# Cooperative Remote Sensing of Ice using a Spatially Indexed Labeled Multi-Bernoulli Filter

Jonatan Olofsson, Edmund Brekke and Tor Arne Johansen
Center for Autonomous Marine Operations and Systems (NTNU-AMOS)
Department of Engineering Cybernetics,
Norwegian University of Science and Technology, Trondheim, Norway.
Corresponding author: jonatan.olofsson@ntnu.no

***Abstract*—In polar region operations, drift ice positioning and tracking is useful for both scientific and safety reasons. Many sensors can be employed to generate detections of sea ice, such as satellite-carried Synthetic Aperture Radar (SAR) and, recently, imagery equipment carried by Unmanned Aerial Systems (UAS). Satellite-carried SAR has the advantage of being able to cover large areas and provide consistent imagery largely independent of weather, albeit at a relatively coarse resolution. Using UAS, the resolution and precision of the tracking can be locally improved.**

**To track the large amount of individual objects present in an area as large as the Arctic, it is necessary to efficiently select and exclusively work with the objects in the relevant field-of-view. In this paper, a Spatially Indexed Labeled Multi-Bernoulli filter is presented and applied to a tracking problem representing a mission setup for field-tests due this year. In the setup, satellite and UAS imagery is combined to provide real-time Multi-Target Tracking of sea ice objects. A brief introduction is given to the implementation of the proposed Spatially Indexed Labeled Multi-Bernoulli filter, which is made available under an Open Source license.**

## I. Introduction

Work in the polar regions of our planet is unavoidably linked with hazards such as drift ice. Increased presence fueled by economic interests in the Arctic, have for several decades [19] called for research in the field of Ice Management. A comprehensive overview of Ice Management in practical use is provided in [5]. The field deals with the forecast, detection and tracking of ice, but also the physical actions taken to avoid collisions [5]. While managing ice is of great importance to polar ventures, predicting ice movement has proven difficult [6], concluding that observations are essential for tracking.

Ice Management is applicable not only for stationary installations, but has also been studied for use in the protection of ship routes in the Arctic [9]. A moving object of protection may be not be able to rely solely on satellite Synthetic Aperture Radar (SAR) imagery, as these generally have limited availability and coverage, with sample times in the order of days. Recent studies have introduced the use of Unmanned Aerial Systems, UAS's, as a supplement to local sensors such as ship- or stationary radar [10, 8, 14].

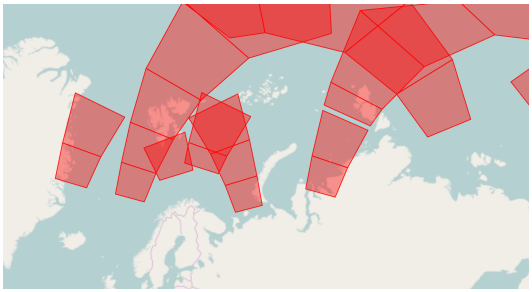

Figure 1: Sample Sentinel-1 satellite observation areas in the Arctic.

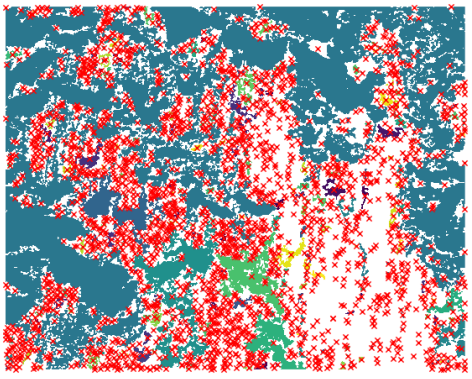

Figure 2: Individual ice detections, even as here from a small subset of a Synthetic Aperture Radar (SAR) imagery, can be overwhelmingly numerous.

In the Arctic, objects are tracked over a geographically vast area. Each observation covers only a relatively limited area — even for satellite observations as seen in Figures 1 and 2 — but can contain a large number of sea ice objects worth tracking. Since each observation's field-of-view is unique, and a majority of the Arctic can be quickly discarded as being outside the observed area, it

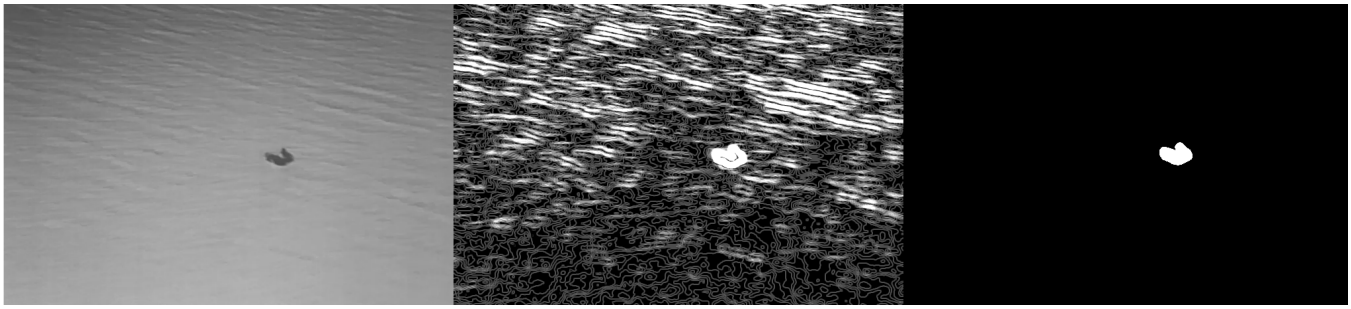

Figure 3: Detections of ice can be extracted from infrared imagery from UAS-carried cameras. Left: Raw thermal image; Middle: Enhanced features; Right: Segmented detected sea ice [14].

is desirable to dynamically limit any estimate correction to the relevant area.

To enhance the observation detail and frequency, we seek to fuse the ice tracking data available from a satellite constellation, with that of imagery — exemplified in Figure 3 — obtained from one or more UAS's through wireless datalinks. This challenge can be expanded to e.g. that of optimal UAS path-planning, decision support etc., but the focus of this paper is on the fusion of large, widely spatially distributed, real-time data from multiple data-sources - SAR and UAS imagery in particular.

After introducing the terminology of the paper in Section II, clustering and spatial indexing is introduced in Section III as it allows for the extraction of partial filters suitable for the limited — as compared to the area of interest — field-of-views provided by e.g. UAS imagery. This allows the full filter to track sea ice over the entire Arctic, while each update may focus on the area within the specific sensor field-of-view.

Furthermore, as both ice motion prediction (sometimes over large time intervals) and measurements are subject to uncertainties, we elect a Bayesian, hypothesis based, multi-target tracking technique in the Labeled Multi-Bernoulli (LMB) filter, introduced in Section IV.

In Section V, we present the paper's main contribution of an Open Source implementation of the Spatially Indexed Labeled Multi-Bernoulli filter, which was developed for this paper and for the use in field-trials due in Svalbard 2017. A scenario demonstrating the fusion of data between satellite and multiple UAS agents, representative for these trials, is presented in Section VI along with simulation results, followed by concluding words in Section VII.

## II. Terminology and Notation

The terminology often employed in common Multi-Target Tracking (MTT) literature, and consequently here, is based on the following definitions;

At any given time instance, a sensor delivers a *scan* — an unordered but enumerated set of reports (measurements) detected at this instance. This set is exhaustive, i.e. the scan set $Z_k$ contains all reports (indexed by $i$), $z_{k,i}$ for time $k$. In the following, time indices are dropped unless needed for clarity.

The purpose of MTT is to estimate a list of tracked *targets*. Each true target follows a "true track", the estimate of which is the result of filtering of data and assignments to the target. These, hypothetical, assignments can indicate association of the target to a specific report, but also a missed detection or even alternative motion models [13]. It is assumed that each true target will yield at most one report in each scan. Each track estimate starts with the event of a newly created target as hypothesized by a target birth model.

The LMB filter follows the common filter predict/correct pattern shown in Figure 4. In the filter loop, the prediction provides an estimate of what inputs are expected during the correction step. As scans become available, each sensor updates the estimate independently, in order of availability.

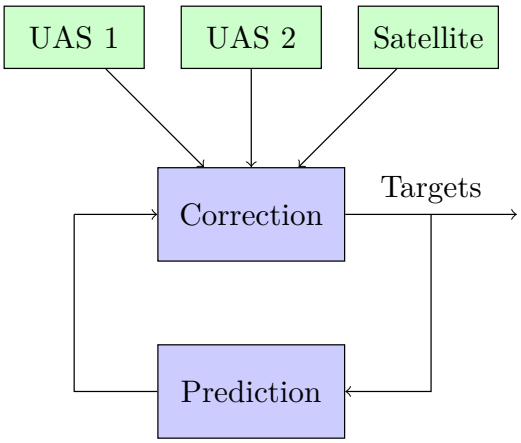


Figure 4: Filter overview. Each timestep in the filter corresponds to one loop, and each loop results in a list of target positions estimated to be present at that time. The sensors displayed here mirrors that of the setup for sea ice tracking described in Section VI.

To accomodate the uncertainty in which report belongs to which target, the LMB filter considers multiple *hypotheses*. The definition for an hypothesis employed in this paper, is that of a set of hypothetical assignments which assigns each previously hypothesized target as either 1) existing and associated with a report in the

incoming scan, 2) existing but missed in the scan ($z_\emptyset$ association), or 3) non-existent (indicated in equations by the letter $F$).

Each hypothesized target is assigned a label at its creation. The space of hypothetical associations (in the sense of 1) and 2) above) — denoted $\Theta_I$ — is parametrized by the set $I$ of labeled targets assumed to be existent. All assignments in $\Theta_I$ assign all targets in $I$ to either of above state 1 or 2. Thus, given a set $\mathcal{L}$ of labeled targets, $I \cap \mathcal{L}$ are considered existing, and $\mathcal{L} - I$ non-existing.

---

*Example:* Given the set of tracked labeled targets $\mathcal{L} = \{\ell_1, \ell_2, \ell_3\}$, hypothesis $n$ at time $k$ assigns target with label $\ell_1$ to report $z_1$, presumes that $\ell_2$ has been missed, and that $\ell_3$ does not exist:

$$\begin{aligned} \theta_{k,n} &= \{(\ell_1,\, z_{k,1}),\quad (\ell_2,\, z_\emptyset)\}, \\ I &= \{\ell_1, \ell_2\}, \\ \mathcal{L} - I &= \{\ell_3\}. \end{aligned}$$

---

The space of hypothetical associations is, for each timestep, combinatorially large. It can, however, be enumerated in order of probability and truncated to the $k$-best hypotheses. For each hypothesis, it is assumed that each track can be modeled as a standard single-target estimation problem, conditioned on the validity of the hypothetical assignments, and that its evolution is described by the single-target transition density $f_k(x_{k,\ell}|x_{k-1,\ell}, \ell)$.

Targets which, as decided through a probability gating procedure [2, 23], no hypothesis will associate to the same reports can be treated in separate *clusters* — independent groups of targets which share ambiguous reports [20].

*Notation:* In the paper, notation common to the LMB literature is used [21], including the inner product,

$$\langle f, g \rangle \triangleq \int f(x)\, g(x)\, dx, \tag{1}$$

as well as the multi-object exponential notation

$$h^X \triangleq \prod_{x \in X} h(x) \tag{2}$$

($h^\emptyset = 1$ by convention).

Further, the Kronecker delta function is defined by

$$\delta_Y(X) \triangleq \begin{cases} 1, & \text{if } X = Y \\ 0, & \text{otherwise} \end{cases}, \tag{3}$$

and is used to select summands relevant to exactly the set $Y$. Similarly, the inclusion function,

$$1_Y(X) \triangleq \begin{cases} 1, & \text{if } X \subseteq Y \\ 0, & \text{otherwise} \end{cases}, \tag{4}$$

is used to select summands relevant to a set $Y$ of which $X$ is a subset. If $X$ is singular, $X = \{x\}$, the notation $1_Y(x)$ is used.

To exhaustively iterate all hypothetical report-target associations, the operator $\mathcal{F}(\mathbb{X})$ is used to denote the collection of all subsets of set $\mathbb{X}$.

## III. Clustering and Spatial Indexing

Hypothetically, any report can be associated with any target, but in practice many of those associations will be very unlikely to be true. As such, if the hypothesis search would be limited, most associations would never be part of any generated hypotheses. That is, targets and reports which are unlikely to be associated are approximately independent. We therefore wish to define a function and a limit, such that — for a given report and target — if the function is below the limit, the assignment need not even be considered. This function is called a *gate*, and the limiting process is called gating [20, 2]. Gating is a standard technique for limiting the number of considered associations, and is further discussed below. Targets which share connections a set of reports are considered as a *cluster*, independent from the other targets.

If a report falls outside the gate for all targets in a cluster, the report will not be associated with that cluster. If a report is associated with several clusters, these need to be merged such that each report is associated with a single cluster [20]. Note that this can have a "snowballing" effect, leading to seemingly unrelated tracks belonging to the same cluster, as they are indirectly connected by ambiguous reports.

### A. Target-Report Matching

The likelihood of association between a report and a target is based on the probability density of a true report being in the (estimated) target probability space [3]. In the case of Gaussian reports and targets, this is synonymous with the probability density function (pdf) of the innovation, itself a Gaussian [17].

As a general result, the area bounded by any Gaussian ($N(\bar{\mu}, \Pi)$) iso-probability limit corresponds to an ellipse

$$\left\{x \quad : \quad (x - \bar{\mu})^T \Pi^{-1} (x - \bar{\mu}) \leq K\right\}, \tag{5}$$

the size and orientation of which can be calculated from the covariance matrix $\Pi$ and the limit parameter $K$ [22].

Hence, gating through a limit of the association probability is referred to as ellipsoidal gating [3], the corresponding ellipse being (with [transformed] target and measurement covariance $P$ and $R$ respectively, in a common space)

$$\left\{x \quad : \quad (x - \hat{x})^T (P + R)^{-1} (x - \hat{x}) \leq K\right\}. \tag{6}$$

In [3], Collins and Uhlmann show that a necessary condition for Eq. (6) is that

$$\exists x : \quad \begin{cases} (x - \hat{x}_\tau)^T P^{-1} (x - \hat{x}_\tau) \leq K \\ (x - \hat{x}_z)^T R^{-1} (x - \hat{x}_z) \leq K \end{cases}, \tag{7}$$

that is, that the ellipses enscribed respectively by the covariances of the state estimate and the report overlap.

This motivates the use of intersections as a method for gating, and since the ellipsoids of the tracks and reports are independent, the gate equation in Eq. (6) does not need to be recomputed for each association pair.

To further simplify the gating, we find the axis-aligned bounding box of the ellipse in Eq. (5):

$$B = \begin{bmatrix} p_{x-}, & p_{x+}, & p_{y-}, & p_{y+} \end{bmatrix}. \tag{8}$$

where the components are given by (for dimension $d \in \{x, y\}$)

$$p_{d\pm} = \mu_d \pm \sqrt{a_d^2 + b_d^2}. \tag{9}$$

and $\bar{a} = (a_x, a_y)^T$, $\bar{b} = (b_x, b_y)^T$ are eigenvectors of the position covariance, describing the ellipse axes of symmetry [22, 17].

Since the association probability at the square bounds of the ellipses in Eq. (7) is lower or equal than at the ellipsoid edge, the intersection of their bounding boxes is a conservative estimate of Eq. (7): if the bounding boxes do not overlap, the ellipses do not either.

The concept may be extended to other unimodal distributions, as well as mixtures by using the minimum bounding box of the desired probability limit.

### B. Cluster Separation

From the matching of reports to targets, a connection graph (visualized in Figure 5) can be formed by connecting pairs of targets which both are a potential match for a common report. The resulting graph will contain one or more groups of connected components which represent the targets which must be kept in the same cluster. Algorithms for finding graph connected components are studied in e.g. [25, 18].

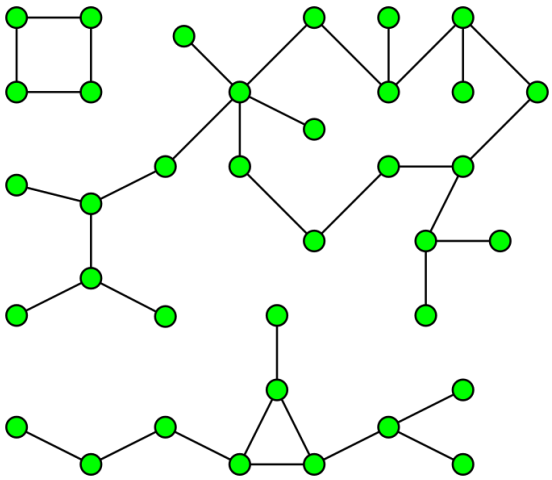

Figure 5: Targets within the same group of connected components must be kept in the same cluster.

The target labels in $\mathcal{L}_+ = \mathcal{L} \cup \mathcal{B}$ (the sets of old and newborn target labels) is therefore correspondingly partitioned into disjoint sets

$$\mathcal{L}_+ = \bigcup_{n=1}^{N} \mathcal{L}_+^{(n)}$$

with $\mathcal{L}_+^{(n)} \cap \mathcal{L}_+^{(m)} = \emptyset$ for $n \neq m$. Similarly, the set of current reports $Z$ can be partitioned into the corresponding clusters

$$Z = Z^{(0)} \bigcup_{n=1}^{N} Z^{(n)}$$

with $Z^{(n)} \cap Z^{(m)} = \emptyset$ for $n \neq m$, and $Z^{(0)}$ is the set of measurements not associated with any previously known target.

### C. Spatial Indexing

Unlike in e.g. the Multiple Hypothesis Tracking (MHT) filter, the targets in the Labeled Multi-Bernoulli filter are not connected to each other through restrictions of track history compatibility — each target is indepedent following each update. This means that for an incoming scan it can be easily be determined, on a target-level, which targets that fall in the sensor field-of-view and thus will be affected by the correction update, simply by looking at the pdf's of the current time-step. Similarly, the associations between reports is a novel matching process between the pdf's of each report and that of each target.

As previously noted, the area in which a report may match a target may be limited through gating, and any gate may be approximated by its axis-aligned bounding box, as exemlified for the Gaussian case in Eq. (8). Such rectangular areas are well suited for e.g. fast intersection lookups, when stored in structures suitable for spatial indexing, such as the R-tree [7]. By storing the targets in a database indexed by their bounding-boxes, the affected targets of each scan can as such be efficiently loaded using bounding box intersections and grouped into clusters as per the previous section — leaving unaffected targets entirely unloaded from the database.

## IV. The Labeled Multi-Bernoulli Filter

The Labeled Multi-Bernoulli filter was proposed in [21] as a simplification of the $\delta$-GLMB-filter [27, 26]. In this section, we review its general formulation, as well as its underlying algorithms and concepts.

### A. Labeled Multi-Bernoulli Filter Outline

The Labeled Multi-Bernoulli filter is defined in the framework of Finite Set Statistics (FISST) [21], of which the Random Finite Set (RFS) is an integral part. An RFS is a set with a probabilistic cardinality distribution, i.e. each potential element is included in the set with a given probability. Specifically, a Bernoulli RFS is a random set which is empty with probability $1 - r$, and

with probability $r$ is a singleton. For an element $x$ with probability $p(\cdot)$, the Bernoulli RFS pdf is given by

$$\pi(X) = \begin{cases} 1-r & \text{if } X = \emptyset \\ r \cdot p(x) & \text{if } X = \{x\} \end{cases}. \tag{10}$$

A multi-Bernoulli RFS is the resulting set of the union of $M$ independent Bernoulli-distributed random finite sets $X^{(i)}$: $X = \bigcup_{i=1}^{M} X^{(i)}$. Consequently, the multi-Bernoulli RFS is parametrized by the set $\left\{\left(r^{(i)}, p^{(i)}\right)\right\}_{i=1}^{M}$, and its pdf is given by [15]

$$\begin{aligned}\pi(\{x_1,\ldots,x_n\}) = &\prod_{j=1}^{M}\left(1-r^{(j)}\right) \\ &\times \sum_{1\le i_1\neq\cdots\neq i_n\le M}\prod_{j=1}^{n}\frac{r^{(i_j)}p^{(i_j)}(x_j)}{1-r^{(i_j)}}.\end{aligned} \tag{11}$$

The labeled multi-Bernoulli is obtained by the augmentation of each Bernoulli RFS with a unique label, $\ell \in \mathcal{L}$. The Labeled Multi-Bernoulli (LMB) RFS can thus be described by the set

$$\left\{\left(r^{(\ell)}, p^{(\ell)}\right)\right\}_{\ell\in\mathcal{L}}.$$

As this set fully describes a multi-target probability density, $\pi(X)$, in the following the shorthand notation $\pi = \left\{\left(r^{(\ell)}, p^{(\ell)}\right)\right\}_{\ell\in\mathcal{L}}$ will be used.

The Labeled Multi-Bernoulli filter follows the classical predict/correct filter recursion, each step outlined below.

*1) LMB Prediction:* Given a LMB pdf $\pi(X)$, the prediction step of the LMB filter and the updated distribution, $\pi_+(X)$, is obtained by the application of the standard prediction update of a Bayesian filter, the Chapman-Kolmogorov equation,

$$\pi_+(X_+) = \int f(X_+)\,\pi(X)\,\delta X. \tag{12}$$

This gives the following set of surviving and new-born targets [21],

$$\pi_+ = \left\{\left(r_{+,S}^{(\ell)}, p_{+,S}^{(\ell)}\right)\right\}_{\ell\in\mathcal{L}} \cup \left\{\left(r_B^{(\ell)}, p_B^{(\ell)}\right)\right\}_{\ell\in\mathcal{B}}, \tag{13}$$

where

$$r_{+,S}^{(\ell)} = \eta_S(\ell)\, r^{(\ell)} \tag{14}$$

$$p_{+,S}^{(\ell)} = \frac{\langle p_S(\cdot,\ell)\, f(x|\cdot,\ell), p(\cdot|\ell)\rangle}{\eta_S(\ell)} \tag{15}$$

$$\eta_S(\ell) = \langle p_S(\cdot,\ell), p(\cdot,\ell)\rangle \tag{16}$$

and $p_S(\cdot,\ell)$ is the distribution of target survival probability. The set $\left\{\left(r_B^{(\ell)}, p_B^{(\ell)}\right)\right\}_{\ell\in\mathcal{B}}$ is given by the birth model, further discussed in Section IV-B.

*2) LMB Correction:* Drawn from the update of $\delta$-GLMB [27], the LMB correction is derived in [21].

In general, as noted in [21], the LMB distribution is not closed under the Bayesian filter update. However, the resulting $\delta$-GLMB distribution — which can represent multiple disjoint hypotheses — may be approximated as in Eq. (18) with an LMB pdf through the collapse of its hypotheses, weighted by their probability. That is, the correction update updates the set

$$\pi_+ = \left\{\left(r_+^{(\ell)}, p_+^{(\ell)}\right)\right\}_{\ell\in\mathcal{L}_+} \tag{17}$$

by the following approximation:

$$\begin{aligned}\pi(\cdot|Z) &\approx \left\{\left(r^{(\ell)}, p^{(\ell)}\right)\right\}_{\ell\in\mathcal{L}_+} \\ &= \bigcup_{i=1}^{N}\left\{\left(r^{(\ell,i)}, p^{(\ell,i)}\right)\right\}_{\ell\in\mathcal{L}_+^{(i)}}\end{aligned}, \tag{18}$$

in which parameters are given by

$$r^{(\ell,i)} = \sum_{(I_+,\theta)\in\mathcal{F}\left(\mathcal{L}_+^{(i)}\right)\times\Theta_{I_+}} w^{(I_+,\theta)}\left(Z^{(i)}\right) 1_{I_+}(\ell) \tag{19}$$

$$\begin{aligned}p^{(\ell,i)}(x) = &\frac{1}{r^{(\ell,i)}} \sum_{(I_+,\theta)\in\mathcal{F}\left(\mathcal{L}_+^{(i)}\right)\times\Theta_{I_+}} w^{(I_+,\theta)}\left(Z^{(i)}\right) \\ &\times 1_{I_+}(\ell)\, p^{(\theta)}(x,\ell)\end{aligned} \tag{20}$$

$$\sum_{(I_+,\theta)\in\mathcal{F}\left(\mathcal{L}_+^{(i)}\right)\times\Theta_{I_+}} w^{(I_+,\theta)}\left(Z^{(i)}\right) = 1 \tag{21}$$

where $\Theta_{I_+}$ is the space of mappings of tracks $\theta : I_+ \to \left\{0,1,\ldots,\left|Z^{(i)}\right|\right\}$ such that $\theta(\iota) = \theta(\iota') > 0$ implies that $\iota = \iota'$ i.e. the mapping is unique for all values except those mapped to zero [21]. Also, for $I_+ \subseteq \mathcal{L}_+^{(i)}$

$$w^{(I_+,\theta)}\left(Z^{(i)}\right) \propto w_{+,i}^{(I_+)}\left[\eta_{Z^{(i)}}^{(\theta)}\right]^{I_+} \tag{22}$$

$$w_{+,i}^{(I_+)} = \prod_{\ell\in\mathcal{L}_+^{(i)}-I_+}\left(1-r_+^{(\ell)}\right)\prod_{\ell'\in I_+} r_+^{(\ell)} \tag{23}$$

$$\eta_{Z^{(i)}}^{(\theta)}(\ell) = \left\langle p_{+,i}(x,\ell), \psi_{Z^{(i)}}(\cdot,\ell;\theta)\right\rangle \tag{24}$$

$$\psi_{Z^{(i)}}(x,\ell;\theta) = \begin{cases} \frac{p_D(x,\ell)p_G g\left(z_{\theta(\ell)}|x,\ell\right)}{\kappa\left(z_{\theta(\ell)}\right)}, \theta(\ell) > 0 \\ q_{D,G}(x,\ell), \theta(\ell) = 0 \end{cases} \tag{25}$$

$$q_{D,G}(x,\ell) = 1 - p_D(x,\ell)\, p_G \tag{26}$$

$$p^{(\theta)}\left(x,\ell|Z^{(i)}\right) = \frac{p_{+,i}(x,\ell)\,\psi_{Z^{(i)}}(x,\ell;\theta)}{\eta_{Z^{(i)}}^{(\theta)}(\ell)} \tag{27}$$

Note that while the $\delta$-GLMB representation in [21] is an intermediate representation in the theoretical derivation of the filter, its construction in implementation is not necessary to reach the collapsed LMB representation of Eq. (18).

Also; In general MTT, if different hypothetical associations assign the same report to different targets, the resulting tracks will generally be incompatible. Through

the approximations of the LMB filter, assignment compatibility is only considered in the hypothesis generation, but then lost in the summations in Eqs. (19)-(20).

To calculate the weight — from Eq. (22) — for an hypothesis we start by making the distinction between associated and non-associated targets by splitting the hypothesis label set $I_+$:

$$I_+^a = \{\ell : \quad \theta(\ell) \neq z_\emptyset\}_{\ell \in I_+}, \tag{28}$$

$$I_+^n = \{\ell : \quad \theta(\ell) = z_\emptyset\}_{\ell \in I_+}, \tag{29}$$

(implying $I_+ = I_+^a \cup I_+^n$ and $I_+^a \cap I_+^n = \emptyset$). We can then rewrite Eqs. (22)-(26) as

$$\begin{aligned} w^{(I_+,\theta)}\left(Z^{(i)}\right) \propto & w_{+,i}^{(I_+)}\left[\eta_{Z^{(i)}}^{(\theta)}\right]^{I_+} \\ = & \prod_{\ell \in \mathcal{L}_+^{(i)} - I_+}\left(1 - r_+^{(\ell)}\right) \\ & \times \prod_{\ell' \in I_+^a} r_+^{(\ell')} \eta_{Z^{(i)}}^{(\theta)}(\ell') \prod_{\ell'' \in I_+^n} r_+^{(\ell'')} \eta_{Z^{(i)}}^{(\theta)}(\ell'') \end{aligned} \tag{30}$$

The product of Eq. (30) can be efficiently expressed using the Negative Log Likelihoods (NLL's), $\Lambda_\ell$;

$$e^{-\Lambda_\ell} = \begin{cases} 1 - r_+^{(\ell)} \text{ if } \ell \in \mathcal{L}_+^{(i)} - I_+ \\ r_+^{(\ell)} \eta_{Z^{(i)}}^{(\theta,a)}(\ell) \text{ if } \ell \in I_+^a \\ r_+^{(\ell)} \eta_{Z^{(i)}}^{(\theta,n)}(\ell) \text{ if } \ell \in I_+^n \end{cases} \tag{31}$$

yielding

$$w^{(I_+,\theta)}\left(Z^{(i)}\right) \propto \exp\left(-\sum_{\ell \in \mathcal{L}_+^{(i)}} \Lambda_\ell\right). \tag{32}$$

### B. Adaptive Birth Model

To include new targets in the tracker, the LMB filter relies on a birth distribution. Following [21], the selected birth model for time $k+1$ is based on the reports of time $k$:

$$\pi_{B,k+1} = \left\{\left(r_B^{(\ell)}, p_B^{(\ell)}\right)\right\}_{\ell \in \mathcal{B}_k} \tag{33}$$

for new labels in $\mathcal{B}_k$ generated for each report in $Z_k$.

The existance probabilities of new targets in this model are proportional to the probability of the report not being associated with any previously known target. For a report, the association probability is given by

$$r_{U,k}(z) = \sum_{(I_+,\theta) \in \mathcal{F}\left(\mathcal{L}_+^{(i)}\right) \times \Theta_{I_+}} w^{(I_+,\theta)}\left(Z^{(i)}\right) 1_\theta(z). \tag{34}$$

Given an expected number of new targets in each scan, $\lambda_{B,k+1}$, the existance probability of new targets is thus given by

$$r_{B,k+1}(z) = \min\left(r_{B,max}, \frac{1 - r_{U,k}(z)}{\sum_{\xi \in Z_k} 1 - r_{U,k}(\xi)} \cdot \lambda_{B,k+1}\right). \tag{35}$$

Note that, for $\lambda_{B,k+1} > 1$, the existance probability may need to be limited by the min()-clause to a maximum value of $r_{B,max} \leq 1$.

### C. Hypothesis Generation

The combinatorial nature of report-target associations employed in the correction step of the LMB filter results in an explosion of possible hypotheses. Each hypothesis weight, $w^{(I_+,\theta)}\left(Z^{(i)}\right)$, controls its impact on Eqs. (19)-(20), so an effective way to limit the computational cost of the algorithm is thus to focus on generating and evaluating only the most relevant association hypotheses, with the highest weights.

The generation of the single best possible assignment is known as the *Linear Assignment Problem* (LAP). The problem can be formulated, using the elements the elements $c_{ij}$ of the cost-matrix $C$, as follows [11] (assuming more columns than rows):

$$\begin{aligned} & \min \sum_{i,j} c_{ij} x_{ij} \\ & \sum_j x_{ij} = 1, \quad \forall i, \quad \sum_i x_{ij} \leq 1, \quad \forall j \\ & \forall x_{ij} \in \{0,1\} \end{aligned} \tag{36}$$

The gist of this problem is to select the lowest sum of costs from an assignment costs matrix $C$, for each row selecting exactly one column value.

In the context of LMB, the cost matrix is created for each target cluster, from the NLL's of the track assignments outlined in Eq. (31), resulting in a matrix exemplified in Eq. (37) for a two-track hypothesis and a two-report scan.

$$C = \begin{pmatrix} {}^{z_1}\Lambda_{\ell_1} & {}^{z_2}\Lambda_{\ell_1} & {}^{n}\Lambda_{\ell_1} & \infty & {}^{F}\Lambda_{\ell_1} & \infty \\ {}^{z_1}\Lambda_{\ell_2} & {}^{z_2}\Lambda_{\ell_2} & \infty & {}^{n}\Lambda_{\ell_2} & \infty & {}^{F}\Lambda_{\ell_2} \end{pmatrix} \tag{37}$$

Each solution to Eq. (36) corresponds to an assignment hypothesis in the filter, and the "cheapest" assignment is the best guess of assignment of the reports received for this given timestep. Several algorithms exist to solve the LAP problem — the auction algorithm [1], the Hungarian algorithm [12] and Jonker-Volgenant [11] being notable mentions.

The task of finding the single best assignment was extended in an algorithm due to Murty [16] to that of finding the $k$ best assignments, for a given assignment matrix $C$. The Murty algorithm, with the underlying LAP solver, is of polynomial complexity in the number of tracks with which the reports can be associated [4], and it is therefore of interest to limit the number of tracks considered for association. Herein lies the value of clustering — as the targets and reports of separate clusters are independent, each cluster may be considered as a separate LAP problem.

Truncation of the hypothesis list is achieved through the termination of Murty's algorithm based on either a maximum number of drawn hypotheses, or a maximum hypothesis NLL sum (minimum hypothesis probability).

## V. Implementation

This section introduces the implementation of the Spatially Indexed Multi-Bernoulli filter implemented for this article. The implementation, in the Python language, is available under a Free and Open-Source Software (FOSS) license at https://github.com/jonatanolofsson/lmb.

A few remarks regarding the implementation;

**Particle Implementation** For now, the implementation is based on particle filter distributions, where the general equations of Eq. (24)–(27) are specialized as in [26].

**Parallelization** Due to restrictions in Python parallelization, clusters are currently updated sequentially. However, particle updates are vectorized using the numpy package.

**Target Storage and Indexing** In the implementation, targets are serialized post-update and stored in an SQLite database. The database is R-tree indexed based on the axis-aligned bounding boxes of each target pdf. This allows for a fast extraction of relevant targets in the initial gating process, allowing large parts of the filter to remain dormant through correction updates.

**Rectangular Gating** The implementation makes use of rectangular gating through the minimum bounding box of reports and targets. As a first stage, all targets within the sensor field-of-view's minimum axis-aligned bounding-box are loaded from the database. Second, the bounding-box of each target is tried for intersection with the bounding-box of each report to establish the clusters, in accordance with Section III-B. Notably, these bounding-boxes need not be axis-aligned.

## VI. Results

As a first result, we study the implementation of the Spatially Indexed Labeled Multi-Bernoulli filter through a linear simulated multi-target tracking scenario. Next we present a simulated sea ice tracking scenario with the collaboration between two UAS agents and a satellite.

### A. Crossing Tracks

This scenario, detailed in Figure 6, shows the filter's ability to detect and track objects through a crossing by considering the multiple likely association possibilities. It also illustrates the track-keeping abilities of the LMB filter and the clustering feature of the implementation.

We see, in Figure 6,

a) The crossing tracks scenario in which in total of five targets are tracked through collision courses. To simplify visual interpretation, each object was given a velocity of 1 in the $x$-direction, to match the time dimension of the following plots. Notably, all tracks retain the correct association troughout the simulation, as indicated by the consistent color of each straight track.

b) Target cardinality, for both true, estimated and verified targets. A target is considered verified if its existance probability exceeds 0.7.

c) Number of clusters which the algorithm separates. Note that due to hypothetical new-born targets, this can in fact exceed the number of estimated targets.

d) Number of hypotheses used. In sequences where multiple associations are possible, more hypotheses are generated before iteration termination due to low hypothesis probability. In particular, note the peak between $t = 5$ and $t = 10$ corresponding to the period where the two initial targets cross tracks.

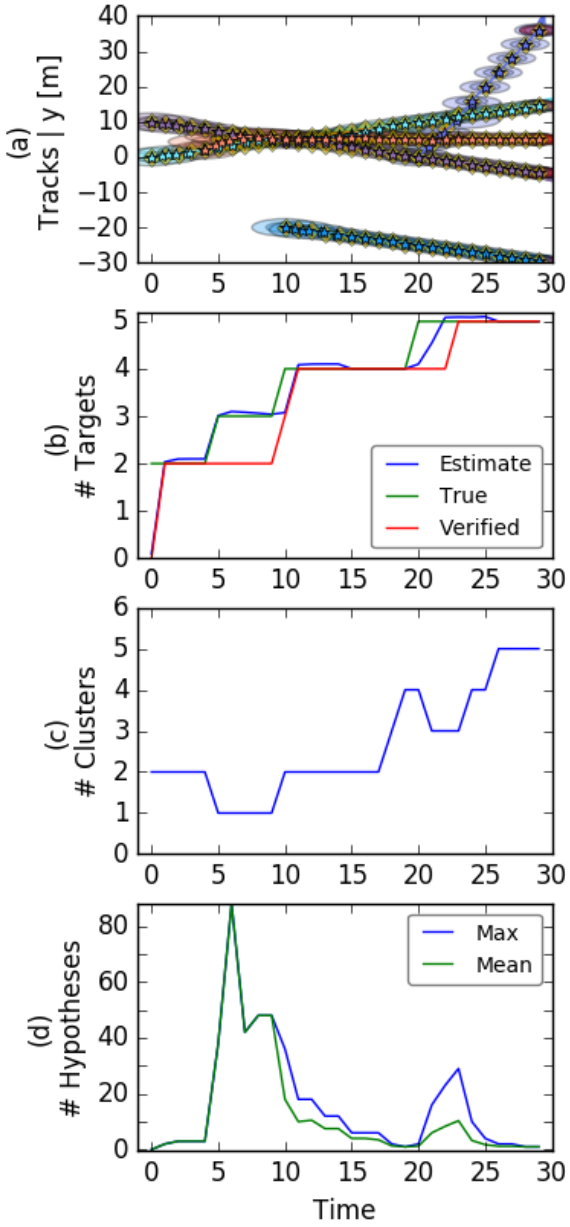


Figure 6: Crossing tracks scenario.

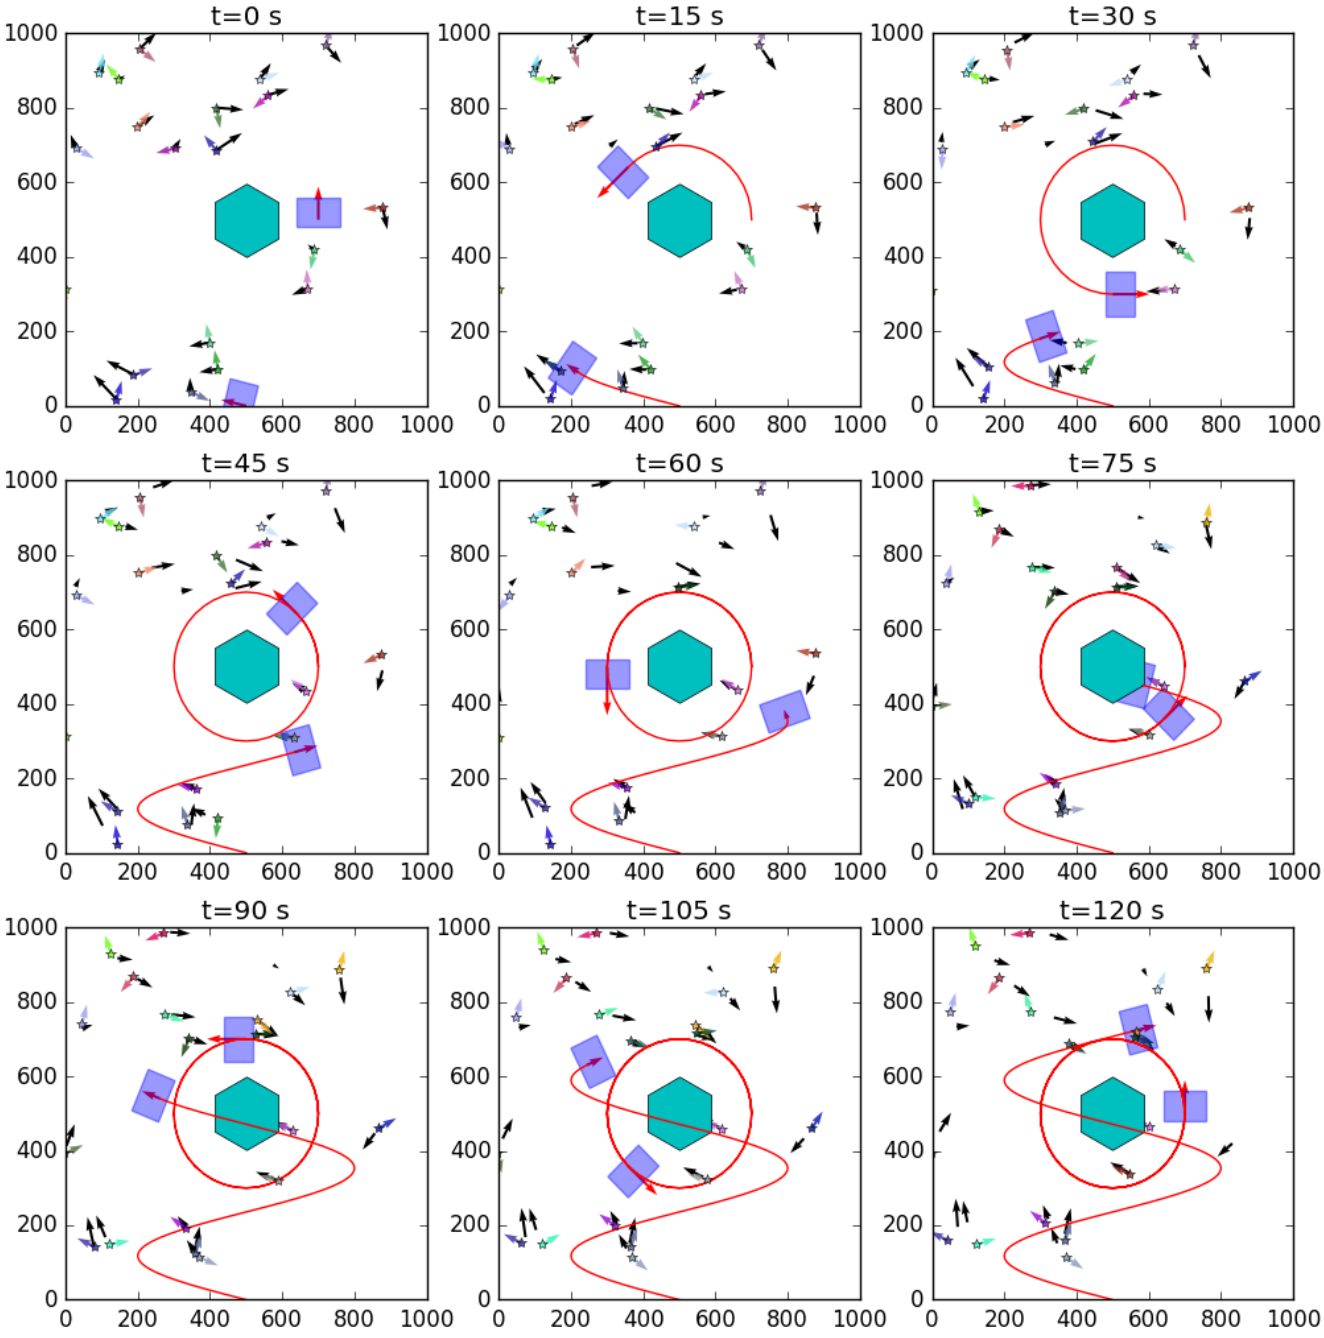


Figure 7: The sea ice tracking scenario is based on the collaboration of two UAS's to loiter and explore respectively, assisted by satellite radar measurements. 20 icebergs are simulated in this scenario. The (black) simulated icebergs are here shown along with the (colored) estimates resulting from the joint observations.

### B. Sea Ice Tracking

In preparation of upcoming sea ice tracking field-tests, a scenario was devised to emulate the anticipated data. In this scenario, shown in Figure 7, sea ice objects are caught in a vortex which is initially observed in full by satellite imagery. In the interest of protecting a fixed installation — in the center of the figure — in the following sequence the sea ice in the vortex is partially observed by two independent UAS agents which report wirelessly to the central filter. The central filter then continously fuses the data received from each sensor into a joint estimate of the sea ice flow field. The observation is also assisted by an incoming processed satellite image at time 75 which span the entire area of interest.

The detections from each UAS are drawn from a field-of-view which in the UAS-local frame of reference corresponds to the bounding box

$$B_{\text{UAS}} = \begin{bmatrix} -60, & 60, & -20, & 60 \end{bmatrix} \tag{38}$$

in each UAS's frame of reference at each time-step.

As sample values, detection probability of UAS reports are set to 0.99, and for the satellite reports to 0.8. 20 icebergs are simulated. As it is unlikely that a verified ice object disappears during the short timeframe of the UAS flight, each track has a survival probability (in the prediction step) of 0.999, uniformly. In practice, disappearing objects are instead likely to be removed by the correction step, as they will no longer yield reports likely to be associated to the target.

The tracks of the UAS agents are a combination of a loitering UAS with a specific area to guard, and a second UAS track with the intention of more broadly exploring the area.

In Figures 7, the ice objects (in black) are plotted with their velocity vectors, togheter with the (colored) filter estimates. UAS tracks are drawn red. In Figure 8, the final output of the filter is displayed as a Gaussian field.

Notably, whereas position information is available from the fully covering satellite imagery, velocity may only be observed through repeated observations and associations. New targets are initiated with a large variance in their velocity estimate, which results in a large variety of velocities considered within the individual filters. Although only the mean velocity estimates are displayed in the plots for clarity, this means a lot of options for future associations will be considered as the uncertainty in velocity is translated to uncertainty in position over time.

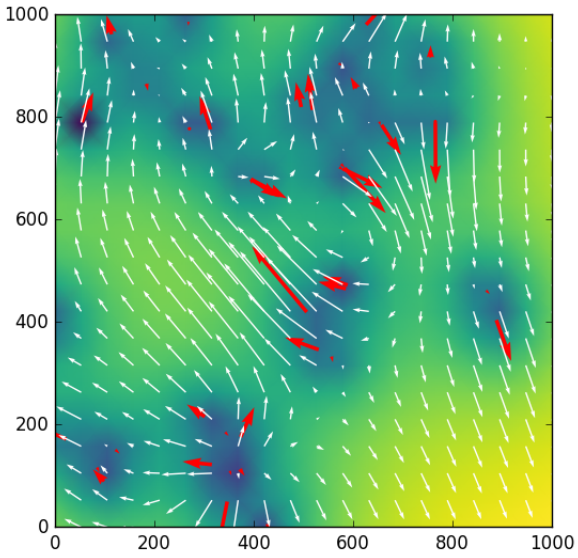


Figure 8: The tracker estimates may be used as an input for current and wind estimates (track estimates shown in red).

## VII. Conclusions

In this paper, an implementation is presented of the Spatially Indexed Labeled Multi-Bernoulli filter, and a preview of its intended application is presented through a simulated iceberg-tracking scenario, combining the information available from both satellite and multiple UAS agents. The scenario is helpful in providing a testbench with which to study the aspects of the LMB filter in preparation for the acquisition of real data.

In this setup, the tracking methodology is limited to point-target tracking. It is of interest, and for future work, to extend this into extended target tracking — in which the restriction of *one report association per target* is leviated through the modelling of geometrical extension. As much data is acquired in the form of imagery, extension information about the tracked ice is easily accessible and as such suits the tracking problem well.

Another issue to be explored further in the case of tracking ice is that of various death- and birth models. Ice in available datasets often exhibits merging and splitting, where several individual detections clump together in what will appear as a single target. This behaviour has been studied for other filters, e.g. in [24], but remains as future research and a future extension in the implementation of the Spatially Indexed Multi-Bernoulli filter.

## Acknowledgments

This project has received funding from the European Union's Horizon 2020 research and innovation programme under the Marie Sklodowska-Curie grant agreement No 642153, as well as the Research Council of Norway through the Centres of Excellence funding scheme, grant number 223254 – NTNU-AMOS.

## References


[1] D. P. Bertsekas, "The auction algorithm: A distributed relaxation method for the assignment problem," *Annals of Operations Research*, vol. 14, no. 1, pp. 105–123, 1988.

[2] S. S. Blackman, "Multiple Hypothesis Tracking For Multiple Target Tracking," *IEEE Aerospace and Electronic Systems Magazine*, vol. 19, no. 1 January, 2004.

[3] J. Collins and J. Uhlmann, "Efficient gating in data association with multivariate Gaussian distributed states," *IEEE Transactions on Aerospace and Electronic Systems*, vol. 28, no. 3, pp. 909 – 916, 1992.

[4] I. J. Cox and S. L. Hingorani, "Efficient implementation of Reid's multiple hypothesis tracking algorithm and its evaluation for the purpose of visual tracking," *IEEE Transactions on Pattern Analysis and Machine Intelligence*, vol. 18, no. 2, pp. 138–150, 1996.

[5] K. Eik, "Review of Experiences within Ice and Iceberg Management," *Journal of Navigation*, vol. 61, no. 04, p. 557, 2008.

[6] ——, "Iceberg drift modelling and validation of applied metocean hindcast data," *Cold Regions Science and Technology*, vol. 57, no. 2-3, pp. 67–90, 2009.

[7] A. Guttman, "R-trees: A Dynamic Index Structure for Spatial Searching," in *Proceedings of the 1984 ACM SIGMOD International Conference on Management of Data*, Boston, Massachusetts, USA, 1984, pp. 47–57.

[8] J. Haugen, "Autonomous Aerial Ice Observation," Ph.D. dissertation, Norwegian University of Science and Technology, 2014.

[9] N. Johannessen, O.M., Alexandrov, V., Frolov, I.Y., Sandven, S., Pettersson, L.H., Bobylev, L.P., Kloster, K., Smirnov, V.G., Mironov, Y.U., Babich, *Remote Sensing of Sea Ice in the Northern Sea Route*, 2007. [Online]. Available: http://www.springerlink.com/index/10.1007/978-3-540-48840-8%5Cnhttp://link.springer.com/10.1007/978-3-540-48840-8

[10] T. A. Johansen and T. Perez, "Unmanned Aerial Surveillance System for Hazard Collision Avoidance in Autonomous Shipping," 2015.

[11] R. Jonker and A. Volgenant, "A shortest augmenting path algorithm for dense and sparse linear assignment problems," *Computing*, vol. 38, no. 4, pp. 325–340, 1987.

[12] H. W. Kuhn, "The Hungarian Method for the assignment problem," *Naval Research Logistics Quarterly*, vol. 2, pp. 83–97, 1955.

[13] T. Kurien, "Issues in the Design of Practical Multitarget Tracking Algorithms," in *Multitarget-Multisensor Tracking: Advanced Applications*, Y. Bar-Shalom, Ed. Artech House, 1990, pp. 43–83.

[14] F. S. Leira, "Object Detection and Tracking With UAVs," Ph.D. dissertation, Norwegian University of Science and Technology, 2017.

[15] R. P. S. Mahler, *Statistical Multisource-Multitarget Information Fusion*. Norwood, MA, USA: Artech House, Inc., 2007.

[16] K. G. Murty, "An Algorithm for Ranking all the Assignments in Order of Increasing Cost," *Operations Research*, vol. 16, no. 3, pp. 682–687, 1968.

[17] J. Olofsson, E. Brekke, T. I. Fossen, and T. A. Johansen, "Spatially Indexed Clustering for Scalable Tracking of Remotely Sensed Drift Ice," in *IEEE Aerospace Conference Proceedings*, Big Sky, MT, USA, 2017.

[18] D. J. Pearce, "An Improved Algorithm for Finding the Strongly Connected Components of a Directed Graph," Victoria University, Wellington, NZ, Tech. Rep., 2005.

[19] C. Randell, F. Ralph, D. Power, and P. Stuckey, "OTC 20264 Technological Advances to Assess , Manage and Reduce Ice Risk in Northern Developments," in *Offshore Tecdhnology Conference*, 2009.

[20] D. B. Reid, "An algorithm for tracking multiple targets," in *1978 IEEE Conference on Decision and Control including the 17th Symposium on Adaptive Processes*, vol. 17, no. 6, San Diego, California, USA, 1978, pp. 843–854.

[21] S. Reuter, B.-T. Vo, B.-N. Vo, and K. Dietmayer, "The Labeled Multi-Bernoulli Filter," *IEEE Transactions on Signal Processing*, vol. 62, no. 12, pp. 3246–3260, 2014.

[22] M. I. Ribeiro, "Gaussian Probability Density Functions : Properties and Error Characterization," Institute for Systems and Robotics, Lisboa, Portugal, Tech. Rep. February, 2004.

[23] J. J. Stein and S. S. Blackman, "Generalized Correlation of Multi-Target Track Data," *IEEE Transactions on Aerospace and Electronic Systems*, vol. AES-11, no. 6, November, pp. 1207–1217, 1975.

[24] C. B. Storlie, T. C. M. Lee, J. Hanning, and D. Nychka, "Tracking of multiple merging and splitting targets: a statistical perspective," *Statistica Sinica*, vol. 19, pp. 1–52, 2009.

[25] R. Tarjan, "Depth-first Search and Linear Graph Algorithms," *SIAM Journal on Computing*, vol. 1, no. 2, pp. 146–160, 1972.

[26] B.-N. Vo, B.-T. Vo, and D. Phung, "Labeled random finite sets and the bayes multi-target tracking filter," *IEEE Transactions on Signal Processing*, vol. 62, no. 24, pp. 6554–6567, 2014.

[27] B.-T. Vo and B.-N. Vo, "Labeled random finite sets and multi-object conjugate priors," *IEEE Transactions on Signal Processing*, vol. 61, no. 13, pp. 3460–3475, 2013.