent with its hypothesis logic, in which essentially all track-measurement combinations are produced, whereas the PDA family merges every hypothesis back at the moment-reduction step. The PHD merger is gate-conditioned and may leave strongly mismatched hypotheses unmerged, recovering split objects. Under linear-Gaussian dynamics with state-independent detection/survival and without explicit birth/spawning, the PHD recursion and the J(I)PDA marginalization coincide.

IV. Summary

This work demonstrates the use of the polar clustering for radar data in MTT applications using real-world maritime X-band radar measurements. The suggested method was evaluated on the performance of different extended-object MTT algorithms. The proposed clustering method tackles the issues inherent to the polar data when converted to Cartesian domain, especially the variance of point cloud spatial density w.r.t. the distance to the sensor, impacting the most sensitive hyper-parameter in density-based clustering algorithms such as DBSCAN.

On the MTT side, the work demonstrates that EOT operating on measurements clusters requires a stabilized extension clutter factor to remove the clutter-power penalty. Furthermore we demonstrated that the IPDA, JIPDA, and PHD can be formulated to guarantee numerical stability with a single log-sum-exp backbone. Finally, the MDB model assigning $P_D = 1$ to newborns reproduces the adaptive-birth recursion within an unmodified PHD update and removes the one-frame latency of the classical PHD filter.

The methods were assessed across frameworks of increasing complexity on real X-band data. Their performance corroborates that the leap in complexity from SPRT to CB-GIW-PHD does not translate into a proportional performance gain on maritime radar data. Future work includes comparison against further clustering algorithms, more advanced MTT such as the Poisson Multi-Bernoulli Mixture (PMBM), a deeper study of maritime-specific parametrization, and the fusion of radar and AIS data for improved state estimation and identity preservation.

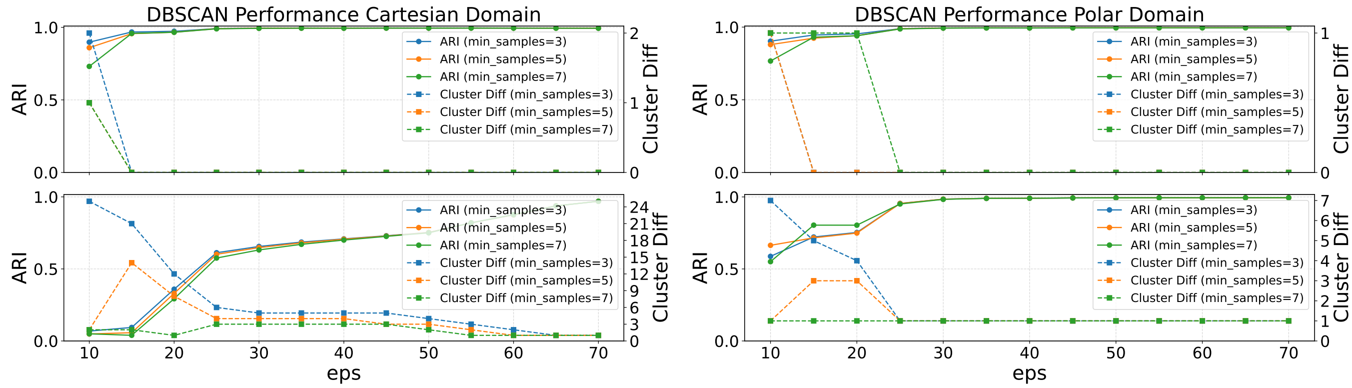

Fig. 3: ARI and cluster difference per ε for Cartesian (left) and polar (right). The plots on the top refer to near-field ROI and the bottom plots for the far-field.

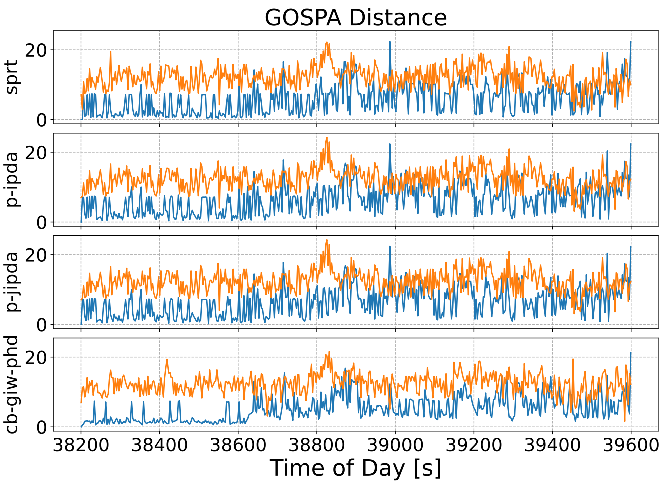

Fig. 4: Overall GOSPA distance per frame on both ROI: near-field (blue) and far-field (orange).

References

[1] J. Han, S. Y. Kim, and J. Kim, "Enhanced target ship tracking with geometric parameter estimation for unmanned surface vehicles," *IEEE Access*, vol. 9, pp. 39 864–39 872, 2021.

[2] M. Feldmann, D. Fränken, and W. Koch, "Tracking of extended objects and group targets using random matrices," *IEEE Transactions on Signal Processing*, vol. 59, no. 4, pp. 1409–1420, 2011.

[3] J. W. Koch, "Bayesian approach to extended object and cluster tracking using random matrices," *IEEE Transactions on Aerospace and Electronic Systems*, vol. 44, no. 3, pp. 1042–1059, 2008.

[4] K. Granström, M. Baum, and S. Reuter, "Extended object tracking: Introduction, overview and applications," *Journal of Advances in Information Fusion*, vol. 12, no. 2, pp. 139–174, 2017.

[5] K. Granström, A. Natale, P. Braca, G. Ludeno, and F. Serafino, "Gamma gaussian inverse wishart probability hypothesis density for extended target tracking using x-band marine radar data," *IEEE Transactions on Geoscience and Remote Sensing*, vol. 53, no. 12, pp. 6617–6631, 2015.

[6] G. Vivone, P. Braca, K. Granström, A. Natale, and J. Chanussot, "Converted measurements bayesian extended target tracking applied to x-band marine radar data," *Journal of Advances in Information Fusion*, vol. 12, no. 2, pp. 189–210, 2017.

[7] B. Ristic, D. Clark, B.-N. Vo, and B.-T. Vo, "Adaptive target birth intensity for phd and cphd filters," *IEEE Transactions on Aerospace and Electronic Systems*, vol. 48, no. 2, pp. 1656–1668, 2012.

[8] S. V. Bordonaro, P. Willett, and Y. Bar-Shalom, "Unbiased tracking with converted measurements," in *2012 IEEE Radar Conference*, 2012, pp. 0741–0745.

[9] D. C. de Ramos, L. R. Ferreira, M. M. D. Santos, E. L. S. Teixeira, L. R. Yoshioka, J. F. Justo, and A. W. Malik, "Evaluation of cluster algorithms for radar-based object recognition in autonomous and assisted driving," *Sensors*, vol. 24, no. 22, 2024. [Online]. Available: https://www.mdpi.com/1424-8220/24/22/7219

[10] J. S. Fowdur, M. Baum, F. Heymann, and P. Banyś, "An overview of the pakf-jpda approach for elliptical multiple extended target tracking using high-resolution marine radar data," *Remote Sens.*, vol. 15, no. 10, p. 2503, 2023.

[11] X. Sun and P. Sajda, "Circular clustering with polar coordinate reconstruction," *IEEE/ACM Transactions on Computational Biology and Bioinformatics*, vol. 21, no. 5, pp. 1591–1600, 2024.

[12] *EUROCONTROL Specification for Surveillance Data Exchange - ASTERIX Category 240: Radar Video Transmission*, EUROCONTROL Std. 978-2-87 497-028-3, Rev. 1.3, 5 2015.

[13] S. Blackman and R. Popoli, *Design and Analysis of Modern Tracking Systems*. Artech House, 1999.

[14] T. Fortmann, Y. bar shalom, and M. Scheffe, "Sonar tracking of multiple targets using joint probabilistic data association." *IEEE Journal of Oceanic Engineering*, vol. 8, pp. 173 – 184, 08 1983.

[15] Y. Bar-Shalom, F. Daum, and J. Huang, "The probabilistic data association filter," *IEEE Control Systems Magazine*, vol. 29, no. 6, pp. 82–100, 2009.

[16] D. Musicki and R. Evans, "Joint integrated probabilistic data association - jipda," in *Proceedings of the Fifth International Conference on Information Fusion. FUSION 2002. (IEEE Cat.No.02EX5997)*, vol. 2, 2002, pp. 1120–1125 vol.2.

[17] D. Musicki, R. Evans, and S. Stankovic, "Integrated probabilistic data association," *IEEE Transactions on Automatic Control*, vol. 39, no. 6, pp. 1237–1241, 1994.

[18] K. Granstrom and U. Orguner, "A phd filter for tracking multiple extended targets using random matrices," *IEEE Transactions on Signal Processing*, vol. 60, no. 11, pp. 5657–5671, 2012.

[19] K. Panta, D. E. Clark, and B.-N. Vo, "Data association and track management for the gaussian mixture probability hypothesis density filter," *IEEE Transactions on Aerospace and Electronic Systems*, vol. 45, no. 3, pp. 1003–1016, 2009.

[20] K. Granström and U. Orguner, "On the reduction of gaussian inverse wishart mixtures," in *2012 15th International Conference on Information Fusion*, 2012, pp. 2162–2169.

[21] A. G. Hem, M. Baerveldt, and E. F. Brekke, "Pmbm filtering with fusion of target-provided and exteroceptive measurements: Applications to maritime point and extended object tracking," *IEEE Access*, vol. 12, pp. 55 404–55 423, 2024.

[22] J. M. Santos, , and M. Embrechts, "On the use of the adjusted rand index as a metric for evaluating supervised classification," in *Artificial Neural Networks – ICANN 2009*, vol. 2, 2009, pp. 175–184.

[23] A. S. Rahmathullah, A. F. Garcia-Fernandez, and L. Svensson, "Generalized optimal sub-pattern assignment metric," in *2017 20th International Conference on Information Fusion (Fusion)*, 2017, pp. 1–8.

[24] A. F. Garcia-Fernandez, J. Gu, L. Svensson, Y. Xia, J. Krejčí, O. Kost, and O. Straka, "Gospa and t-gospa quasi-metrics for evaluation of multi-object tracking algorithms," 2025. [Online]. Available: https://arxiv.org/abs/2507.13706