Doctoral theses at NTNU, 2024:218

Audun Gullikstad Hem

Maritime Target Tracking With Exteroceptive Sensors and Target-Provided Information

Doctoral thesis

NTNU
Norwegian University of Science and Technology
Thesis for the Degree of
Philosophiae Doctor
Faculty of Information Technology and Electrical Engineering
Department of Engineering Cybernetics

NTNU
Norwegian University of
Science and Technology

Audun Gullikstad Hem

Maritime Target Tracking With Exteroceptive Sensors and Target-Provided Information

Thesis for the Degree of Philosophiae Doctor

Trondheim, June 2024

Norwegian University of Science and Technology
Faculty of Information Technology and Electrical Engineering
Department of Engineering Cybernetics

NTNU
Norwegian University of
Science and Technology

NTNU
Norwegian University of Science and Technology

Thesis for the Degree of Philosophiae Doctor

Faculty of Information Technology and Electrical Engineering
Department of Engineering Cybernetics

ISBN 978-82-326-8028-3 (printed ver.)
ISBN 978-82-326-8027-6 (electronic ver.)
ISSN 1503-8181 (printed ver.)
ISSN 2703-8084 (online ver.)

ITK-report: 2024-03-W

Doctoral theses at NTNU, 2024:218

Printed by NTNU Grafisk senter

Summary

Autonomous and unmanned operations at sea are becoming increasingly common. However, as the operational goals increase in ambition, the existing methods show their weaknesses. The maritime domain can be complex and unpredictable, whereas the system demands in unmanned and autonomous operations are high. One of the most important pieces needed for maritime autonomy is situational awareness. Identifying targets and tracking their movements is a key component of situational awareness. This thesis concerns target tracking, with a focus on maritime applications.

The main contribution of this thesis is a methodology for the utilization of target-provided information in target tracking. In many applications, targets may transmit messages with valuable information regarding their identity, position, direction, and more. At sea, the Automatic identification system (AIS) protocol specifies the transmission of such messages and is mandatory for commercial ships and widely used by other vessels. Target-provided information has played a role in situational awareness since it became available, but often in cursory manners such as by plotting the position of the transmitting vessel on a map. This thesis combines the usefulness of target-provided information with advanced target tracking methods and explores how they, together with exteroceptive sensor measurements, can be an invaluable addition to the situational awareness of a vessel.

Another contribution is the development, testing, and evaluation of the advanced trackers used to process the target-provided information together with exteroceptive measurements. The thesis presents an extension of the joint integrated probabilistic data association filter that includes support for multiple kinematic target models and modeling of target visibility. Both with and without the addition of target-provided measurements, the tracker is tested on maritime data and shown to be robust when faced with challenging problems. Furthermore, the thesis presents a variant of the Poisson multi-Bernoulli mixture filter with multiple kinematic models that includes target-provided measurements, with applications to both point target and extended object tracking. The results show that target-provided information such as position, speed, course, and vessel dimension generally improve the tracking results, and in

certain situations, performance improvements can be considerable.

The thesis also considers how to deploy such methods in real-life situations and presents a method for validating AIS message information. Validating the messages ensures their safe use when estimating the target states and enables their use in other parts of a large system. Furthermore, the trackers are tested in several fully autonomous collision avoidance scenarios, where the autonomous vessel has to safely avoid other vessels. The results show that the methods enable the vessel to make correct decisions and solve issues present with previous tracking methods used in similar situations.

Lastly, some specific problems related to maritime target tracking are considered, namely time inaccuracies in the received measurements and false alarms from the wakes of targets. The presented solutions improve the tracking performance when the problems are encountered.

Preface

This thesis is submitted in partial fulfillment of the requirements for the degree of Philosophiae Doctor (Ph.D.) at the Norwegian University of Science and Technology (NTNU) in Trondheim. The work has been carried out at the Department of Engineering Cybernetics (ITK). Professor Edmund F. Brekke has been the main supervisor of this work, with co-supervision from professor Thor I. Fossen, both from the Department of Engineering Cybernetics. The research was conducted as part of the Autosit project, and was funded by the Research Council of Norway, DNV, Maritime Robotics, and Kongsberg.

Acknowledgements

First, I want to underline how thorough and helpful supervision, coupled with an expressed belief that I could power through most obstacles if I only tried hard enough, has greatly helped me achieve my goals with this thesis. Furthermore, this acknowledgment is about all the people in my life. It is about all my brilliant colleagues at Gløshaugen and Nyhavna, many of whom are better described as friends, who have made my workdays fun, productive, and interesting. It is about partners in the research project and colleagues at ITK who have stepped up when called upon to provide my research with resources that would be hard to obtain otherwise. This acknowledgment is also about all my friends in Trondheim and elsewhere, who have made me an expert in explaining my work with a single sentence and leaving it at that before we go do something else. It is about my roommates who made a couple of years of a pandemic a passable experience, and my other good childhood friends who have reluctantly accepted that I have stayed in Trondheim while they have moved elsewhere. Lastly, this acknowledgment is about my family, who I'm very lucky to have and do not take for granted. Many names could be mentioned, but I'll rather apply the tired Norwegian phrase: "ingen nevnt, ingen glemt". You all know who you are.

Audun Gullikstad Hem
Trondheim, January 2024

Contents

List of Abbreviations

Commonly used abbreviations are listed here alphabetically. The abbreviations are also explained in each chapter at their first use.

ADS-B	Automatic dependent surveillance-broadcast
AIS	Automatic identification system
ASV	Autonomous surface vehicle
BP	Belief propagation
COLAV	Collision avoidance
COLREGS	Collision avoidance regulations
CT	Coordinated turn
CV	Constant velocity
EKF	Extended Kalman filter
FISST	Finite set statistics
GOSPA	Generalized optimal sub-pattern assignment
GP	Gaussian process
IMM	Interacting multiple models
IMO	International Maritime Organization
IOU	Intersection over union
IPDA	Integrated probabilistic data association
JIPDA	Joint integrated probabilistic data association

JPDA	Joint probabilistic data association
MAP	Maximum a posteriori
MB	Multi-Bernoulli
MBM	Multi-Bernoulli mixture
MHT	Multiple hypothesis tracking
MMSE	Minimum mean square error
MMSI	Maritime mobile service identity
MPC	Model predictive control
OSPA	Optimal sub-pattern assignment
p.g.fl.	Probability generating functional
PDA	Probabilistic data association
pdf	Probability density function
PF	Particle filter
PHD	Probability hypothesis density
PMBM	Poisson multi-Bernoulli mixture
PPP	Poisson point process
RFS	Random finite set
RIB	Rigid inflatable boat
ROC	Receiver operating characteristic
SB-MPC	Scenario based model predictive control
SLRT	Sequential likelihood ratio test
TFAR	Track false alarm rate
TFR	Track fragmentation rate
TLE	Track localization error
TPD	Track probability of detection

USV	Unmanned surface vehicle
VIMMJIPDA	Visibility interacting multiple-model joint integrated probabilistic data association

List of Symbols

Commonly used symbols are listed here, grouped by topic. The symbols are also explained in each chapter where they are used.

General symbols

N	A number of elements
n	Number of tracks
m	Number of measurements
T	Threshold
dt	Time period between time steps
$(\cdot)_k$	A quantity at time step k
$(\cdot)_{k\|k-1}$	A predicted quantity at time step k
$(\tilde{\cdot})$	Latent variables that are marginalized away
$(\cdot)^0$	An initial quantity
α	Significance level
t	A point in time

Indexes, subscripts, and superscripts

t	Track index
k	Time step index
j	Measurement index *or* PMBM global hypothesis index
i	PMBM component index

s	Model index
v	Visibility state
τ	ID index
ex	Exteroceptive measurement related
tp	Target-provided measurement related
R	Radar related
A	AIS related

States and state estimates

$\mathbf{y}$	Hybrid state vector
$\mathbf{x}$	Continuous state vector
$\mathbf{x}^c$	Kinematic part of the state vector
x, y	Positional elements in the kinematic state vector
v_x, v_y	Velocity elements in the kinematic state vector
ϕ	Heading angle
ω	Turn rate *or* visibility transition probabilities
$\hat{\mathbf{x}}_0^s$	Initial state estimate
$\mathbf{P}_0^s$	Covariance of initial state estimate
$\mathbf{P}_v$	Initial velocity covariance

Measurements

Z	Measurement set
$\mathbf{z}$	Measurement vector
z	Measurement scalar
$\mathbf{p}$	Kinematic measurement information vector
L	Length of a vessel
W	Width of a vessel

Association hypotheses and data association

a	Association hypothesis
$\mathcal{A}$	Set of association hypotheses
w	Weight of association hypothesis
l	Track-to-measurement likelihood

Densities, distributions, and functions

$\mathcal{N}$	Gaussian probability density function
Φ	Gaussian cumulative distribution function
$\mathcal{G}$	Gamma probability density function
Γ	Gamma function
f	Generic (single-target) probability density function
$f_{\mathbf{y}}$	Transition density for hybrid state
$f_{\mathbf{z}}$	Measurement density conditional on hybrid state
$G[h]$	Probability generating functional with test function
L	Index function relating track numbers and measurement numbers
1_{Ω}	Indicator function
d	GOSPA distance function
s	GOSPA-T track switch function

Poisson intensities

λ	Poisson intensity for false alarms *or* generic Poisson intensity
u	Unknown target intensity *or* superscript for unknown target values
U	Unknown target rate
b	Target birth intensity *or* superscript for birth intensity values
B	Target birth rate

Probabilities

o	Visibility probabilities
P_{S}	Survival probability
P_{D}	Detection probability
P_C	Received ID correctness probability
p	Association hypothesis probability
ξ	ID probabilities
r	Existence probability
μ	Mode probabilities
η	Probability of a target being visible
ω	Visibility transition probability *or* turn rate
π	Mode transition probability

Kalman filtering

$\hat{(\cdot)}$	A Kalman filter estimate
$\mathbf{f}$	Kinematic prediction function
$\mathbf{F}$	Process model transition matrix
$\mathbf{H}$	Measurement matrix
$\mathbf{H}^*$	Complementary measurement matrix
$\mathbf{H}^{\mathrm{vel}}$	Velocity measurement matrix
$\mathbf{J}$	Polar-to-cartesian Jacobian
$\mathbf{w}$	Measurement noise
$\mathbf{R}$	Measurement noise covariance matrix
$\mathbf{R}_c$	Cartesian measurement noise covariance
$\mathbf{R}_p$	Polar measurement noise covariance
$\mathbf{S}$	Innovation covariance matrix

$\mathbf{v}$	Process noise
σ	Noise standard deviation
q	Process noise intensity
$\mathbf{Q}$	Process noise covariance matrix

Extended object tracking

$\mathbf{K}(\Theta, \Theta)$	Gaussian process covariance matrix
θ	Measurement angle on target contour *or* surveillance area coverage
C	Cell index subscript
$(\cdot)^l$	Index of measurement in a measurement cell
η_γ	Forgetting factor
α, β	Shape parameters of the gamma distribution

Other

g	Gate size
Ω	Surveillance region
$\mathcal{V}$	Number of possible IDs
p	OSPA metric order
c	OSPA metric cutoff
γ	GOSPA-T metric switch penalty
$\mathcal{E}$	Target extent area
H	Hypothesis in a statistical hypothesis test
ζ	Z-test statistic
Λ	Weighted sum of test statistics
$\mathcal{L}$	Likelihood ratio
ψ	Radar angle
$\boldsymbol{\rho}$	Wake position

Chapter 1

Introduction

This thesis is written as part of the Autosit project at NTNU [110]. Short for *Autonomous ships, intentions and situational awareness*, the project aims to develop methods for improving autonomous navigation at sea. Autonomous and unmanned maritime operations are becoming increasingly common, and the technology is evolving rapidly. Above water, such operations involve autonomous surface vehicles (ASVs) or unmanned surface vehicles (USVs). ASVs and USVs are relevant for many different applications and are in the near future expected to operate in a multitude of environments. Some examples include pedestrian ferries in urban environments [21], cargo transport on inland waterways [127], and bulk carriers at sea [84]. Figure 1.1 shows three examples of ASVs and USVs that are already in use. The first are the Mariner and Otter from Maritime Robotics, two multi-purpose drones. The Otter has, among other things, been used for surveying and mapping purposes [34], whereas the Mariner has been used for, for example, transport and surveying in demanding environments [37, 101]. The third example is the MF Estelle, a passenger ferry from Torghatten and Zeabuz that operates in Stockholm [38]. However, for such autonomous operations to be safe and efficient, the vessels must understand and be aware of their surroundings.

1.1 Situational awareness in the maritime domain

In broad terms, situational awareness denotes the problem of understanding the environment surrounding a vessel. For this purpose, the use of sensor technology is useful, and sometimes strictly necessary. Whenever a ship is steered by a captain, additional information from sensors can improve safety by, for example, warning the captain of potential collisions. Nevertheless, the decision-making will mainly rely on the captain's ability to see and interpret their surroundings. Autonomous vessels, however, should be able to make sense of their surroundings and act upon

(a) A Mariner USV behind an Otter USV. Photo: Maritime Robotics.

(b) MS Estelle, an autonomous passenger ferry. Photo: Brødrene Aa.

Figure 1.1: Examples of unmanned and autonomous surface vessels.

them without human intuition. If we want to bridge the gap between the physical and the digital world, a key problem is how to make the digital system understand what is happening around it. An obvious approach to enable this understanding is to use sensors to measure the physical world and then process the data received from the sensors.

As part of a larger situational awareness framework, the target tracking system's contribution involves detecting and estimating the kinematic states of surrounding targets based on sensor information. Above and below the surface, target tracking is important when monitoring maritime activity. It helps to track the movements of ships and other vessels, providing information to coastal authorities and enhancing safety at sea. In subsea environments, target tracking is vital for monitoring submarines, identifying and categorizing underwater objects, and keeping a watchful eye on infrastructure [3, 81, 124]. Furthermore, target tracking is used for monitoring the airspace in the context of air traffic control [93], for tracking vehicles and people on land [55, 56], and much more.

The methods in this thesis focus on the maritime domain, and maritime data form the basis of the results. Because better autonomy at sea is the overarching goal of the Autosit project, the most relevant use case is ASVs. Figure 1.2 shows an example of an ASV system and outlines the role of target tracking in autonomous navigation. Roughly, the system works as follows. From the path planning module, the guidance and collision avoidance modules are informed of the next destination. Given a clear path, the guidance module will inform the control module of the desired speed and heading. However, if dynamic objects in the vicinity need consideration, the ASV can not simply follow the path. The target tracking module informs the collision avoidance module of the positions and velocities of the dynamic objects. The collision avoidance module uses the information to determine the collision risk and potential courses of action. If the ASV must make a maneuver, the collision

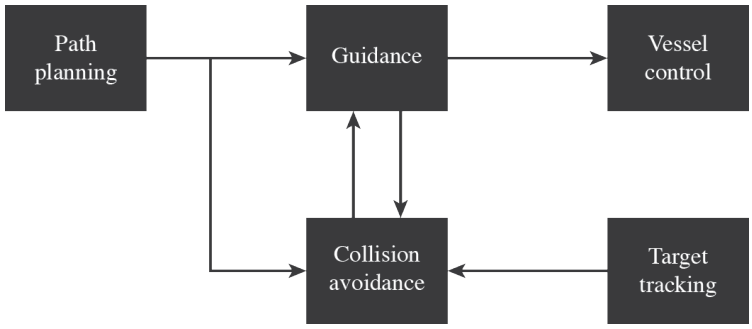

Figure 1.2: Overview of a system for autonomous navigation.

avoidance module tells the guidance module to change the desired speed and heading.

1.2 Exteroceptive sensors and target-provided information

When trying to establish an understanding of the dynamic objects inhabiting the area surrounding a vessel, the perhaps most common approach is to use exteroceptive sensors. Exteroceptive sensors placed on the vessel inform the system about what happens outside of the system itself. Proprioceptive sensors, on the other hand, measure internal states in the system. In maritime situational awareness applications, widely used exteroceptive sensors are radar, LiDAR, and cameras [62]. Furthermore, exteroceptive sensors can be either active or passive. Active sensors transmit signals and retrieve information by how the environment reacts to these signals. Radars and LiDARs are active sensors because they transmit radio and light waves, respectively. Cameras are passive sensors, as they only receive information from their environment. The methods in this thesis consider the use of active sensors, primarily radar and LiDAR.

Radar, whose name comes from *radio detection and ranging*, was developed and found its use throughout the first half of the 20th century [17]. Today, we find maritime radars on most boats above a certain size, and radars can be seen on both larger vessels pictured in Figure 1.1 in the form of white discs. Most maritime radars rotate at some fixed frequency, where each revolution corresponds to a radar scan. The radar output contains the range and bearing of the surfaces that reflect the transmitted radio waves. Both the range and bearing are limited in resolution; the resolution of the range decides the precision of the distance from the radar to the detection, and the bearing resolution decides the precision of the angle of the detection. The radar divides the surveillance area into several cells, in which detections can occur. If their view is unobstructed, radars can detect targets several tens of kilometers away. However, long ranges mean that the detections become less

precise. The detections are often clustered before they are used in the target trackers to provide a single position measurement for each target. The clustering avoids that a single target gives rise to several detections, depending on the number of sensor cells a target inhabits. The above is just a simplified explanation, sufficient to understand how target trackers use the measurements. Signal processing techniques such as beamforming make the inner workings of a radar significantly more complex, see, for example, [116].

Apart from operating with a different wave frequency, maritime LiDARs work conceptually similar to radars, with detections that fall within cells decided by the resolution in range and bearing. They also provide detections by sweeping the area. Whereas radars typically have relatively low resolution and update frequency, LiDARs have higher resolution and operate at a higher frequency. The trade-off is that LiDARs have a limited range, typically around a hundred meters. Due to their different strengths and weaknesses, LiDAR is well suited to complement radars. When operating in environments with small distances but high precision demands, LiDAR is a good choice. However, radar is a prerequisite in applications where targets must be detected well in advance at long ranges.

At sea and in the air, the targets themselves can provide information through standardized systems. In this thesis, such information is denoted as target-provided information. Many ships transmit messages through the Automatic identification system (AIS) [73]. The International Maritime Organization stipulates what has to be transmitted and which ships must transmit messages [72]. The messages can contain a wide range of information, such as position, speed, course, ship dimensions, and rate of turn. In addition, the messages contain the ID of the target. Aircraft provide similar information through the Automatic dependent surveillance-broadcast (ADS-B) protocol [39]. In both cases, on-board equipment is usually the source of the information.

The types of measurements, exteroceptive and target-provided, are inherently very different. Because exteroceptive sensors are placed at some distance from the targets, the precision of their measurements may vary. Furthermore, for LiDAR and radar, the provided information is often limited to the position and possibly speed of the reflecting surface. On the other hand, the system that utilizes the measurements usually knows how well the sensor works at any given time. Target-provided information is not dependent on the distance from the receiver to the transmitter, except for limitations in transmission range. The possible information in the messages is only limited to the messaging protocol, and what the target itself can measure. However, the transmission of the messages can be infrequent, and the receiver has no guarantee of the correctness of the messages. Furthermore, not all targets transmit any information. As such, it is unwise to use only target-provided information in situational awareness systems, but as a complement to exteroceptive

sensors, it can give invaluable insight.

1.3 Target tracking methods and concepts

The above sections explain the role of target tracking in situational awareness and the relevant sensors. However, making sense of the sensor information is not a trivial problem and has been the focus of significant research efforts. The early methods mainly considered filtering out the measurement noise, whereas later methods also considered artifacts such as false alarms and missed detections.

1.3.1 Filtering

The detections provided by the sensors will have some degree of inaccuracy relative to the true target positions. The inaccuracies, also called measurement noise, can be mitigated with filtering. The filtering step aims to solve the Bayesian filtering recursion given by the two equations

$$\begin{aligned} f_{k|k-1}(\mathbf{x}_k|\mathbf{z}_{1:k-1}) &= \int f_{\mathbf{x}}(\mathbf{x}_k|\tilde{\mathbf{x}}) f_{k-1}(\tilde{\mathbf{x}}|\mathbf{z}_{1:k-1}) \mathrm{d}\tilde{\mathbf{x}}, \\ f_k(\mathbf{x}_k|\mathbf{z}_{1:k}) &= \frac{f_{\mathbf{z}}(\mathbf{z}_k|\mathbf{x}_k) f_{k|k-1}(\mathbf{x}_k|\mathbf{z}_{1:k})}{\int f_{\mathbf{z}}(\mathbf{z}_k|\tilde{\mathbf{x}}) f_{k-1}(\tilde{\mathbf{x}}|\mathbf{z}_{1:k-1}) \mathrm{d}\tilde{\mathbf{x}}} \end{aligned} \tag{1.1}$$

where the first is the Champman-Kolmogorov equation, and the second is the Bayes update. The state the filter wants to estimate is $\mathbf{x}$, and the subscript indicates the time step of the state. For the densities f, the subscript $k|k-1$ indicates that the density is conditioned on the measurements up to the previous time step, and the subscript k indicates that the density is conditioned on the measurements up to the current time step. The measurements are denoted as $\mathbf{z}$, where the subscripts indicate from which time step the measurements originated. Furthermore, $f_{\mathbf{x}}(\cdot)$ is the state transition density, and $f_{\mathbf{z}}(\cdot)$ is the measurement likelihood density. The state transition density describes how the state evolves, and the measurement likelihood density describes how likely a measurement is given the state. Equation (1.1) is the basis of most of the methods presented in the following chapters.

Throughout this thesis, the Kalman filter [80] is the main tool used for filtering. For linear Gaussian models, it is an optimal closed-form solution to (1.1). The Kalman filter is a recursive method that estimates the state of a linear system based on noisy measurements. It is optimal in the sense that it minimizes the mean squared error of the estimates. The extended Kalman filter (EKF) is an extension of the Kalman filter that supports non-linear models. The EKF linearizes the system around the current estimate and then applies the Kalman filter to the linearized system. The

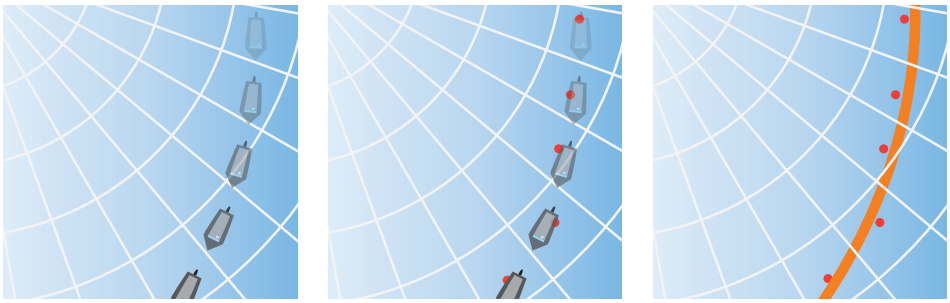

Figure 1.3: Use of a sensor to detect a target, and use of a target tracker to estimate the target state.

error from the linearization means that it is no longer optimal, but it is nevertheless a good option when the non-linearities are small. The unscented Kalman filter (UKF) is another extension of the Kalman filter, which instead of linearizing the system, uses a deterministic sampling approach to propagate the state distribution through the non-linear system. This results in more accurate estimates than the EKF can provide, especially for highly non-linear systems, but at a computational cost. Particle filters are another popular way of handling non-linearities. They represent the distributions as weighted particles, and the individual particles are propagated according to the system model. This approach can provide good estimates for highly non-linear systems, but the often large number of particles required for good performance can make it computationally expensive. See, for example, [130], for a more comprehensive description of the use of filters in target tracking.

Filtering solves the problem of noisy measurements but not the problem of false alarms, missed detections, and ambiguous associations between targets and measurements. The ability of target trackers to handle such problems is what distinguishes them from what we consider regular filters.

1.3.2 Target tracking

Figure 1.3 shows an area covered by a sensor, then the detections from five scans, and then an estimated trajectory of a hypothesized target from which the detections arose. In such a situtation, it is relatively easy to see how the target and the measurements correspond. The problem becomes more difficult when false alarms and missed detections are present, such as in Figure 1.4, where the wake of a ship has caused false alarms and the sensor is unable to detect the target for two out of the five scans. In addition, several targets can be present, and the association between the measurements and the targets may be ambiguous. To estimate the target states, the tracker must associate the hypothesized targets with their measurements, a problem

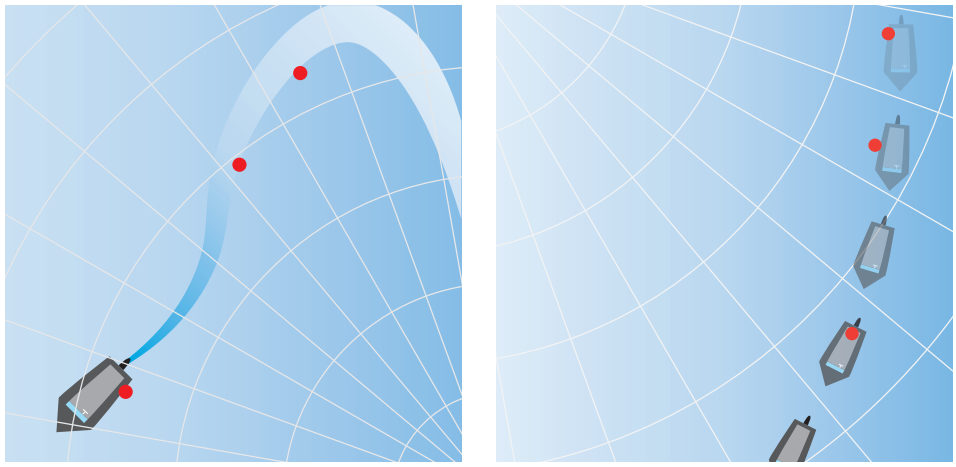

(a) Illustration of clustered radar detections with false alarms created by the wake of a boat.

(b) The same situation as in Figure 1.3, but with two missed detections.

Figure 1.4: False alarms and missed detections when tracking a boat.

known as data association. The fact that the number of targets is unknown, and may change over time, introduces further complications. Only when the measurements have been associated with a target can a filtering method such as the Kalman filter be used to estimate the target state. In (1.1), this amounts to finding the measurements in $\mathbf{z}_{1:k}$ that belong to a specific target state $\mathbf{x}_k$. The joint estimation of the number of targets and their states is denoted as multi-target tracking [130]. Several measurements can also originate from the same target, which requires either measurement clustering or a method that allows multiple measurements to be associated with the same target. The two approaches amount to what is known as point target tracking and extended object tracking, respectively. This section briefly explains common methods for point target tracking, whereas extended object tracking is not considered further until Chapter 4.

The probabilistic data association (PDA) filter [9] is one of the early methods made to solve the data association problem. It calculates the probability of association between the measurements and the target by considering the predicted position of the target, and comparing this to the received measurements while taking the uncertainty of the prediction and the measurements into account. In this way, it can account for the possibility of false alarms and missed detections.

Further development resulted in the joint PDA (JPDA) filter [41], which extended the PDA filter to account for the existence of multiple targets and the possibility of different measurements originating from different targets. Whereas the PDA filter assumes that only a single target is present, the JPDA filter assumes that some known

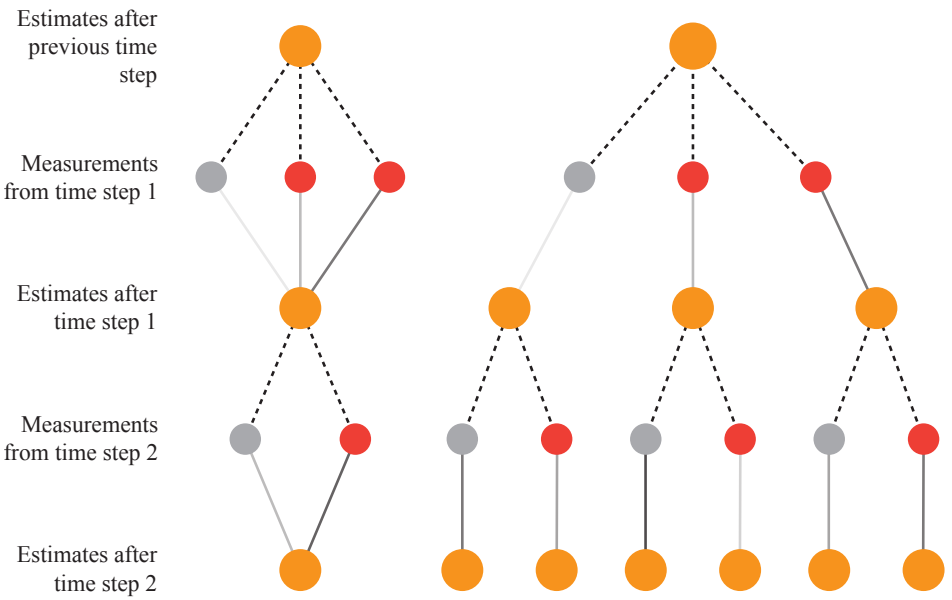

(a) A hypothesis structure similar to how they appear in PDA and its derived methods. After calculating the target estimate conditioned on the associations, the estimates are combined into a single estimate.

(b) A hypothesis structure similar to those of MHT and PMBM. A new target estimate is created for each possible association hypothesis, and the estimates are assigned a weight based on how likely the association hypothesis is.

Figure 1.5: A situation where one target estimate (the uppermost orange dot) was initialized at the previous time step. Two measurements are received (red dots) in the next time step and one in the time step after that. At each time, the track is either associated to one of the measurements or not associated to any measurement (the gray dots). The opacity of the lines indicates how likely the different hypotheses are.

number of targets are present. However, the actual number of targets is usually not known, and several methods have been developed to initialize tracks based on the received measurements [142]. One of the solutions is the integrated PDA (IPDA) filter [106] and its multi-target version, the joint integrated PDA (JIPDA) filter [105]. Here, target existence is not assumed, and the existence probability is calculated for each track that estimates a hypothesized target. Furthermore, the methods have been extended to work with the Interacting multiple models (IMM) filter, which uses several kinematic models for predicting the target state and switches between them depending on the estimated behavior of the target [15, 108].

For each time step, the PDA filter and its derived methods estimate the current state of a target based on the estimate from the previous time step and the measurements from the current time step. If the methods receive several measurements, their impact on the updated estimate depends on the association probability between the measurement and the estimate. If one represents the estimates as Gaussian distributions, this means that the updated estimate will be a Gaussian mixture with components representing the estimate conditioned on different measurements. Usually, this mixture is reduced to a single Gaussian before the next time step. This approximation ensures that the target trackers are computationally tractable but also means that temporal information is lost. Figure 1.5a shows an example of the hypothesis structure of the PDA filter and its derived methods over two time steps.

Multiple hypothesis tracking (MHT) [115] is a different approach to the data association problem. As in PDA, MHT considers all possible associations between the measurements and the estimates. For each possible association, the estimate is updated based on the measurement. However, the different combinations are used as the basis for new associations the next time new measurements arrive, and no mixture reduction takes place. This process results in a tree structure starting with the first detection and branches to the possible successive detections. A track hypothesis in the tree is a path from the root to one of the leaves, where said leaf is the estimated target state at the current time. Such a tree structure can be seen in Figure 1.5b. This way of solving the multi-target tracking problem is denoted as multi-scan tracking, as opposed to the single-scan methods described above. However, the tree structure grows exponentially and quickly becomes intractable to maintain. To mitigate this, branches of the hypothesis are pruned to remove the most unlikely hypotheses.

The Poisson multi-Bernoulli mixture (PMBM) filter is another multi-scan multi-target tracking method, that is based on the theory of Random finite sets (RFSs). It models the potentially detected targets as Bernoulli RFSs, and the not yet detected targets as a Poisson point process (PPP). As opposed to what is the case for MHT, this allows for more mathematically rigorous procedures for initialization and termination of tracks. Each time new measurements are received, all possible associations between the new measurements and the Bernoullis from the previous

time step give rise to new Bernoullis. Bernoullis representing potential new targets are created based on incoming measurements with a basis in a Poisson point process (PPP). As for the MHT algorithm, this results in an exponentially growing number of Bernoullis, and this is mitigated by a combination of pruning and only keeping a fixed number of possible global hypotheses. A global hypothesis is a set of Bernoullis which together form a complete set of possible associations between the measurements and the targets.

Furthermore, target trackers often utilize input from more than one sensor. When conducting target tracking with input from different sensors, a fusion scheme is needed to combine the information. One approach is to process the measurements from each sensor separately and then fuse the resulting estimates. This method is known as track-level fusion. By considering the individual sensors in isolation, the problem only lies in the fusing of the estimates from the different sensors. This task is complicated by the unknown correlation between estimates and the need for associating the estimates from the different sensors with each other. Nevertheless, methods such as covariance intersection [78] have proven robust when fusing estimates, and the data association problem is usually relatively easy to solve if the individual estimates are of high quality. However, the loss of information introduced by the separate processing of the measurements is a theoretical limitation in any track-level fusion scheme. Furthermore, when data from a sensor are received with long intervals between updates, it can struggle to produce good estimates on its own. This is often the reality for target-provided information. Another widely applied approach is measurement-level fusion. In measurement-level fusion, the tracker updates its estimates based on measurement data from different sensors. This is, theoretically, an approach that allows for better estimates, but the different measurements have to be carefully modeled. The methods must consider how to handle differences in, for example, update frequencies, detection probabilities, and false alarm rates.

1.3.3 Related work on the fusion of exteroceptive sensor measurements and target-provided information

Several methods for measurement-level fusion of target-provided information and exteroceptive measurements have been proposed in recent years. In [59], Habtemariam et al. propose a measurement-level fusion method for combining AIS and radar measurements in the context of the JPDA filter. The authors account for the transmitted ID information and present a method for applying the IDs to targets using Bayesian inference. Furthermore, they consider the physical nature of AIS messages, with their infrequent transmission and absence of false alarms. The results show the benefit of utilizing AIS messages and the increased performance of the measurement-level fusion approach relative to track-level fusion approaches.

A somewhat different approach is found in [44], where Gaglione et al. present a tracker that utilizes belief propagation and a particle filter. The problem is formulated as a factor graph, with calculations consisting of passing messages between nodes in the graph. The work also considers the initialization of tracks using both exteroceptive and target-provided measurements, in addition to the time discrepancy between the different measurement types. The method estimates the ID information probabilistically and allows for handling incorrect ID information.

Furthermore, MHT and the PMBM filter have previously been used together with target-provided information. Liland describes an AIS-guided MHT in [94]. Here, the AIS information is modeled as if it was provided by another exteroceptive sensor, with artificial models for false alarms and missed detections. The transmitted IDs are assumed to always be correct, and any errors are presumed removed in the pre-processing stage. The method improves the tracking performance relative to an MHT without AIS information. In [102], Miao et al. present a method for using AIS information in the PMBM filter. They take a similar approach to the MHT variant and model the AIS information similarly to exteroceptive measurements. Furthermore, the transmitted IDs are not explicitly considered in the calculations but are used to label the tracks corresponding to AIS-transmitting targets.

1.4 Research objectives

The Autosit project is a continuation of the Autosea project [26], a previous research project at NTNU that looked into autonomous operations at sea. As such, many of the research objectives for this thesis are inspired by the findings of that project. That includes specific problems encountered during the experimental testing in the project and more general problems, such as the development of target tracking methods that better utilize AIS information. The research objectives are condensed into the following goals:

1. Provide solutions to problems encountered in the maritime domain regarding occlusions, target maneuvering, wake clutter, and potential time offset effects.
2. Use of target-provided information in target tracking:
   - Create a model that allows for the use of target-provided measurements in target tracking.
   - Develop methods that utilize target-provided measurements in target tracking.
3. Use of the developed methods in real-life applications:
   - Evaluate the demands for the safe use of AIS in maritime target tracking.

- Validate the developed methods in field tests.

The goals are divided into three groups, where the first demands research into more general problems encountered in the maritime domain for several different sensors and platforms. The second group considers the inclusion of target-provided measurements in target trackers. Both demand the development of robust target trackers suitable for use in maritime environments. The third involves using the methods in real-world experiments.

1.5 Contributions and outline

The thesis is built up of chapters corresponding to research papers, with a structure that roughly follows the research objectives. There are three parts: The first considers the development of target trackers for the maritime domains, with and without the use of target-provided information. The second part expands upon how to use the methods from the previous chapters in autonomous marine operations. The last part catches up with some of the problems that can be encountered in maritime target tracking when using exteroceptive sensors, that were not covered explicitly in the first part. A summary of each chapter, with their respective publications listed, follows.

Part I: Including Target-Provided Measurements in Multi-Target Tracking

Chapter 2: Multi-Target Tracking With Multiple Models and Visibility

Publication:

[24] E. F. Brekke, A. G. Hem, and L.-C. N. Tokle. Multitarget Tracking With Multiple Models and Visibility: Derivation and Verification on Maritime Radar Data. *IEEE Journal of Oceanic Engineering*, 46(4):1272–1287, July 2021

We demonstrate how a variation of the joint integrated probabilistic data association (JIPDA) filter with interacting multiple models (IMM) and a visibility state can be derived as a special case of the Poisson multi-Bernouli mixture (PMBM) filter with a hybrid state representation and standard approximations. The proposed method is tested on two radar data sets which were recorded during maritime collision avoidance experiments.

Chapter 3: Variations of Joint Integrated Probabilistic Data Association With Radar and Target-Provided Measurements

Publication:

[68] A. G. Hem and E. F. Brekke. Variations of joint integrated data association with radar and target-provided measurements. *Journal of Advances in Information Fusion*, 17(2):97–115, Dec. 2022

We present a multi-target tracking algorithm utilizing target-provided information using the framework of joint integrated probabilistic data association (JIPDA). The use case we consider is maritime target tracking using radar measurements combined with messages from the Automatic identification system (AIS). The full details of the tracking algorithm are presented, including implementation-specific considerations to account for the different natures of the incoming measurements. We detail three different methods of handling the target-provided measurements, one processing them as they arrive, i.e., sequentially, and the others collecting and processing them at fixed intervals. The results show that all improve over the pure radar tracking algorithm.

Chapter 4: Poisson Multi-Bernoulli Mixture Filtering With Fusion of Target-Provided and Exteroceptive Measurements

Publication:

[65] 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. To appear in *IEEE Access*. doi: 10.1109/ACCESS.2024.3389824

We present a method for including target-provided measurements in the Poisson multi-Bernoulli mixture (PMBM) filter, both when using it to track extended objects and point targets. We use messages from the Automatic identification system as an example of target-provided measurements, and radar and LiDAR as examples of exteroceptive sensors. In the point target case, we utilize several different kinematic models in parallel through the interacting multiple models framework, and compare the presented method to several common trackers and other PMBM filter configurations. The results show that our method outperforms similar methods when target-provided measurements are available. For the extended object tracking case, we expand upon the Gaussian process PMBM filter. The extended object method is evaluated on both simulated and experimental data, and is shown to improve the tracking performance when including target-provided measurements in comparison to when it only uses exteroceptive measurements.

Part II: Use of AIS in Maritime Autonomy Applications

Chapter 5: Autonomous Marine Collision Avoidance with Sensor Fusion of AIS and Radar

Publication:

[70] A. G. Hem, E. F. Brekke, G. D. K. M. Kufoalor, and I. H. Kingman. Autonomous Marine Collision Avoidance With Sensor Fusion of AIS and Radar. Submitted to the 15th IFAC Conference on Control Applications in Marine Systems, Robotics and Vehicles (CAMS 2024)

We present the results of a series of autonomy experiments conducted in the Trondheimsfjord, Norway, to evaluate novel target tracking methods that use both exteroceptive sensors and messages from the Automatic identification system (AIS). The experiments used a 6-meter-long autonomous surface vessel equipped with a radar and an AIS antenna, in combination with two target vessels. The tracker output was used to inform a collision avoidance method about the surrounding targets, which allowed it to operate in the intended safe and regulation-abiding manner. The results show that the tracker provided target estimates of high enough quality to be used directly in the collision avoidance method without post-processing, and that the surface vessel was able to navigate safely through the scenarios.

Chapter 6: Validation of AIS Information With Exteroceptive Sensor Fusion in Autonomous Operations

Publication:

[66] A. G. Hem and E. F. Brekke. Validation of AIS Information With Exteroceptive Sensor Fusion in Autonomous Operations. To appear in *IEEE Intelligent Transportation Systems Magazine*. doi: 10.1109/MITS.2024.3389869

The Automatic identification system (AIS) can improve situational awareness at sea, but its protocol is simple and does not guarantee message integrity, authentication, and proper use. The lack of safety measures creates problems when AIS messages are used for, for example, tracking a target or predicting a target trajectory. We present a methodology for validation of AIS messages, a prerequisite for their safe use in maritime situational awareness applications. The validation method relies on target trackers which fuse AIS data and exteroceptive sensor data, and detects errors in position, speed and course, and rate of turn. By use of simulated data, we show that the proposed methods effectively detect errors in the position and velocity data received through AIS messages and is also able to detect errors in turn rate

data. The effectiveness of the methods is demonstrated on a real-world dataset with injected false AIS data.

Part III: Selected Topics in Target Tracking With Exteroceptive Sensors

Chapter 7: Compensating Radar Rotation in Target Tracking

Publication:

[67] A. G. Hem and E. F. Brekke. Compensating radar rotation in target tracking. In *Proc. Sensor Data Fusion: Trends, Solutions, Applications*, Bonn, Germany, Oct. 2022

For applications where the vessel carrying the radar is highly maneuverable, or the radar rotation frequency is low, it may be unreasonable to approximate the radar measurements as arriving simultaneously. This chapter identifies effects caused by radar rotation and shows that these effects can degrade tracking performance. We present methods for mitigating the undesirable effects and apply these to a joint integrated probabilistic data association (JIPDA) tracker. Results on simulated data show that the proposed methods give significant performance benefits in certain situations.

Chapter 8: Target Tracking With Existence Modeling in the Presence of Wakes

Publication:

[64] A. G. Hem, H.-G. Alvheim, and E. F. Brekke. WakeIPDA: Target Tracking With Existence Modeling in the Presence of Wakes. In *Proc. IEEE 26th International Conference on Information Fusion*, Charleston, SC, USA, June 2023

We present a novel target tracking algorithm, which is designed to track a target in the presence of wake clutter. What distinguishes the method from previous wake-compensating trackers is that it also models the existence of the target and exploits the information provided by the wake measurements for this purpose. We present two ways of modeling the wake, and we evaluate the algorithm's performance on simulated data. Results show that the method improves upon comparable target tracking methods when wake measurements are present.

Co-authored papers

In addition to the papers listed above and included in this thesis, I have also contributed to the following papers:

[121] J. Å. Sagild, A. G. Hem, and E. F. Brekke. Counting technique versus single-time test for track-to-track association. In *Proc. IEEE 24th International Conference on Information Fusion*, Sun City, South Africa, Nov. 2021

[22] E. F. Brekke and A. G. Hem. A long simulation scenario for evaluation of multi-target tracking methods. In *Proc. 3rd International Conference on Electrical, Computer, Communications and Mechatronics Engineering (ICECCME)*, Tenerife, Canary Islands, Spain, July 2023

[5] M. Baerveldt, A. G. Hem, and E. F. Brekke. Comparing Multiple Extended Object Tracking with Point Based Multi Object Tracking for LiDAR in a Maritime Context. *Journal of Physics: Conference Series*, 2618(1):012011, Oct. 2023

[43] D. Gaglione, P. Braca, G. Soldi, F. Meyer, A. G. Hem, E. F. Brekke, and F. Hlawatsch. Comments on "Variations of Joint Integrated Data Association with Radar and Target-Provided Measurements". *Journal of Advances in Information Fusion*, 18(2):93–101, Dec. 2023

[23] E. F. Brekke, A. G. Hem, and L.-C. N. Tokle. The VIMMJIPDA: Hybrid state formulation and verification on maritime radar benchmark data. In *Proc. IEEE OCEANS*, Virtual conference, Oct. 2020

Part I

Including Target-Provided Measurements in Multi-Target Tracking

Chapter 2

Multi-Target Tracking With Multiple Models and Visibility

This chapter is based on the publication:

[24] E. F. Brekke, A. G. Hem, and L.-C. N. Tokle. Multitarget Tracking With Multiple Models and Visibility: Derivation and Verification on Maritime Radar Data. *IEEE Journal of Oceanic Engineering*, 46(4):1272–1287, July 2021

Changes from the original publication involve shortening of the introduction section. My work in the chapter mainly involves implementation, result acquisition, analysis, and the writing regarding these topics.

2.1 Introduction

Target tracking, including track initialization and track termination, for autonomous ship situational awareness poses high requirements to both accuracy and reliability, which straighforward joint probabilistic data association (JPDA) or multiple hypothesis tracking (MHT) may not be capable of fulfilling. It is important to establish tracks as fast as possible to enable early action in collision avoidance (COLAV) scenarios. Furthermore, track-loss is highly undesirable, because it can take several seconds to re-establish a track with typical off-the-shelf radars. Velocity and course estimates are likely to suffer from significant uncertainties. When tracking a maneuvering recreational boat, the root mean square error (RMSE) of the course estimate is likely to reach values around 10°, and may occasionally hit 45° or even 90°, leading to significant confusion in a COLAV system [35].

These challenges lead to three extensions of the basic JPDA formalism. First, modeling of track existence can yield significantly faster and more reliable track

initialization [142]. This leads to joint integrated probabilistic data association (JIPDA). Second, using multiple motion models in parallel improves the accuracy of velocity estimation. This leads to interacting multiple model JIPDA (IMMJIPDA) [108]. IMM-JPDA techniques have previously been used in the COLAV system reported in [123]. Third, including a model of target visibility in addition to the existence model can make the tracker less susceptible to track-loss [141]. We refer to the resulting tracking method as a VIMMJIPDA.

Modern theory of multi-target tracking is heavily influenced by finite set statistics (FISST) and its key concepts such as random finite sets, set densities, probability generating functionals (p.g.fl.'s) and the multi-target Bayes filter [98]. The structure inherent in the multi-target Bayes filter can, under standard assumptions, be exploited to derive a solution with finite computational requirements, which has become known as the PMBM filter [138]. It was shown in [138] that the PMBM filter reduces to a more refined version of the JIPDA when similar approximations are being made (that is, merging of all association hypotheses after every estimation cycle). This refined filter was termed track oriented multi Bernoulli/Poisson (TOMB/P) in [138] and is in the broader class of Poisson multi-Bernoulli (PMB) filters.

The first contribution of the present chapter is to strengthen the PMBM-JIPDA link by also establishing versions of IMMJIPDA with basis in the PMBM filter. This is done by introducing a hybrid continuous-discrete state vector that accounts for all discrete-natured uncertainty except for existence uncertainty, which on the other hand is accounted for by the random finite set framework. This differs from how existence and visibility were treated in the seminal IPDA paper [107], where both were considered as different state values in a common Markov Chain.

A second contribution of the chapter is to publish two benchmark data sets for maritime radar tracking which recently have been recorded as part of the Autosea project [26], and to verify the proposed algorithm on these data sets. There is in general a scarcity of publicly available data sets for verification of multi-target tracking methods when compared to related fields such as simultaneous localization and mapping (SLAM), where several established data sets are regularly used for verification of new algorithms (see for exampleSection 5 in [79]). In particular, the authors are not aware of any public benchmark data set for maritime radar tracking. The closest would be the *simulated* benchmark data reported in [29]. In the future, the Stone Soup framework aims to enable users to set up experiments with real-world data [90] [12], but currently it does not contain any data set comparable to those we present in this chapter.

Evaluation of tracking methods is typically focused on performance measures such as variations of the Optimal sub-pattern assignment (OSPA) metric [122] [113] or various measures of track quality (track lifetime, track fragmentation rate, etc.) [109]. Most of these measures rely on a ground truth, and report some kind of average

behavior, which is hard to assess except through extensive Monte-Carlo simulations. In contrast we take a more qualitative approach in this chapter. We study real data with the purpose of gaining a better understanding of likely causes of trouble that can arise in a multi-target tracking system used for maritime COLAV, rather than conducting a systematic evaluation using standard performance measures.

The chapter is organized as follows. In Section 2.2 we present some minimal background on p.g.fl. and FISST, and the notation that will be used. In Section 2.3 we present the hybrid state model as a special case of standard assumptions in multi-target tracking. In Section 2.4 we derive the VIMMJIPDA. In Section 2.5 we provide details on the implementation of the tracking method. Section 2.5.2 presents the data sets used and tracking results. Finally, a conclusion follows in Section 2.6.

2.2 Notation and terminology

We will address the PMBM filter from the p.g.fl. perspective that was used in [138]. The reader is referred to [98] for a standard reference on p.g.fl.'s in multi-target tracking. The main rationale for this approach is that it arguably leads to neater mathematics than the alternative formulation in terms of set densities, which has been used in for example [47]. The price to pay for this is a higher level of abstraction. In general, a p.g.fl. $G[h]$ is a machine that takes test functions h as arguments, and returns a number for every test function. If $f(\mathbf{x})$ is a function, we define the corresponding linear functional by

$$f[h] = \int f(\tilde{\mathbf{x}})h(\tilde{\mathbf{x}})\,\mathrm{d}\tilde{\mathbf{x}}. \tag{2.1}$$

This notation will frequently be used to describe inner products of functions. Also notice that we consistently use a tilde notation for variables that are to be integrated out as part of such an expression. These notations follow the general notational framework laid out in [25].

Whenever $\tilde{\mathbf{x}}$ is a hybrid continuous-discrete variable, the infinitesimal measure $\mathrm{d}\tilde{\mathbf{x}}$ above is adopted to be the product measure between the Lebesgue measure and counting measure[1] [40] for the continuous and discrete parts, respectively. Integrals with respect to the counting measure become summations, so that the above can equivalently be written

$$f[h] = \sum_{\tilde{\mathbf{x}}^{\mathrm{d}}} \int f(\tilde{\mathbf{x}})h(\tilde{\mathbf{x}})\,\mathrm{d}\tilde{\mathbf{x}}^{\mathrm{c}}, \tag{2.2}$$

[1]Measure is the mathematical term for function that takes a set and gives back its size, or "volume" if you like, and which additionally obeys certain rules in order to be well behaved.

where $\mathbf{x}^{\mathrm{d}}$ and $\mathbf{x}^{\mathrm{c}}$ refers to the discrete part and continuous parts, respectively. See also Section II.B in [97].

This generalizes to functions of sets $X = \{\mathbf{x}^1, \ldots, \mathbf{x}^n\}$, with

$$G[h] = \int f(\tilde{X}) \prod_{\tilde{\mathbf{x}} \in \tilde{X}} h(\tilde{\mathbf{x}}) \, \delta\tilde{X} \tag{2.3}$$

$$= \sum_{n=0}^{\infty} \frac{1}{n!} \int f(\{\tilde{\mathbf{x}}^1, \ldots, \tilde{\mathbf{x}}^n\}) \prod_{i=1}^{n} h(\tilde{\mathbf{x}}^i) \, \mathrm{d}\tilde{\mathbf{x}}^1 \cdots \mathrm{d}\tilde{\mathbf{x}}^n \tag{2.4}$$

where the set integral [98, §11.3.3] in the first line is naturally taken with respect to the set measure written out and simplified to integrals on the state space in the second line [131, §VII].

We will occasionally talk about intensities of Poisson processes. Such an intensity is a function $f(\mathbf{x})$ such that the expected number of objects found inside a region S is given by $\int_S f(\mathbf{x})\mathrm{d}\mathbf{x}$, ie. a scaled pdf. A random finite set for which this is true and the number of objects found within all pairwise disjoint regions is independent, is a Poisson process. The number of objects in such a process is Poisson distributed with their states independently identically distributed (i.i.d.). Its p.g.fl. is of the form

$$G^{\mathrm{ppp}}[h] = \exp(\lambda[h-1]) \tag{2.5}$$

where $\lambda(\mathbf{x})$ is its intensity.

We will also talk about Bernoulli, multi-Bernoulli (MB) and MB mixtures (MBM) p.g.fl.'s. The Bernoulli comes from a set that has a Bernoulli trial for being either empty or singleton and no other outcomes possible on the number of elements. The trial has probability r for being singleton, in which case the element has some state pdf f. The MB comes from having N independent, not necessarily identical, such sets, and the MBM from having a mixture of MB's. The MBM p.g.fl. is

$$G^{\mathrm{mbm}}[h] = \sum_{j=1}^{M} p^j \prod_{i=1}^{N} (1 - r^{i,j} + r^{i,j} f^{i,j}[h]), \tag{2.6}$$

where p^j is the probability that some latent hypothesis j is correct, with $\sum_{j=1}^{M} p^j = 1$. Under hypothesis j, $r^{i,j}$ and $f^{i,j}$ are the singleton probabilities and state pdf's of the i'th Bernoulli in the MB, respectively. This reduces to an MB whenever $M = 1$ and $N > 1$ (the summation disappears and $p^1 = 1$, making the superscript j unnecessary), and to a Bernoulli when $N = 1$ and $M \in \mathbb{N}$ (the product is removed, leaving a single factor that renders the superscript i unnecessary).

The random finite set framework hinges on the fundamental theorem of multi-object calculus [98, p. 384] which states that the representation can be transformed

back and forth between set densities and p.g.fl.'s by means of integration and differentiation. The proof of this theorem can be found in [146] and relies on the property that delta functions obey $\int \delta_{\mathbf{x}_0}(\tilde{\mathbf{x}}) f(\tilde{\mathbf{x}}) \, \mathrm{d}\tilde{\mathbf{x}} = f(\mathbf{x}_0)$. A combination of Dirac's and Kroenecker's delta for the continuous and discrete parts, respectively, satisfies this, and the fundamental theorem continues to hold for the hybrid spaces considered here. More generally, the hybrid spaces considered here are locally compact, Hausdorff, and second countable (LCHS), as required by standard FISST theory [100].

Remark 2.1. What is a JIPDA and what is a PMB filter? The distinction is not immediately clear because the main working principle for both is the approximation of a MBM of the form (2.6) with a MB. One distinction between the two methods is largely a matter of historical origin: Musicki's term JIPDA was used in the (non-FISST) literature until around 2015, when researchers building on Williams' work started to talk about the PMB filter. A second distinction is that a PMB might use other types of mixture reduction where the notion of a track identity in the JPDA sense is lost, eg. [140].

To suggest a more precise terminology, we may call the method a variation of JIPDA if it 1) does *not* have a dynamic Poisson component, 2) approximates the joint distribution over the track to measurement association variables with its (approximate) marginals, and 3) uses moment matching for mixture reduction of the single target pdf's to arrive at the MB form. Thus, we will refer to our method as a variation of JIPDA.

In relation to this question one can also bring up the issue of track identities and track labels: Most researchers familiar with JIPDA will probably think of the tracks in a JIPDA as having distinct identities, while this is less clear for a PMB filter. The reader may consult references such as [132], [57], [20] and [126] for different viewpoints in this debate, which is beyond the scope of this chapter. Although the derivations here are based on the PMBM filter with the track oriented MB approximation, one can arrive at the same equations by adapting the IMM-JIPDA of [108]. From this point of view, the reader can interpret identities as they do for the JIPDA, if preferred. ◁

2.3 Model assumptions

The PMBM filter is a direct solution of the multi-target Bayes filter under the assumption of no track-before-detect, no merged measurements and no extended objects. This *standard model* of multitarget tracking can be summarized by the following general assumptions:

M1 New targets are born according to a Poisson process with intensity $b(\mathbf{y})$.

M2 Existing targets survive from time step $k-1$ to k with probability $P_{\mathrm{S}}(\mathbf{y}_{k-1})$.

M3 The motion of a surviving target is given by $f_{\mathbf{y}}(\mathbf{y}_k|\mathbf{y}_{k-1})$.

M4 A target with state $\mathbf{y}_k$ generates a measurement $\mathbf{z}_k$ with probability $P_{\mathrm{D}}(\mathbf{y}_k)$.

M5 Clutter measurements occur according to a Poisson process with intensity $\lambda(\mathbf{z})$.

M6 The measurement of a detected target is related to the state according to $f_{\mathbf{z}}(\mathbf{z}_k|\mathbf{y}_k)$.

In addition, independence is assumed whenever it is not contradicted by M1-M6.

In this chapter, the state vector is of the form

$$\mathbf{y}_k = \begin{bmatrix} \mathbf{x}_k \\ s_k \\ \mathrm{v}_k \end{bmatrix} \tag{2.7}$$

where $\mathbf{x}_k$ is the kinematic state vector, s_k is a discrete-valued random variable that attains one of M possible models, and v_k is a binary random variable that attains 1 if the target is visible, and otherwise attains 0. Based on this, we re-write assumptions M1-M6 in more elaborate terms taking the hybrid state into account.

M1: The Birth model

We assume the birth intensity to be stationary, and of the form

$$b(\mathbf{y}) = B V_\Omega f_b(\mathbf{x}) \mu^{0s} o^{0\mathrm{v}} \tag{2.8}$$

$$f_b(\mathbf{x}) = \frac{1_\Omega(\mathbf{H}^{(s)}\mathbf{x})}{V_\Omega} \mathcal{N}(\mathbf{H}^{*(s)}\mathbf{x}\,;\, \mathbf{0}, \mathbf{P}_v^{(s)}) \tag{2.9}$$

where B is the single parameter quantifying the overall rate in terms of *expected number of objects per unit position volume* entering per time step. The notation Ω denotes the surveillance region, while $1_\Omega(\cdot)$ denotes the corresponding indicator function, which is one everywhere in Ω and otherwise zero. This region has volume V_Ω. The prior probabilities μ^{0s} and $o^{0\mathrm{v}}$ are of the model state s and the visibility state v, respectively. We use the notation $\eta^0 = o^{01}$ to denote the a priori probability that an unknown target is visible (see Assumption M4* and Section 2.4.2). The matrix $\mathbf{P}_v^{(s)}$ is the tunable a priori velocity covariance, typically of the form $\mathbf{I}\sigma_v^2$. The matrix $\mathbf{H}^{(s)}$ picks out the positional components of $\mathbf{x}$, while $\mathbf{H}^{*(s)}$ picks out all other components of $\mathbf{x}$.

M2: The survival probability

We use a constant survival probability P_{S}.

M3: The single-target Markov model

Instead of the purely kinematic Markov model typically used in the tracking literature, we use a more complex state evolution model of the form

$$f_{\mathbf{y}}(\mathbf{y}_k \,|\, \mathbf{y}_{k-1}) = f_{\mathbf{x}}^{s_k}(\mathbf{x}_k \,|\, \mathbf{x}_{k-1}) \pi^{s_{k-1} s_k} \omega^{\mathrm{v}_{k-1} \mathrm{v}_k}. \tag{2.10}$$

Here $\pi^{s_{k-1} s_k}$ and $\omega^{\mathrm{v}_{k-1} \mathrm{v}_k}$ are probabilities from the Markov chain transition matrices of the model and visibility states, respectively. We assume that the kinematic Markov model is given by a Gaussian

$$f_{\mathbf{x}}^{s_k}(\mathbf{x}_k \,|\, \mathbf{x}_{k-1}) = \mathcal{N}(\mathbf{x}_k \,;\, \mathbf{f}^{(s_k)}(\mathbf{x}_{k-1}), \mathbf{Q}^{(s_k)}). \tag{2.11}$$

For the sake of generality, the Gaussian assumption will only be explicitly invoked when necessary.

M4: The detection probability

The detection probability is state dependent in accordance with

$$P_{\mathrm{D}}(\mathbf{y}) = \begin{cases} P_{\mathrm{D}} & \text{if } \mathrm{v} = 1 \\ 0 & \text{if } \mathrm{v} = 0. \end{cases} \tag{2.12}$$

We say that the target is visible whenever $\mathrm{v} = 1$, although such a target may still have missed detections.

M5: The clutter intensity

Clutter is uniformly distributed according to a Poisson process with constant intensity λ. We remark that generalization to measurement-dependent clutter intensity is straightforward.

M6: The kinematic likelihood

The likelihood of a measurement conditional on the target state can be written

$$f_{\mathbf{z}}(\mathbf{z} \,|\, \mathbf{y}) = f_{\mathbf{z}}^{s}(\mathbf{z} \,|\, \mathbf{x}) = \mathcal{N}(\mathbf{z} \,;\, \mathbf{H}^{(s)}\mathbf{x}, \mathbf{R}^{(s)}). \tag{2.13}$$

It is only defined when $\mathrm{v} = 1$. It will never be used when $\mathrm{v} = 0$ because in such a case we will not have any detection according to M4 above. Notice that the measurement matrix $\mathbf{H}^{(s)}$ is the same as the matrix used in (2.9).

2.4 From PMBM to VIMMJIPDA

The main result of [138] states that both the predicted and the posterior p.g.fl. can be factorized into a Poisson p.g.fl. and an MBM p.g.fl. Omitting temporal subscripts, this factorization can be written

$$G[h] = G^{\mathrm{ppp}}[h]G^{\mathrm{mbm}}[h]. \tag{2.14}$$

A key idea of the PMBM filter is that targets that exist, but have not been observed, constitute a set of *unknown targets* whose random properties can be quantified by the Poisson factor in (2.14) through a corresponding unknown target intensity [138]. In Section 2.4.1 we will see that it is reasonable to take this intensity as spatially uniform and stationary in the VIMMJIPDA.

The MBM component $G^{\mathrm{mbm}}[h]$ is a mixture over a collection of association hypotheses, and for every hypothesis it is a product of un-normalized Bernoulli p.g.fl.'s given by hypothesis weight contributions, existence probabilities (possibly zero) and state pdf's (possibly left undefined for Bernoulli p.g.fl.'s with zero probability of existence) for every hypothetical track. We postpone the detailed exposition of this structure to Section 2.4.4. It was shown in [138] that the PMBM form is closed under the Bayes recursion, although the number of hypotheses will grow exponentially and the number of Bernoulli components will grow linearly with time.

2.4.1 The unknown target intensity

From [138], we know that the Poisson component, that isunknown target intensity, here denoted by the scaled pdfs $u_k(\mathbf{y})$ and $u_{k|k-1}(\mathbf{y})$, is measurement updated and time predicted according to

$$u_k(\mathbf{y}) = (1 - P_{\mathrm{D}}(\mathbf{y}))u_{k|k-1}(\mathbf{y}) \quad \text{and} \tag{2.15}$$

$$u_{k|k-1}(\mathbf{y}) = b(\mathbf{y}) + u_{k-1}[f_{\mathbf{y}}(\mathbf{y} \mid \tilde{\mathbf{y}})P_{\mathrm{S}}(\tilde{\mathbf{y}})] \tag{2.16}$$

respectively.

We do not want the spatial parameterization of $u_k(\mathbf{y})$ to be any more complicated than that of $b(\mathbf{y})$. For this reason we make the approximation that (2.9) is not altered by a prediction using the kinematic Markov model. This means that multiplying by (2.11) and integrating does not change its form, so that

$$f_b(\mathbf{x}) \approx f_b[f^s_{\mathbf{x}}(\mathbf{x} \mid \tilde{\mathbf{x}})] = \int f^s_{\mathbf{x}}(\mathbf{x} \mid \tilde{\mathbf{x}}) f_b(\tilde{\mathbf{y}}) \, \mathrm{d}\tilde{\mathbf{x}}. \tag{2.17}$$

For this assumption to be reasonable, two approximations must hold. 1) The velocity uncertainty must not increase noticeably between the time a target arrived and

the time it was first observed[2], and 2) The effect of the kinematic Markov model predicting objects to leave Ω is relatively small.

Proposition 2.1. For stationary birth and a Markov transition model where all states have a positive (that is, non-zero) probability of reaching a region of the state space where $P_{\mathrm{D}}(\mathbf{y}) > 0$ and/or $P_{\mathrm{S}}(\mathbf{y}) < 1$ in finite time, the intensity $u_{k|k-1}(\mathbf{y})$ will converge, and furthermore under (2.17) and M1* – M4* to

$$u(\mathbf{y}) = \lim_{k\to\infty} u_{k|k-1}(\mathbf{y}) = UV_\Omega f_u(\mathbf{x})\mu_u^s o_u^{\mathrm{v}}, \tag{2.18}$$

with overall intensity

$$U = B + UP_{\mathrm{S}}(1-\eta_u P_{\mathrm{D}}) = \frac{B}{(1-P_{\mathrm{S}}(1-\eta_u P_{\mathrm{D}}))}, \tag{2.19}$$

kinematic state pdf

$$f_u(\mathbf{x}) \approx f_b(\mathbf{x}) = \frac{1_\Omega(\mathbf{H}^{(s)}\mathbf{x})}{V_\Omega}\mathcal{N}(\mathbf{H}^{*(s)}\mathbf{x}\,;\,\mathbf{0},\mathbf{P}_v^{(s)}), \tag{2.20}$$

mode probabilities

$$\boldsymbol{\mu}_u = \frac{B}{U}\boldsymbol{\mu}^0 + P_{\mathrm{S}}(1-P_{\mathrm{D}}\eta_u)\boldsymbol{\pi}^\top\boldsymbol{\mu}_u, \tag{2.21}$$

$$= \frac{B}{U}(\mathbf{I} - P_{\mathrm{S}}(1-P_{\mathrm{D}}\eta_u)\boldsymbol{\pi}^\top)^{-1}\boldsymbol{\mu}^0, \tag{2.22}$$

and visibility state probabilities

$$\begin{bmatrix} o_u^1 \\ o_u^0 \end{bmatrix} = \frac{B}{U}\begin{bmatrix} o^{01} \\ o^{00} \end{bmatrix} + P_{\mathrm{S}}\begin{bmatrix} \omega^{11} & \omega^{01} \\ \omega^{10} & \omega^{00} \end{bmatrix}\begin{bmatrix} 1-P_{\mathrm{D}} & 0 \\ 0 & 1 \end{bmatrix}\begin{bmatrix} o_u^1 \\ o_u^0 \end{bmatrix} \tag{2.23}$$

$$= \frac{B}{U}\left(\mathbf{I} - \begin{bmatrix} P_{\mathrm{S}}\omega^{11}(1-P_{\mathrm{D}}) & P_{\mathrm{S}}\omega^{01} \\ P_{\mathrm{S}}\omega^{10}(1-P_{\mathrm{D}}) & P_{\mathrm{S}}\omega^{00} \end{bmatrix}\right)^{-1}\begin{bmatrix} o^{01} \\ o^{00} \end{bmatrix}, \tag{2.24}$$

where $\eta_u = o_u^1$ is defined, and the mode and visibility state probability equations are in matrix form. Here, $\boldsymbol{\mu}_u$ and $\boldsymbol{\mu}^0$ are the column vectors with μ_u^s and μ^{0s} as its s'th elements, respectively, and $\boldsymbol{\pi}$ is the matrix with $\pi^{\tilde{s}s}$ at the $\tilde{s}$'th row and s'th column.

Proof. See Appendix 2.6. □

VIMMJIPDA, therefore, directly specifies and uses the converged undetected intensity instead of the birth intensity. This leaves the overall rate U, initial mode probabilites μ_u^s, initial visibility probability η_u and the velocity covariance matrix $\mathbf{P}_v^{(s)}$ as tuning parameters. From one point of view, μ^{0s} and $o^{0\mathrm{v}}$ should be set to the limit values of their respective Markov chains. From another viewpoint

[2]Opposed to the prediction of many Markov models used in target tracking, we should expect an unobserved target to have a limited stationary velocity covariance for all practical purposes, which should describe our uncertainty upon target arrival, and $\mathbf{P}_v^{(s)}$ should be designed accordingly.

the Markov chains are just approximations and some other values might be more appropriate. The former will lead to $\mu_u^s = \mu^{0s}$, whereas $o_u^{\mathrm{v}} \neq o^{0\mathrm{v}}$ due to the detection dependency. We note that not all values for μ_u^s and η_u could come from valid μ^{0s} and η^0, respectively. Nevertheless, as the Markov chains are in reality approximations, one might wish to ignore this in tuning for applications.

Remark 2.2. There can be potential performance gains in track initialization by appropriate modeling of a dynamic and/or spatially varying undetected intensity [139], not further investigated here.

The intensity is not a probability, so it is not a proper Markov chain. Its convergence properties are probably best analysed when including visibility, and therefore from the eigen-properties of the resulting matrix in (2.23). The convergence is on the order of a few 10s of time steps for many realistic parameter sets, such as the ones used in this chapter. The convergence transient can, however, be long, and as VIMMJIPDA neglects the transient it will in those cases have more suboptimal track initialization either initially or after convergence, depending on chosen parameters.

The equations above can provide insights into which birth intensity one is actually using, or be a tool in parameter selection. For tuning purposes, it should be noted from (2.19) that the unknown target intensity is always larger than the birth intensity when $P_{\mathrm{D}} < 1$ or $P_{\mathrm{S}} > 0$, so it is slightly misleading to tune the unknown target intensity in terms of how many objects that is expected to enter at a given time step. ◁

2.4.2 The VIMMJIPDA prior

In a PMB filter the mixture over association hypotheses is approximated by a single MB after every estimation cycle, so that the prior MBM component is of the form

$$G_{k-1}^{\mathrm{mbm}}[h] = \prod_{t=1}^{n_{k-1}} (1 - r_{k-1}^t + r_{k-1}^t f_{k-1}^t[h]). \tag{2.25}$$

This is the same as happens in the JIPDA, where the approximation is done by approximating the joint track to measurement association probabilities by the product of their marginals. Here, r_{k-1}^t is the existence probability and, recalling our hybrid state definition, the state densities in (2.25) are of the form

$$f_{k-1}^t(\mathbf{y}) = f_{k-1}^{ts}(\mathbf{x}) \mu_{k-1}^{ts} o_{k-1}^{\mathrm{tv}} \tag{2.26}$$

where $f_{k-1}^{ts}(\mathbf{x})$ is the prior kinematic density conditional on model s, μ_{k-1}^{ts} is the prior probability of model s and o_{k-1}^{tv} is the prior probability of the visibility state v. Reducing of the association hypotheses in this way, makes the form of (2.25) attained after both prediction and update.

We also introduce the notation $\eta^t = o^{t1}$ to denote the probability that target t is in visibility state 1, that is, the probability that it is visible. $f^{ts}_{k-1}(\mathbf{x})$ will typically be approximated as Gaussian through moment matching in the prediction and the update, as will be done in Section 2.5.

2.4.3 The VIMMJIPDA prediction

The PMBM prediction formulas from [138] state that the predicted MBM component remains an MB of the same form as (2.25), with existence probabilities and state densities given by

$$r^t_{k|k-1} = r^t_{k-1} f^t_{k-1}[P_{\mathrm{S}}(\tilde{\mathbf{y}})] \tag{2.27}$$

$$f^t_{k|k-1}(\mathbf{y}) = \frac{f^t_{k-1}[f_{\mathbf{y}}(\mathbf{y}|\tilde{\mathbf{y}})P_{\mathrm{S}}(\tilde{\mathbf{y}})]}{f^t_{k-1}[P_{\mathrm{S}}(\tilde{\mathbf{y}})]}. \tag{2.28}$$

Proposition 2.2. For any track t the predicted existence probability is

$$r^t_{k|k-1} = r^t_{k-1} P_{\mathrm{S}} \tag{2.29}$$

while the predicted model probabilities, visibility probabilities and kinematic pdfs are given by

$$\mu^{ts}_{k|k-1} = \textstyle\sum_{\tilde{s}} \pi^{\tilde{s}s} \mu^{t\tilde{s}}_{k-1} \tag{2.30}$$

$$\eta^t_{k|k-1} = \omega^{01}(1 - \eta^t_{k-1}) + \omega^{11}\eta^t_{k-1} \tag{2.31}$$

$$f^{ts}_{k|k-1}(\mathbf{x}) = \int f_{\mathbf{x}}(\mathbf{x}|\tilde{\mathbf{x}}) f^{ts,0}_{k-1}(\tilde{\mathbf{x}}) \mathrm{d}\tilde{\mathbf{x}}. \tag{2.32}$$

where the mode-conditional prior is given by

$$f^{ts,0}_{k-1}(\tilde{\mathbf{x}}) = \frac{1}{\sum_{\tilde{s}} \pi^{\tilde{s}s} \mu^{t\tilde{s}}_{k-1}} \sum_{\tilde{s}} \pi^{\tilde{s}s} \mu^{t\tilde{s}}_{k-1} f^{t\tilde{s}}_{k-1}(\tilde{\mathbf{x}}). \tag{2.33}$$

Proof. See Appendix 2.6. □

Remark 2.3. The mode-conditional prior $f^{ts,0}_{k-1}(\tilde{\mathbf{x}})$ is in the IMM framework approximated as a single Gaussian by means of moment-matching over the mixture in (2.33). The details are well known, and can be found in, for example, [8]. Assuming $f^{ts,0}_{k-1}(\tilde{\mathbf{x}})$ and $f_{\mathbf{x}}(\mathbf{x}|\tilde{\mathbf{x}})$ to be Gaussian-linear, it is straightforward to evaluate the Chapman-Kolmogorov integral in (2.32). Linearization is typically used when mild nonlinearities are present. ◁

To summarize, the predicted pdf of the hybrid state for any track t is given by

$$f^t_{k|k-1}(\mathbf{y}) = o^{\mathrm{tv}}_{k|k-1} \mu^{ts}_{k|k-1} f^{ts}_{k|k-1}(\mathbf{x}). \tag{2.34}$$

2.4.4 The VIMMJIPDA posterior

The PMBM update formulas expand the MB given by (2.27) and (2.28) into an MBM. In the standard PMBM formalism a track (also sometimes known as a potential target [42] or single-target hypothesis [138]) is established for every measurement ever observed, so that the number of tracks at time k is $n_k = m_k + n_{k-1} = \sum_{l=1}^{k} m_l$ where m_l is the number of measurements at time l. In practice, several of these tracks will be eliminated by pruning.

The MBM ranges over a collection $\mathcal{A}_k$ of the possible single-scan association hypotheses. In this chapter, the notation a_k^t denotes the measurement claimed by track t according to the hypothesis a_k at step k. Every hypothesis $a_k = (a_k^1, \dots, a_k^{n_k}) \in \mathcal{A}_k$ is a mapping from the set of tracks to the set of measurements in scan number k, including the zeroth measurement (missed detection), under the restrictions that non-zero measurements cannot be repeated and that all non-zero measurements must be covered. That is

$$\begin{aligned} \mathcal{A}_k = \{a_k \in \{0, \dots, m_k\}^{n_k} \text{ such that} \\ a_k^t = j > 0 \implies a_k^{t'} \neq j \text{ for all } t' \neq t, \\ \text{and if } a_k^t \neq j > 0 \text{ for all } t \leq n_{k-1} \text{ then } a_k^{n_{k-1}+j} = j\}. \end{aligned} \tag{2.35}$$

In the p.g.fl. formulation the MBM is of the form

$$G_k^{\text{mbm}}[h] \propto \sum_{a \in \mathcal{A}_k} \prod_{t=1}^{n_k} w_k^{ta_k^t} \left(1 - r_k^{ta_k^t} + r_k^{ta_k^t} f_k^{ta_k^t}[h]\right). \tag{2.36}$$

Here $w_k^{ta_k^t}$ is the weight contribution of track number t to the probability of the hypothesis a_k, $r_k^{ta_k^t}$ is the existence probability of track t under the hypothesis a, and $f_k^{ta_k^t}[h]$ is the linear functional corresponding to the state pdf of track t under the hypothesis a. For bookkeeping purposes we define the linear index function $L(j;k) = j + \sum_{l=1}^{k-1} m_l$ with inverse $L^{-1}(t;k) = t - \sum_{l=1}^{k-1} m_l$. What L does is to give us the index of the track with origin in the measurement $\mathbf{z}_k^j$, while its inverse

indicates that track t was created in an earlier time step if $L^{-1}(t;k) \leq 0$, or a later time step if $L^{-1}(t;k) > m_k$. Otherwise (that is, when $L^{-1}(t;k) \in \{1, \dots, m_k\}$), it gives the identity/index of the measurement that track t was created from [25].

The expressions for weights, existence probabilities and state pdf's vary between 4 cases: Empty track, new target, missed detection and detection. Below we recapitulate the PMBM expressions from [138] for these cases.

Track $t = L(j;k)$ is **empty** if measurement j is claimed by another track under hypothesis a_k, that is, when $\exists\, t' \neq t$ such that $a_k^{t'} = j$. In such a case $w_k^{ta_k^t} = 1$, $r_k^{ta_k^t} = 0$ and $f_k^{ta_k^t}(\mathbf{y})$ is left undefined.

We have a **new target** whenever $a_k^t > 0$ and $L^{-1}(t;k) > 0$, meaning that track t originated in the current scan. In such a case

$$w_k^{ta_k^t} = \lambda + u[f_{\mathbf{z}}(\mathbf{z}_k^{a_k^t} \mid \tilde{\mathbf{y}}) P_{\mathrm{D}}(\tilde{\mathbf{y}})] \tag{2.37}$$

$$r_k^{ta_k^t} = \frac{u[f_{\mathbf{z}}(\mathbf{z}_k^{a_k^t} \mid \tilde{\mathbf{y}}) P_{\mathrm{D}}(\tilde{\mathbf{y}})]}{\lambda + u[f_{\mathbf{z}}(\mathbf{z}_k^{a_k^t} \mid \tilde{\mathbf{y}}) P_{\mathrm{D}}(\tilde{\mathbf{y}})]} \tag{2.38}$$

$$f_k^{ta_k^t}(\mathbf{y}) = \frac{f_{\mathbf{z}}(\mathbf{z}_k^{a_k^t} \mid \mathbf{y}) P_{\mathrm{D}}(\mathbf{y}) u(\mathbf{y})}{u[f_{\mathbf{z}}(\mathbf{z}_k^{a_k^t} \mid \tilde{\mathbf{y}}) P_{\mathrm{D}}(\tilde{\mathbf{y}})]}. \tag{2.39}$$

We have a **missed detection** whenever $a_k^t = 0$ and $L^{-1}(t;k) \leq 0$, meaning that track t originated in an earlier scan. In such a case

$$w_k^{ta_k^t} = 1 - r_{k|k-1}^t + r_{k|k-1}^t f_{k|k-1}^t[1 - P_{\mathrm{D}}(\tilde{\mathbf{y}})] \tag{2.40}$$

$$r_k^{ta_k^t} = \frac{r_{k|k-1}^t f_{k|k-1}^t[1 - P_{\mathrm{D}}(\tilde{\mathbf{y}})]}{1 - r_{k|k-1}^t + r_{k|k-1}^t f_{k|k-1}^t[1 - P_{\mathrm{D}}(\tilde{\mathbf{y}})]} \tag{2.41}$$

$$f_k^{ta_k^t}(\mathbf{y}) = \frac{(1 - P_{\mathrm{D}}(\mathbf{y})) f_{k|k-1}^t(\mathbf{y})}{f_{k|k-1}^t[1 - P_{\mathrm{D}}(\tilde{\mathbf{y}})]}. \tag{2.42}$$

We have a **detection** whenever $a_k^t > 0$ and $L^{-1}(t;k) \leq 0$. In such a case

$$w_k^{ta_k^t} = r_{k|k-1}^t f_{k|k-1}^t[f_{\mathbf{z}}(\mathbf{z}_k^{a_k^t} \mid \tilde{\mathbf{y}}) P_{\mathrm{D}}(\tilde{\mathbf{y}})] \tag{2.43}$$

$$r_k^{ta_k^t} = 1 \tag{2.44}$$

$$f_k^{ta_k^t}(\mathbf{y}) = \frac{f_{\mathbf{z}}(\mathbf{z}_k^{a_k^t} \mid \mathbf{y}) P_{\mathrm{D}}(\mathbf{y}) f_{k|k-1}^t(\mathbf{y})}{f_{k|k-1}^t[f_{\mathbf{z}}(\mathbf{z}_k^{a_k^t} \mid \tilde{\mathbf{y}}) P_{\mathrm{D}}(\tilde{\mathbf{y}})]}. \tag{2.45}$$

In the subsequent propositions we translate these hybrid-state expressions into expressions for the existence probability, mode probability, visibility probability and kinematic density for each of the 3 non-empty cases under the assumptions M1*-M6*. Also after the measurement update it is possible to write the hybrid state pdfs in the factorized form

$$f_k^{tj}(\mathbf{y}) = o_k^{tvj} \mu_k^{tsj} f_k^{tsj}(\mathbf{x}) \tag{2.46}$$

for all state variables whose pdf or pmf is not identically zero.

Proposition 2.3. In the new target case we have

$$w_k^{tj} = \lambda + U\eta_u P_{\mathrm{D}} \tag{2.47}$$
$$r_k^{tj} = \frac{U\eta_u P_{\mathrm{D}}}{\lambda + U\eta_u P_{\mathrm{D}}} \tag{2.48}$$
$$\mu_k^{tsj} = \mu_u^s \tag{2.49}$$
$$\eta_k^{tj} = 1 \tag{2.50}$$
$$f_k^{tsj}(\mathbf{x}) \approx \mathcal{N}(\mathbf{x}\,;\, \hat{\mathbf{x}}_0^s, \mathbf{P}_0^s) \tag{2.51}$$

where $\hat{\mathbf{x}}_0^s = {\mathbf{H}^{(s)}}^\top \mathbf{z}_k^{a_k^t}$ and $\mathbf{P}_0^s = {\mathbf{H}^{(s)}}^\top \mathbf{R}^s \mathbf{H}^{(s)} + {\mathbf{H}^{*(s)}}^\top \mathbf{P}_v \mathbf{H}^{*(s)}$ are the new target initial state mean and covariance, respectively.[3]

Proof. See Appendix 2.6. □

Remark 2.4. Typically in a JIPDA implementation, one will only initialize tracks in unclaimed measurements, that is, measurements that have not been gated to existing tracks. This gives a formidable reduction in complexity compared to initializing new tracks in all measurements. Furthermore, one will typically use a linkage-based track clustering scheme, so that only tracks which share measurements are processed together. With these simplifications, every new track will constitute a separate cluster, containing only a single association hypothesis. The weight expression (2.47) will then never be needed, but the other expressions from Proposition 2.3 remain relevant. ◁

Proposition 2.4. In the missed detection case we have

$$w_k^{t0} = 1 - r_{k|k-1}^t + r_{k|k-1}^t(1 - \eta_{k|k-1}^t P_{\mathrm{D}}) \tag{2.52}$$
$$r_k^{t0} = \frac{r_{k|k-1}^t(1 - \eta_{k|k-1}^t P_{\mathrm{D}})}{1 - r_{k|k-1}^t + r_{k|k-1}^t(1 - \eta_{k|k-1}^t P_{\mathrm{D}})} \tag{2.53}$$
$$\mu_k^{ts0} = \mu_{k|k-1}^{ts0} \tag{2.54}$$
$$\eta_k^{t0} = \frac{(1 - P_{\mathrm{D}})\eta_{k|k-1}^t}{1 - P_{\mathrm{D}}\eta_{k|k-1}^t} \tag{2.55}$$
$$f_k^{ts0}(\mathbf{x}) = f_{k|k-1}^{ts}(\mathbf{x}) \tag{2.56}$$

Proof. See Appendix 2.6. □

[3] When the state vector consists of measured position coordinates first and non-measured velocity states last this becomes $\hat{\mathbf{x}}_0^s = [\mathbf{z}_k^{a_k^t}; \mathbf{0}]$ and $\mathbf{P}_0^s = \mathrm{diag}(\mathbf{R}^s, \mathbf{P}_v)$.

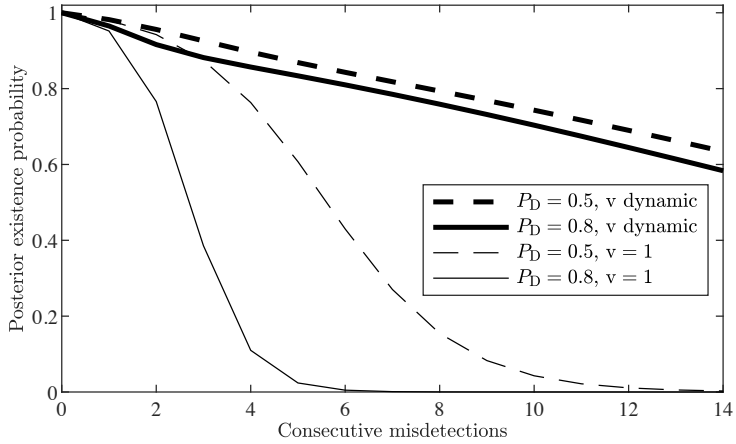

Figure 2.1: Impact of visibility model on existence probability for increasing number of consecutive missed detections.

Figure 2.1 compares the impact of the formulas (2.53) and (2.55) with enforcing unity visibility $\mathrm{v} = 1$. The transition probabilities from Table 2.1 have been used. For a target with a moderately high detection probability (for example, $P_{\mathrm{D}} = 0.8$), the existence probabilities differ by more than a decade after only 5 consecutive missed detections. Notice that the value of P_{D} has very little impact when the visibility model is used. In practice, the visibility probability can be seen as a dynamic detection probability that is estimated from the data. A non-unity value of P_{D} can then be seen as an upper bound on this dynamic detection probability.

Proposition 2.5. In the detection case we have

$$w_k^{tj} = P_{\mathrm{D}} r_{k|k-1}^t \eta_{k|k-1}^t \textstyle\sum_{\tilde{s}} \mu_{k|k-1}^{t\tilde{s}} l^{t\tilde{s}j} \tag{2.57}$$

$$r_k^{tj} = 1 \tag{2.58}$$

$$\mu_k^{tsj} = \mu_{k|k-1}^{ts} l^{tsj} / \textstyle\sum_{\tilde{s}} \mu_{k|k-1}^{t\tilde{s}} l^{t\tilde{s}j} \tag{2.59}$$

$$\eta_k^{tj} = 1 \tag{2.60}$$

$$f_k^{tsj}(\mathbf{x}) = f_{\mathbf{z}}^s(\mathbf{z}_k^j \mid \mathbf{x}) f_{k|k-1}^{ts}(\mathbf{x}) / l_k^{tsj} \tag{2.61}$$

where the track-to-measurement likelihood is given by

$$l_k^{tsj} = \int f_{\mathbf{z}}^s(\mathbf{z}_k^j \mid \tilde{\mathbf{x}}) f_{k|k-1}^{ts}(\tilde{\mathbf{x}}) \mathrm{d}\tilde{\mathbf{x}}. \tag{2.62}$$

Proof. See Appendix 2.6. □

2.4.5 Mixture reduction

The goal of the mixture reduction is to reduce the posterior MBM to an MB with hybrid state pdfs of the form $f_k^t(\mathbf{y}) = f_k^{ts}(\mathbf{x})\mu_k^{ts}o_k^{\mathrm{tv}}$. Following along the lines of a conventional JIPDA, and using $W = \sum_{a\in\mathcal{A}_k}\prod_{t=1}^{n_k} w_k^{ta^t}$, the MBM p.g.fl. after update is approximated as the MB

$$G_k^{\mathrm{mbm}}[h] = \sum_{a\in\mathcal{A}_k}\frac{1}{W}\prod_{t=1}^{n_k} w_k^{ta^t}\left(1 - r_k^{ta^t} + r_k^{ta^t} f_k^{ta^t}[h]\right) \tag{2.63}$$

$$\approx \prod_{t=1}^{n_k}\sum_{a\in\mathcal{A}_k}\frac{\prod_{t'=1}^{n_k} w_k^{t'a^{t'}}}{W}\left(1 - r_k^{ta^t} + r_k^{ta^t} f_k^{ta^t}[h]\right). \tag{2.64}$$

In other words, the joint is approximated by the product of its marginals. One can identify the hypothesis probabilities

$$\Pr(a_k) = \frac{1}{W}\prod_{t=1}^{n_k} w_k^{ta_k^t} \qquad a_k \in \mathcal{A}_k \tag{2.65}$$

used to calculate the marginal association probabilities

$$p_k^{tj} = \sum_{\substack{a_k\in\mathcal{A}_k \\ \text{s.t. } a_k^t=j}} \Pr(a_k), \qquad t \in \{1,\dots,n_k\} \tag{2.66}$$

which again are used in moment-based mixture reduction [138].

Proposition 2.6. Enforcing independence between the visibility state and the other states, the factors in the MB in (2.64) have marginal existence probability, visibility

probability, mode probabilities and mode conditional kinematic pdf's

$$r_k^t = \sum_{j=0}^{m_k} p_k^{tj} r_k^{tj} = 1 - p_k^{t0}(1 - r_k^{t0}), \tag{2.67}$$

$$\eta_k^t = \sum_{j=0}^{m_k} \underbrace{\frac{1}{r_k^t} p_k^{tj} r_k^{tj}}_{\beta_k^{tj}} \eta_k^{tj} = \frac{1 - p_k^{t0}(1 - \eta_k^{t0} r_k^{t0})}{r_k^t}, \tag{2.68}$$

$$\mu_k^{ts} = \sum_{j=0}^{m_k} \underbrace{\frac{1}{r_k^t} p_k^{tj} r_k^{tj}}_{\beta_k^{tj}} \mu_k^{tsj}, \quad \text{and} \tag{2.69}$$

$$f_k^{ts}(\mathbf{x}) = \sum_{j=0}^{m_k} \underbrace{\frac{1}{r_k^t \mu_k^{ts}} p_k^{tj} r_k^{tj} \mu_k^{tsj}}_{\beta_k^{tsj}} f_k^{tsj}(\mathbf{x}), \tag{2.70}$$

respectively. In the above, one can identify the existence conditional association probabilities

$$\beta_k^{tj} = \frac{p_k^{tj} r_k^{tj}}{r_k^t} = \begin{cases} \dfrac{p_k^{tj}}{r_k^t}, & j > 0 \text{ (that is, } r_k^{tj} = 1) \\ \dfrac{r_k^{t0} p_k^{t0}}{r_k^t}, & j = 0. \end{cases} \tag{2.71}$$

and the existence and mode conditional association probabilities

$$\beta_k^{tsj} = \frac{p_k^{tj} r_k^{tj} \mu_k^{tsj}}{r_k^t \mu_k^{ts}} = \beta_k^{tj} \frac{\mu_k^{tsj}}{\mu_k^{ts}} \tag{2.72}$$

Proof. See Appendix 2.6. □

Assuming that each $f_k^{tsj}(\mathbf{x})$ is Gaussian, we can approximate each $f_k^{ts}(\mathbf{x})$ by a Gaussian using standard moment-matching techniques with weights β_k^{tsj} [8, §1.4.16].

Remark 2.5. None of the discrete variables are part of the data association problem in (2.65) and (2.66), which is a natural consequence of treating them as part of the state space. They still clearly influence the data association, but only through the weights. To use a hybrid state space within the general PMBM framework to get to these equations contrasts earlier approaches with the unnecessary computational complexity of [28] and [14], and the rather lengthy derivations of [108]. When not taking visibility into account and disregarding the fact that JIPDA neglects unknown targets in the data association [138], the data association and mixture reduction

equations here are equivalent to [108, §V.B] We point out that augmenting with other target state variables can be achieved in a similar manner, and without the need to complicate the data association more than changing the weights.

Other parameters, such as clutter, unknown targets, detection probability etc. also influence the association through the weights in (2.65), just as in JPDA [10] and JIPDA [108]. ◁

Remark 2.6. The visibility v and the mode s are not *a posteriori* independent after the mixture reduction has been performed. This is so because the mixture reduction is equivalent to marginalization over a third variable, namely the association hypotheses. Nevertheless, for convenience we approximate them as independent, given by their respective marginals (2.68) and (2.69).

Maintenance of the full joint distribution would lead to greater complexity, needing M additional mode probabilities and kinematic state distributions if both are to be conditioned on the visibility.

It is possible to investigate the quality of this approximation using the forward Kullback Leibler divergence between their true joint posterior and its M-projection given by the marginals [103]. In the majority of test cases we have found that it is quite small (≤ 0.02), although it can also reach large values (≥ 0.2) in other cases, especially when the mode probabilities are very different for different association hypotheses. ◁

2.4.6 Implementation

The tracking results reported in this chapter are obtained with an implementation of the VIMMJIPDA that builds upon a real-time PDAF/IPDA implementation described in [143] [26]. In particular, preprocessing of the radar data follows the pipeline described in [143].

To mitigate complexity, validation gating with gate threshold $g = 3.5$ standard deviations is used. The gating is performed for each kinematic model, and the gated measurements become the union of the measurements gated by the individual kinematic model validation gates. As such, the probability of gating a true measurement under the given model is not constant, but we nevertheless assume that this probability is close to unity and therefore negligible. See, for example, [136] for other approaches to gating with multiple models. Only tracks that share measurements in their validation gates are processed together. Thus, single-linkage track clustering is used. For each track cluster, brute force hypothesis enumeration is performed if there is less than 4 targets or less than 2 measurements in the cluster. Otherwise, Murty's method with maximally 8 hypotheses is used. To solve the 2D assignment problems required by the Murty framework the auction method is used.

Tracks are divided into preliminary and confirmed tracks. Preliminary tracks are initialized using one-point initialization in all measurements that have not been validated by confirmed or preliminary tracks. Both preliminary and confirmed tracks are allowed to compete for the same measurements, although the implementation also includes an option only allowing preliminary tracks to access measurements not validated by confirmed tracks. Termination is done in the same manner as described in [143]. The termination procedure consists of three steps; firstly, all tracks with existence probability below an existence termination threshold T_d are terminated. Furthermore, tracks which are deemed to be identical through the hypotheses test described in [143] are found, and the youngest of the the identical tracks are terminated. Lastly, tracks which have not associated to any measurement for more than six time steps are terminated.

Kinematic models

We use 3 kinematic models in the IMM framework: A low-noise constant velocity (CV) model, a coordinated turn (CT) model, and a high-noise CV model. The CV models have state vector $\mathbf{x} = [x, y, v_x, v_y]^\top$ (that isposition and velocity) and are of the form $\mathbf{x}_k = \mathbf{F}^{(s)}\mathbf{x}_{k-1} + \mathbf{v}_k$, $\mathbf{v}_k \sim \mathcal{N}(\mathbf{0}, \mathbf{Q}^{(s)})$ where

$$\mathbf{F}^{(s)} = \begin{bmatrix} \mathbf{I}_2 & \mathrm{dt}\mathbf{I}_2 \\ \mathbf{0} & \mathbf{I}_2 \end{bmatrix}, \quad \mathbf{Q}^{(s)} = \begin{bmatrix} \mathrm{dt}^3/3\mathbf{I}_2 & \mathrm{dt}^2/2\mathbf{I}_2 \\ \mathrm{dt}^2/2\mathbf{I}_2 & \mathrm{dt}\mathbf{I}_2 \end{bmatrix} \sigma_{a,s}^2$$

where $s \in \{1, 2\}$. The CT model has state vector $\mathbf{x} = [x, y, v_x, v_y, \omega]^\top$ where ω is the turn rate, and is of the form $\mathbf{x}_k = \mathbf{F}(\mathbf{x}_{k-1})\mathbf{x}_{k-1} + \mathbf{v}_k$, $\mathbf{v}_k \sim \mathcal{N}(\mathbf{0}, \mathbf{Q})$ where

$$\mathbf{F}^{(3)}(\mathbf{x}) = \begin{bmatrix} 1 & 0 & \frac{\sin \mathrm{dt}\omega}{\omega} & \frac{-1+\cos \mathrm{dt}\omega}{\omega} & 0 \\ 0 & 1 & \frac{1-\cos \mathrm{dt}\omega}{\omega} & \frac{\sin \mathrm{dt}\omega}{\omega} & 0 \\ 0 & 0 & \cos \mathrm{dt}\omega & -\sin \mathrm{dt}\omega & 0 \\ 0 & 0 & \sin \mathrm{dt}\omega & \cos \mathrm{dt}\omega & 0 \\ 0 & 0 & 0 & 0 & 1 \end{bmatrix}$$

and

$$\mathbf{Q}^{(3)} = \begin{bmatrix} \mathbf{Q}_{a,1} & \mathbf{0} \\ \mathbf{0} & \sigma_\omega^2 \end{bmatrix}. \tag{2.73}$$

The CT model is linearized in the standard manner [8, §11.7.2].

Remark 2.7. The turn rate of the CT model might at first glance seem to stay forever independent of the other states. However, the linearization provides the needed correlation between the turn rate and the other states through the Jacobian $\mathbf{F}^{(3)} + [\mathbf{0}_{5\times 4}, (\partial_\omega \mathbf{F}^{(3)})\mathbf{x}]$ used in the covariance prediction, allowing the filter to estimate the turn rate. ◁

Measurement models

Measurements contain 2-dimensional positions converted to Cartesian coordinates:

$$\mathbf{z}_k = \mathbf{H}\mathbf{x}_k + \mathbf{w}_k, \quad \mathbf{w}_k \sim \mathcal{N}(\mathbf{0}, \mathbf{R}) \tag{2.74}$$

where $\mathbf{H} = [\mathbf{I}_2, \mathbf{0}]$ and the size of the zero matrix in $\mathbf{H}$ depends on the model. the measurement noise matrix consists of a Cartesian component $\mathbf{R}_c = \sigma_c^2 \mathbf{I}_2$ and a polar component $\mathbf{R}_p = \mathrm{diag}([\sigma_r^2, \sigma_\theta^2])$ so that

$$\mathbf{R} = \mathbf{R}_c + \mathbf{J}\mathbf{R}_p\mathbf{J}^\top \tag{2.75}$$

where $\mathbf{J}$ is the Jacobian of the polar-to-Cartesian mapping [10, §1.7.2].

2.5 Results on maritime radar data

We have implemented the VIMMJIPDA on two maritime radar data sets. In both these experiments, a Simrad Navico 4G broadband radar was mounted on top of the semi-autonomous surface craft Telemetron, which is an 8.45 m long rigid inflatable boat (RIB). The files can be found at CodeOcean through the URL `https://doi.org/10.24433/CO.3351829.v1`.

2.5.1 The Joyride data

In this data set, consisting of 238 radar scans, Telemetron was following a small motorboat that performed several volatile maneuvers. The motorboat was equipped with a smartphone GPS receiver, making ground truth available. The tracking results when using the full VIMMJIPDA tracker can be seen in Figure 2.2a. When the visibility state is removed, with no additional changes, the track following the motorboat becomes fragmented at the end of the scenario, as can be seen in Figure 2.2b. This is due to six consecutive missed detections. Two challenges this data set poses, namely course estimation and avoiding track loss, are discussed below.

To estimate the motorboat's course angle is difficult due to its high maneuverability. The course estimates for the motorboat can be seen in the top figure in Figure 2.5. Overall, the results are acceptable with regard to covariance consistency [8 §5.4]: The course RMSE was $16.6°$, while the corresponding standard deviation obtained by linearization from the posterior covariance matrices was $29.1°$. Further studies reveal that the course error exceeded $15°$, $30°$ and $45°$ in respectively 37%, 16% and 6% of the time steps. In a of total 3 scans, the actual error was above $45°$ while the standard deviation was estimated to be less than half of the error. These can be identified with three different situations, which are shown in Figure 2.2c. The first two (time steps 45 and 53) found place in sharp maneuvers. The last event (time

Table 2.1: Tracking system parameters

Quantity	Symbol	Unit	Value
Radar sample interval	T	[s]	2.5
Model 1 process noise	$\sigma_{a,1}$	[m/s²]	0.1
Model 2 process noise	$\sigma_{a,2}$	[m/s²]	1.5
Turn rate process noise	σ_{ω}	[rad/s²]	0.02
Cartesian noise std.	σ_c	[m]	6.6
Polar range std.	σ_r	[m]	8.0
Polar bearing std.	σ_{θ}	[°]	1.0
Detection probability	P_{D}	[%]	92
Survival probability	P_{S}	[%]	99.9
Initial visibility probability	η_u	[%]	90
Visibility Markov probability	$\omega^{\tilde{\mathrm{v}}\mathrm{v}}$	[%]	$\left[\begin{smallmatrix} 90 & 10 \\ 52 & 48 \end{smallmatrix}\right]$
Gate size	g	[−]	3.5
Clutter intensity	λ	[1/m²]	5×10^{-7}
Unknown target rate	U	[1/m²]	10^{-7}
Initial velocity std.	σ_v	[m/s]	10
Initial model probability	μ_u^s	[%]	$\left[\begin{smallmatrix} 80 & 10 & 10 \end{smallmatrix}\right]$
Existence confirmation threshold	T_c	[%]	99.9
Existence termination threshold	T_d	[%]	1
IMM transition probability	$\pi^{\tilde{s}s}$	[%]	$\left[\begin{smallmatrix} 99 & .5 & .5 \\ .5 & 99 & .5 \\ .5 & .5 & 99 \end{smallmatrix}\right]$

step 220) happened during re-acquisition after the target had been invisible for 6 scans, while executing two significant maneuvers. The speed estimates, which can be seen in the bottom figure of Figure 2.5, also have an acceptable consistency. The speed RMSE error was $1.77\,\mathrm{m/s}$, while the standard deviation was $5.34\,\mathrm{m/s}$.

The most obvious role of the visibility model is that of keeping the existence probability from becoming unnaturally low. This helps with the challenge of avoiding track loss. Without the visibility model, the existence probability drops to 2.5×10^{-5} during the 6 scans without detections, while it only drops to 0.75 with the visibility model, which arguably is more in line with common sense. This is due to the effects seen in Figure 2.1, and the existence and visibility probabilities for the situation with the 6 scans can be seen in Figure 2.4.

It is perhaps natural to ask what practical ramifications a low existence probability can have. Can one not simply set a very low termination threshold to remove the fragmentation problem? A low termination threshold would be a solution to the fragmentation problem, but data association in a JIPDA where the targets have very low existence probability can behave in an undesirable manner. As the different β_k^{tsj}

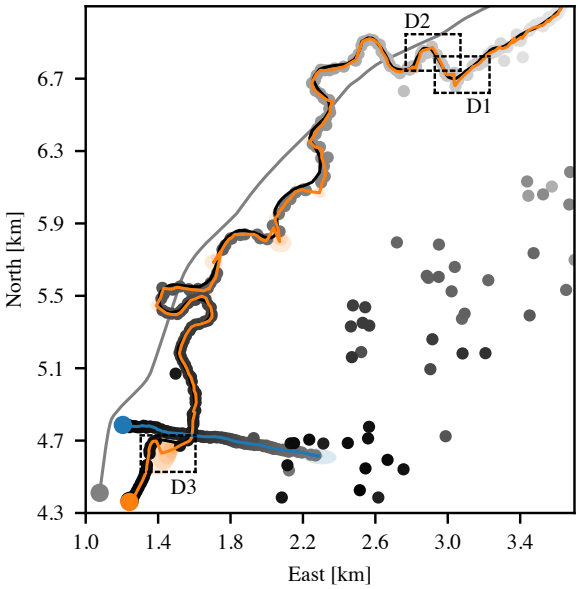

(a) With visibility state.

Figure 2.2: Figure continues on next page.

are conditioned on the existence of track t, they can, if the existence probability is small, be large for $j > 0$ even though the p^{tj} for $j > 0$ are very low. This happens when we have $p^{t0}r^{t0} \ll p^{tj} \ll p^{t0}$, where the variables in the former inequality is proportional to β^{t0} and β^{tj} from (2.71), in which we see the clear dependence on the existence probability.

The situation depicted in Figure 2.3 has been created to illustrate this. Here, the measurements from time steps 212, 213 and 214 are removed, and the termination threshold is reduced to 10^{-9}. This leads to both tracks being invisible for 6 consecutive time steps without being terminated. When the visibility model is not in use this gives both tracks a very low existence probability. Furthermore, when the next measurement for the blue track arrives at time step 217, the blue track is correctly assigned the measurement with high probability, and the orange track is assigned the missed detection with high probability. The blue track is correctly updated with the measurement to get an increased existence probability and corrected state estimate, and the orange track correctly gets a reduced posterior existence probability. However, under the missed detection hypothesis the orange track is most likely to be non-existent, and its existence conditioned state estimate is therefore updated with the measurement with a high weight. Thus, tracks with low existence

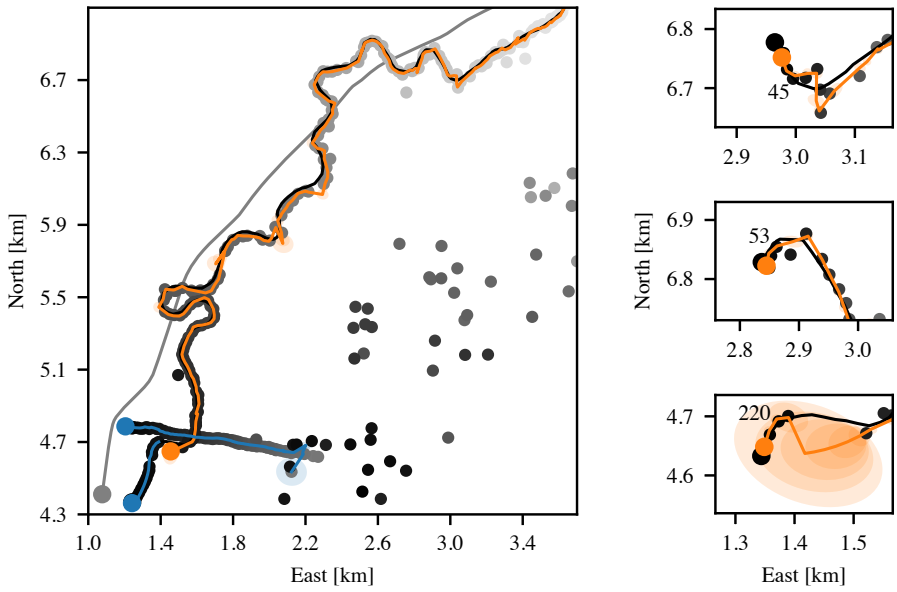

(b) Without visibility state.

(c) Data sections.

Figure 2.2: Track trajectories for the Joyride data with (left) and without (right) using the visibility model. Telemetron's trajectory is in grey, and the tracker results are in orange (motorboat) and blue (fishing vessel). The measurements are the grayscale dots with darker signifying later measurements. The dotted rectangles are data sections, numbered sequentially. Numbers within the data section plots are the corresponding time step, placed next to the measurement arriving at that specific time.

probability can start to follow other targets in the case of missed detection and without modelling visibility.

Drawing from this, at the time step when a new measurement appears to a cluster where the tracks have very low probability, the individual tracks in the cluster will largely ignore the presence of other tracks with regards to the updating of their kinematic pdfs. As such, the tracker output becomes similar to what one would expect from a single-target tracker. The visibility model mitigates this problem by ensuring a more realistic modeling of the existence probability.

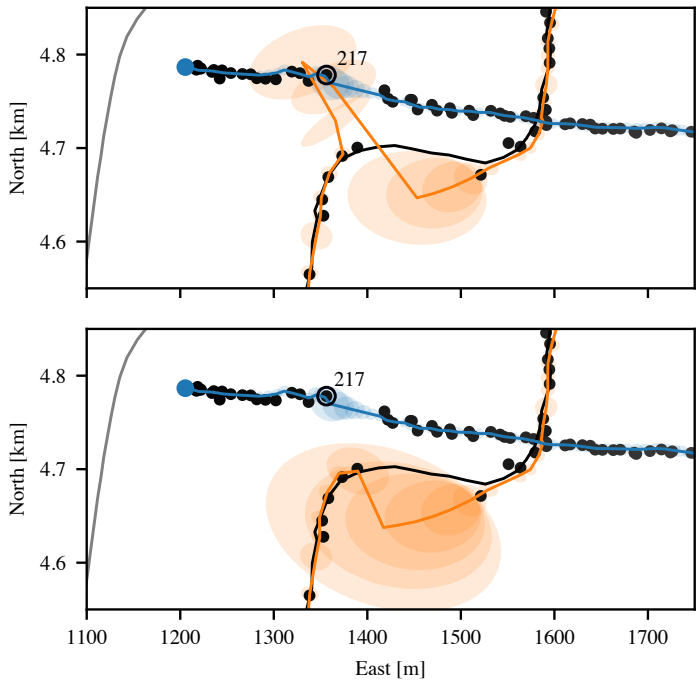

Figure 2.3: Above: Track jump as a result of low existence probabilities for both tracks. Below: The same behavior is avoided when using the visibility model.

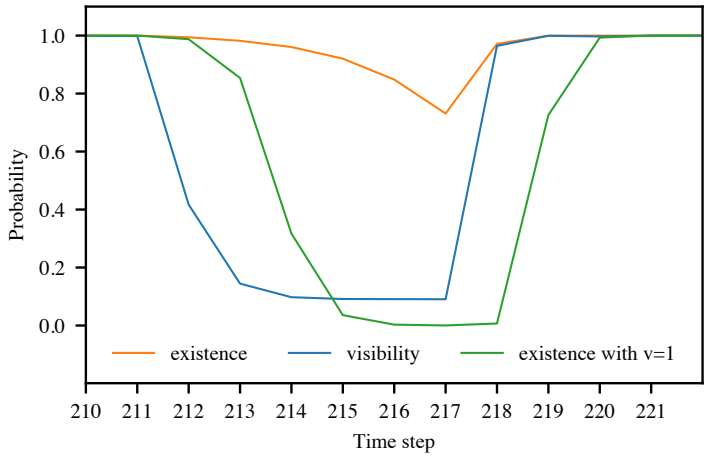

Figure 2.4: Existence and visibility probabilities for the motorboat during the six consecutive time steps without detection. The existence is shown both when using the visibility model, and when assuming the target is always in a visible state.

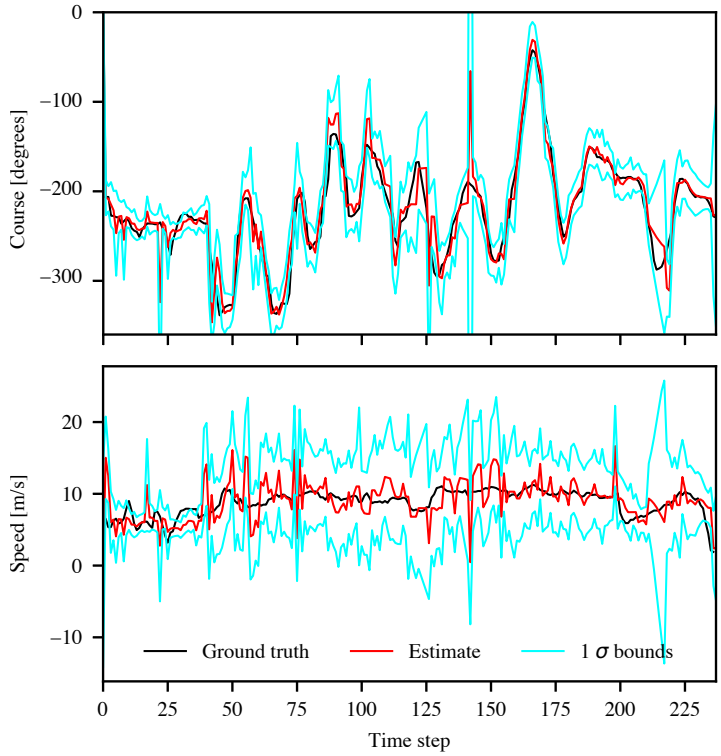

Figure 2.5: Course and speed estimates for the motorboat in the Joyride scenario.

2.5.2 The Autosea demonstration data

This data set, consisting of 1471 radar scans, stems from a COLAV experiment conducted in Trondheimsfjorden as a final demonstration of the Autosea project [26]. The experiment involved Telemetron together with two mid-size ships, the research vessel Gunnerus and the tugboat Munkholmen 2, which it had to avoid in autonomous mode. However, as the experiment was conducted in open waters, occasional traffic entered the experimental area. This included 3 high-speed RIBs with very different maneuvering characteristics than the intended target ships. Drone images of this scenario are shown in Figure 2.6. The tracking results for the scenario with the three RIBs, when using the full VIMMJIPDA tracker, can be seen in Figure 2.8a. The scenario consists of the data from the 314th to the 418th radar scans. Below we discuss track initialization, and kinematic modeling of targets with different maneuvering characteristics, two challenges that this data set poses for a multi-target tracking system.

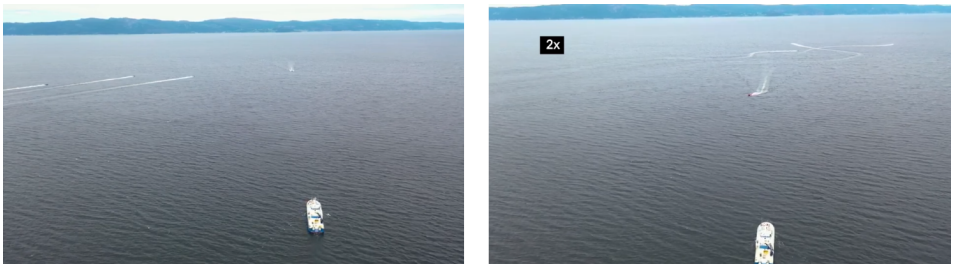

Figure 2.6: Drone video of the Autosea demonstration.Three high speed RIBs can be seen entering to the left, and playing in Telemetron's wake to the right. At the bottom is Gunnerus which is on collision course with Telemetron. Telemetron is supposed to discover and avoid Gunnerus autonomously.

Track initialization is a challenge because of the speed, maneuverability and closeness of the RIBs. In Figure 2.7a we see the results from the VIMMJIPDA with the default gate size $g = 3.5$. The first measurement of the green track arrives at time step 333, while the first measurement of the orange track arrives at time step 338. The validation gates of the tracks at these time steps are shown, and the first measurement originating from the RIBs are highlighted whenever they lead to the initialization of a track. The course and speed estimates for for the westernmost RIB, that isthe orange track in the left section of Figure 2.8a, can be seen in Figure 2.9. The high speed ($\approx 18\,\mathrm{m/s} \approx 65\,\mathrm{km/h}$) makes it necessary to initialize the tracks with high initial velocity standard deviations. This increases the susceptibility to false tracks. The maneuverability means that a track initialized just before a maneuver may turn into a lost track before it has the time to establish much confidence. This tendency can be seen in the blue track in Figure 2.7b. Even though the blue track is not ultimately lost, it misses its first measurement after initialization due to the too small validation gate. The closeness tends to cause the two eastern RIBs to share measurements between their validation gates, with the consequence that it may take extra time to establish a track on the middle RIB if the validation gates are too large, as the easternmost RIB tends to get its track established first. This can be seen in Figure 2.7c. Here, the validation gates of the easternmost track blocks the middle track from initializing until seven time steps after the first measurement from its target arrived. Careful tuning is needed to handle this, as smaller validation gates can lead to a track missing measurements due to them falling outside its validation gate, while too large validation gates prevent proper initialization.

Closeness also leads to challenges with track-loss, swaps and coalescence, depending on the tuning parameters. After having safely passed Telemetron, the RIBs execute severe maneuvers similar to those seen in the Joyride data (for example, accelerations of more than $3\,\mathrm{m/s^2}$ over a time of 10 s). The most exciting phase of

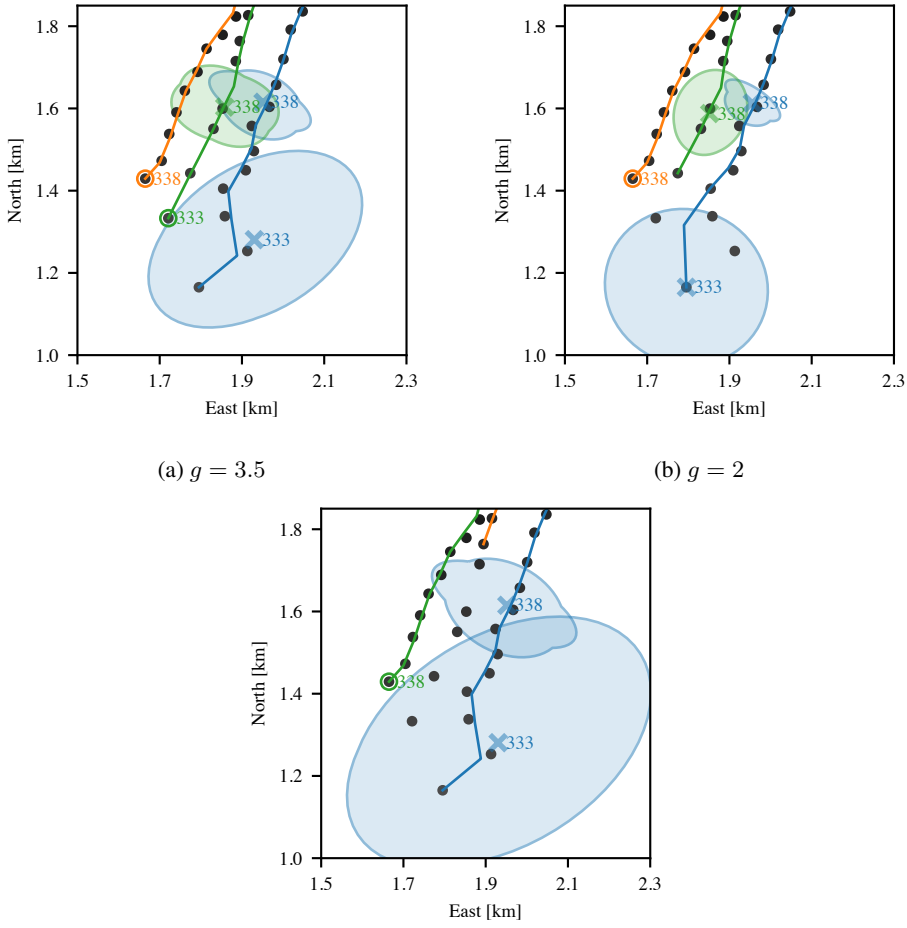

(a) $g = 3.5$

(b) $g = 2$

(c) $g = 5$

Figure 2.7: Initialization of tracks on the RIBs with different validation gate sizes. The validation gates for existing tracks are shown at the times of the appearance of the first measurements from the other targets. The first measurements are highlighted if an initialization has taken place. The time step of validation gates and highlighted measurements are shown.

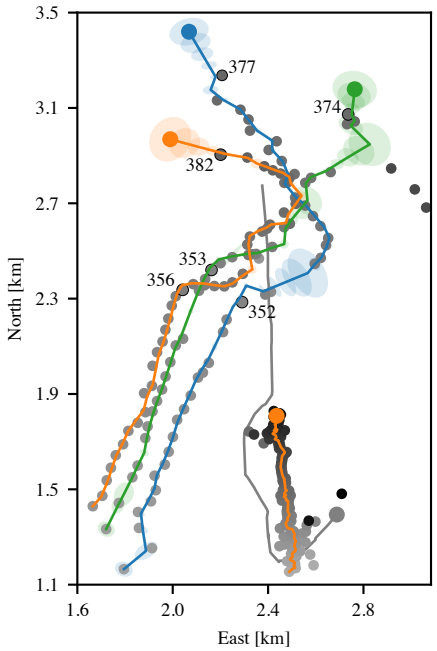

(a) IMM framework with two CV models and a CT model.

Figure 2.8: Figure continues on next page.

the scenario happens during the time steps 352-377. The measurements originating from the RIBs during these time steps can be seen in the upper halves of the plots in Figure 2.8. The estimated changes in both course and speed for the leftmost RIB can be seen in Figure 2.9. The closeness and high maneuverability of the RIBs and the low maneuverability of Gunnerus highlights the usefulness of the IMM framework. In Figure 2.8 the tracking results are shown when using the IMM framework, when using a single CV model with $\sigma = 1.5$, and when using a single CV model with $\sigma = 0.5$. Furthermore, the IMM model probabilities for Gunnerus and the leftmost RIB can be seen in Figure 2.10. When only using a single CV model with $\sigma = 1.5$, as is shown in Figure 2.8b, the RIBs are tracked in a similar manner as when using the IMM framework. The RIBs have a high maneuverability, and are well suited for a kinematic model with high process noise. The track for Gunnerus, however, becomes noisy compared to when using the IMM framework. The CV model with high process noise does not fit the low maneuverability of the larger ship, and the IMM framework assigns more than 90% of its probability mass to the low-noise model. If one, when still using a single model, tries to amend this by reducing the

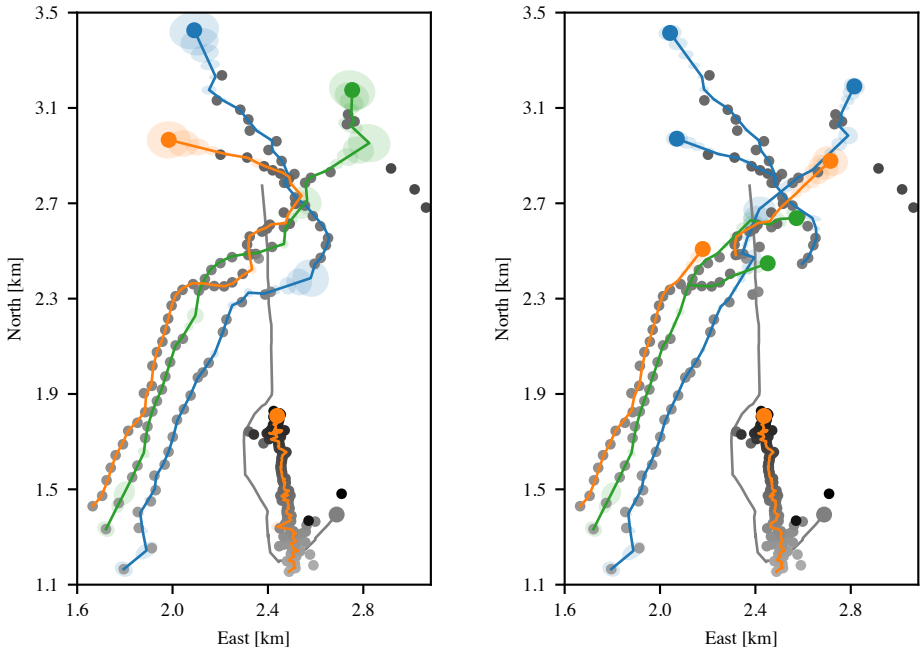

(b) Single CV model with $\sigma = 1.5$.

(c) Single CV model with $\sigma = 0.5$.

Figure 2.8: Track trajectories for scenario with the three RIBs (radar scans 314-418), when using the IMM framework, and when using only a single kinematic model.

process noise, a result as the one shown in Figure 2.8c is typical. Here, the track for Gunnerus is less noisy, albeit more noisy than when using the IMM framework, but the tracker is not able to properly follow the RIBs. Thus, in the presence of targets with highly different maneuverability the kinematic modeling has to be more flexible than what a single kinematic model with constant process noise can provide.

2.6 Chapter summary

We have presented a JIPDA-based algorithm for detection and tracking of multiple maneuvering targets, and studied its performance on two real-world data sets recorded by a low-cost maritime radar. We deem its performance satisfactory due to reasonable consistency properties and avoidance of track-loss. The data sets exhibit several challenges, and can be used as a benchmark test also for other tracking methods.

The most important challenge posed by these data is perhaps the difficulty of obtaining accurate course estimates. The large course uncertainties imply that safety

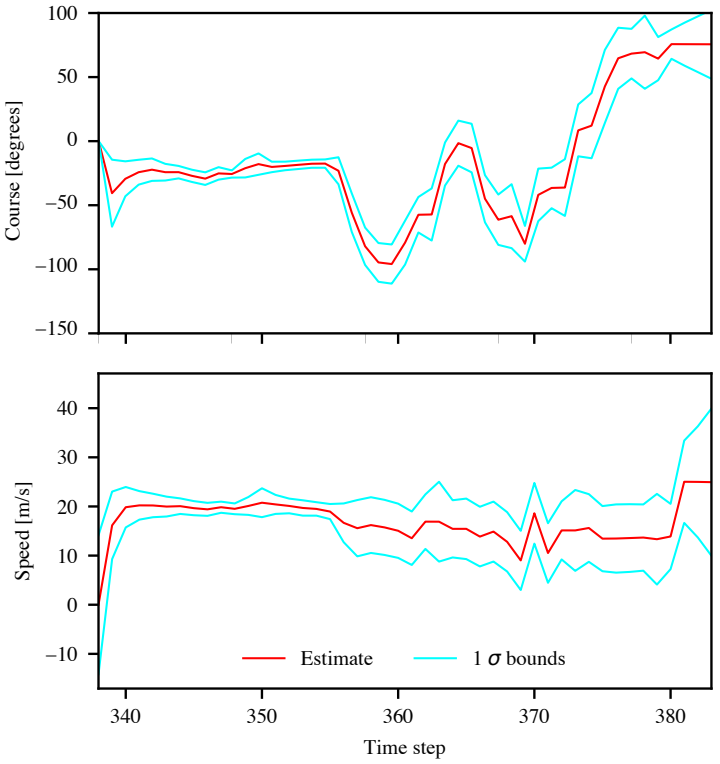

Figure 2.9: Course and speed estimates for the leftmost RIB in the Autosea demonstration scenario. The RIB changes its course by $\approx 90^\circ$ within $12\mathrm{s}$, while maintaining a speed of $\approx 15\mathrm{m/s}$.

margins must be large when a maritime COLAV relies on radar tracking alone. The uncertainties can probably be somewhat mitigated by models for wake clutter, such as those proposed in [119] and [19]. One may also hope that more advanced multiple model schemes than the standard IMM framework could improve performance. However, the fundamental limitations of radar resolution and sampling rate can ultimately only be beaten by additional measurements, such as Doppler and camera measurements. Another direction of future research is fusion with AIS data, possible along the lines of [42], which can improve the track quality for larger vessels with AIS transponders [63].

Further future investigations also include more in-depth studies of different independence approximations briefly remarked here, keeping Gaussian mixtures for each track using [2, 30, 120], both homogeneous and heterogeneous multi sensor scenarios [61], analysis of performance bounds from a parameter perspective, and

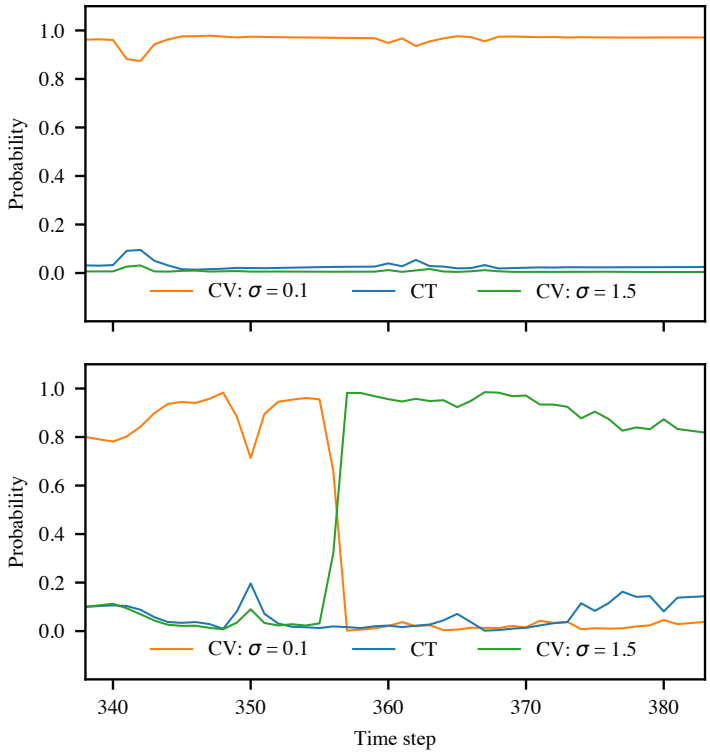

Figure 2.10: IMM model probabilities for Gunnerus (above) and the leftmost RIB (below). We see how the starboard turn of the RIB after $k = 356$ triggers a redistribution of probability mass onto the high-noise model, while the track on the larger Gunnerus vessel always prefers the low noise model. The CT model is struggling to accrue probability mass, possibly because the maneuvers are too short compared to the sampling time.

comparisons to what could be termed VIMM-PMBM/GLMB/MHT, with the main algorithm contributions already laid out in this chapter.

Acknowledgment

The radar data were recorded by Erik Wilthil, Andreas Flåten, Bjørn-Olav Eriksen and Giorgio D. K. M. Kufoalor, with assistance from Maritime Robotics and Kongsberg. The term "visibility" was suggested by Roslyn Lau.

Appendix I

In this appendix we prove Proposition 2.1, which describes the steady-state unknown target intensity. The convergence to a steady-state intensity is due to two simple facts. 1) Birth is acting additive and death and observation are acting multiplicative; the multiplicative decrease grows with the intensity, whereas the additive increase does not, so the intensity must be bounded above when all states have a positive probability of reaching a region of decrease in finite time. 2) Any oscillatory behavior present in the Markov model must die out due to the stationary "fill in" from birth and "drain" due to death and detection.

Except for the overall rate and visibility, the states are independent of each other in birth, prediction and detection and will therefore be independent upon convergence. The first equality in the equations for U, μ_u^{s} and o_u^{v} follow directly by marginalization of the other variables after inserting $u_{k|k-1}(\mathbf{y}) = u_{k-1|k-2}(\mathbf{y}) = u(\mathbf{y})$ in (2.16), and the second by solving algebraically.

In order to find $f_u(\mathbf{x})$, we note that all states can reach a point where $P_{\mathrm{D}}(\mathbf{y}) < 1$ or $P_{\mathrm{S}}(\mathbf{y}) < 1$ in *finite* time, so any initial pdf will phase out and leave the impact of the birth pdf under the Markov model. Hence, with the approximation (2.17), $f_u(\mathbf{x}) \approx f_b(\mathbf{x})$.

Appendix II

In this appendix we prove Proposition 2.2, which describes the track-wise prediction in the VIMMJIPDA. The result for the existence prediction follows from simply noting that $f_{k-1}^t[P_{\mathrm{S}}(\tilde{\mathbf{y}})] = P_{\mathrm{S}} f_{k-1}^t[1] = P_{\mathrm{S}}$ since P_{S} is constant and the linear functional of a pdf evaluated at the unit function is unity. For the hybrid state, we can decompose the prediction into a visibility part and a mode-kinematics part:

$$f_{k|k-1}^t(\mathbf{y}) = \left(\sum_{\tilde{\mathrm{v}}} \omega^{\tilde{\mathrm{v}}\mathrm{v}} \eta_{k-1}^{t\tilde{\mathrm{v}}} \right) p_{k|k-1}^t(\mathbf{x}, s). \tag{2.76}$$

The sum inside the parentheses is equal to (2.31). The joint density of kinematic state and mode is

$$\begin{aligned} p^t_{k|k-1}(\mathbf{x}, s) &= \sum_{\tilde{s}} \pi^{\tilde{s}s} \mu^{t\tilde{s}}_{k-1} \int f_{\mathbf{x}}(\mathbf{x}|\tilde{\mathbf{x}}) f^{t\tilde{s}}_{k-1}(\tilde{\mathbf{x}}) \mathrm{d}\tilde{\mathbf{x}} \\ &= \overbrace{\left(\textstyle\sum_{\tilde{s}} \pi^{\tilde{s}s} \mu^{t\tilde{s}}_{k-1}\right)}^{\mu^{ts}_{k|k-1}} \int f_{\mathbf{x}}(\mathbf{x}|\tilde{\mathbf{x}}) \underbrace{\frac{1}{\sum_{\tilde{s}} \pi^{\tilde{s}s} \mu^{t\tilde{s}}_{k-1}} \sum_{\tilde{s}} \pi^{\tilde{s}s} \mu^{t\tilde{s}}_{k-1} f^{t\tilde{s}}_{k-1}(\tilde{\mathbf{x}})}_{f^{ts,0}_{k-1}(\tilde{\mathbf{x}})} \mathrm{d}\tilde{\mathbf{x}}. \end{aligned} \tag{2.77}$$

The over- and underbraces in (2.77) illustrate how the predicted mode probability and mode-conditional prior both are readily extracted from this joint density.

Appendix III

In this appendix we prove Propositions 2.3, 2.4 and 2.5, which describe the measurement updates for the three cases of new target, missed detection and detection, respectively.

For the **new target** case in Proposition 2.3, the key term that governs the new target expressions is the product $f_{\mathbf{z}}(\mathbf{z}^{a^t_k}_k \mid \mathbf{y}) P_{\mathrm{D}}(\mathbf{y}) u(\mathbf{y})$ in (2.37)-(2.39). Under the assumptions M1*-M6* it is given by

$$P_{\mathrm{D}} U \mu^s_u \eta_u 1_\Omega(\mathbf{H}^{(s)}\mathbf{x}) \mathcal{N}(\mathbf{z}^{a^t_k}_k; \mathbf{H}^{(s)}\mathbf{x}, \mathbf{R}^{s)}) \mathcal{N}(\mathbf{H}^{*(s)}\mathbf{x}; \mathbf{0}, \mathbf{P}^{(s)}_v) \tag{2.78}$$

insofar as $\mathrm{v} = 1$. Otherwise it is 0. Because the vertical stacking of $\mathbf{H}^{(s)}$ and $\mathbf{H}^{*(s)}$ is a permutation matrix, we can rest assured that the product of Gaussians is proportional to the Gaussian $\mathcal{N}(\mathbf{x}\,; \hat{\mathbf{x}}^s_0, \mathbf{P}^s_0)$. If the region Ω is large enough, then the term $1_\Omega(\mathbf{H}^{(s)}\mathbf{x})$ has no significant impact on the product as a function of $\mathbf{x}$. Consequently, the integral over $\mathbf{x}$ that is latent in the weight and existence expressions is approximated as 1. This leads to

$$u[f_{\mathbf{z}}(\mathbf{z}^{a^t_k}_k \mid \tilde{\mathbf{y}}) P_{\mathrm{D}}(\tilde{\mathbf{y}})] \approx U \eta_u P_{\mathrm{D}} \tag{2.79}$$

and the expressions for w^{tj}_k and r^{tj}_k follow from this. This also establishes (2.51). We can furthermore see that the dependence on the model state s is encapsulated by μ^s_u. For the visibility, the term η_u is canceled by a similar term from the normalization integral in the denominator whenever $\mathrm{v} = 1$, leading to unity probability of $\mathrm{v} = 1$.

For the **missed detection** case in Proposition 2.4, we start by noticing that the inner product $f^t_{k|k-1}[1 - P_{\mathrm{D}}(\tilde{\mathbf{y}})]$ appears several places in (2.40)-(2.42). Under the assumptions M1*-M6* it becomes

$$\begin{aligned}
&f^t_{k|k-1}[1 - P_{\mathrm{D}}(\tilde{\mathbf{y}})] \\
&\quad = \sum_{\mathrm{v}} \sum_{s} \int o^{t\mathrm{v}}_{k|k-1} \mu^{ts}_{k|k-1} f^{ts}_{k|k-1}(\mathbf{x})(1 - P_{\mathrm{D}}(\mathrm{v}))\, \mathrm{d}\mathbf{x} \\
&\quad = \sum_{\mathrm{v}} o^{t\mathrm{v}}_{k|k-1}(1 - P_{\mathrm{D}}(\mathrm{v})) \\
&\quad = (1 - \eta^t_{k|k-1}) \cdot 1 + \eta^t_{k|k-1}(1 - P_{\mathrm{D}}) \\
&\quad = 1 - \eta^t_{k|k-1} P_{\mathrm{D}}.
\end{aligned} \tag{2.80}$$

Based on this, the expression for the hybrid state pdf in (2.42) becomes

$$f^{ts0}_k(\mathbf{y}) = \begin{cases} \dfrac{f^{ts}_{k|k-1}(\mathbf{x}) \mu^{ts}_{k|k-1}(1 - \eta^t_{k|k-1})}{1 - \eta^t_{k|k-1} P_{\mathrm{D}}} & \text{if } \mathrm{v} = 0 \\ \dfrac{(1 - P_{\mathrm{D}}) f^{ts}_{k|k-1}(\mathbf{x}) \mu^{ts}_{k|k-1} \eta^t_{k|k-1}}{1 - \eta^t_{k|k-1} P_{\mathrm{D}}} & \text{if } \mathrm{v} = 1 \end{cases}$$

We can verify that the posterior mode probabilities are the same as the predicted mode probabilities by marginalization over the kinematic state and the visibility:

$$\mu^{ts0}_k = \mu^{ts}_{k[k-1} \frac{1 - \eta^t_{k|k-1} + \eta^t_{k|k-1}(1 - P_{\mathrm{D}})}{1 - \eta^t_{k|k-1} P_{\mathrm{D}}} = \mu^{ts}_{k[k-1}. \tag{2.81}$$

In the same manner we can show that the kinematic density remains the same. The visibility probability is also found by marginalization over the other parts of the hybrid state:

$$\eta^{t0}_k = \frac{(1 - P_{\mathrm{D}}) \eta^t_{k|k-1}}{(1 - \eta^t_{k|k-1}) + \eta^t_{k|k-1}(1 - P_{\mathrm{D}})} = \frac{(1 - P_{\mathrm{D}}) \eta^t_{k|k-1}}{1 - P_{\mathrm{D}} \eta^t_{k|k-1}}. \tag{2.82}$$

Notice that the posterior mode and visibility are independent in this case.

For the **detection** case in Proposition 2.5, the existence probability is trivial. The weight expression is a straightforward evaluation of the integral in (2.43). The remaining entities are found by marginalization of the hybrid state pdf from (2.45), which can be written

$$\frac{P_{\mathrm{D}} \eta^t \mu^{ts}_{k|k-1} f^s_{\mathbf{z}}(\mathbf{z}^j_k \mid \mathbf{x}) f^{ts}_{k|k-1}(\mathbf{x})}{\sum_s P_{\mathrm{D}} \eta^t_k \mu^{ts}_{k|k-1} l^{tsj}_k} = \underbrace{\frac{\mu^{ts}_k l^{tsj}_k}{\sum_s \mu^{ts}_k l^{tsj}_k}}_{\mu^{tsj}_k} \underbrace{\frac{f^s_{\mathbf{z}}(\mathbf{z}^j_k \mid \mathbf{x}) f^{ts}_{k|k-1}(\mathbf{x})}{l^{tsj}_k}}_{f^{tsj}_k(\mathbf{x})} \tag{2.83}$$

in the case that $\mathrm{v} = 1$, while it is zero when $\mathrm{v} = 0$. Consequently, the posterior visibility probability is unity, which corresponds well with our intuition.

Appendix IV

In this appendix we prove Proposition 2.6, which describes the mixture reduction in the VIMMJIPDA. Using (2.66), the sums in the Bernoullis in the MB (2.64) is reduced to a sum over $a_k^t \in \{0, \dots, m_k\}$ and becomes

$$\sum_{j=0}^{m_k} p_k^{tj} \left(1 - r_k^{tj} + r_k^{tj} f_k^{tj}[h]\right) = 1 - r_k^t + r_k^t \sum_{j=0}^{m_k} \beta_k^{tj} f_k^{tj}[h], \tag{2.84}$$

where $\sum_{j=0}^{m_k} p_k^{tj} = 1$, (2.67) and (2.71) have been used. We now recognise $f_k^t(\mathbf{y}) = \sum_{j=0}^{m_k} \beta_k^{tj} f_k^{tj}(\mathbf{y})$.

To achieve the wanted conditioning we expand f_k^{tj}, and multiply and divide with the normalization constants of the kinematic-mode-visibility joint and the mode-visibility joint, respectively, to get

$$\begin{aligned}
\sum_{j=0}^{m_k} \beta_k^{tj} f_k^{tj}(\mathbf{y}) &= \sum_{j=0}^{m_k} \beta_k^{tj} f_k^{tsj}(\mathbf{x}) \mu_k^{tsj} o_k^{\mathrm{tv}j} \\
&= \frac{\sum_{j=0}^{m_k} \beta_k^{tj} f_k^{tsj}(\mathbf{x}) \mu_k^{tsj} o_k^{\mathrm{tv}j}}{\sum_{j=0}^{m_k} \beta_k^{tj} \mu_k^{tsj} o_k^{\mathrm{tv}j}} \frac{\sum_{j=0}^{m_k} \beta_k^{tj} \mu_k^{tsj} o_k^{\mathrm{tv}j}}{\sum_{j=0}^{m_k} \beta_k^{tj} o_k^{\mathrm{tv}j}} \sum_{j=0}^{m_k} \beta_k^{tj} o_k^{\mathrm{tv}j} \\
&\approx \overbrace{\sum_{j=0}^{m_k} f_k^{tsj}(\mathbf{x}) \underbrace{\frac{\beta_k^{tj} \mu_k^{tsj}}{\sum_{j=0}^{m_k} \beta_k^{tj} \mu_k^{tsj}}}_{\beta_k^{tsj}}}^{f_k^{ts}} \underbrace{\sum_{j=0}^{m_k} \beta_k^{tj} \mu_k^{tsj}}_{\mu_k^{ts}} \underbrace{\sum_{j=0}^{m_k} \beta_k^{tj} o_k^{\mathrm{tv}j}}_{o_k^{\mathrm{tv}}}.
\end{aligned} \tag{2.85}$$

The approximation in the last line enforces independence of the visibility and the other states. The statement then follows with $\eta_k^t = o_k^{t1}$.

For the last equalities in (2.67) and (2.68), notice that $r_k^{tj} = \eta_k^{tj} = 1$ for all $j > 0$ and $\sum_{j=1}^{m_k} p_k^{tj} = 1 - p_k^{t0}$. So $\sum_{j=0}^{m_k} p_k^{tj} r_k^{tj} = \sum_{j=1}^{m_k} p_k^{tj} + p_k^{t0} r_k^{t0} = 1 - p_k^{t0} + p_k^{t0} r_k^{t0}$, and exactly the same for η_k^t with $\eta_k^{t0} r_k^{t0}$ inserted for r_k^{t0}.

Chapter 3

Variations of Joint Integrated Probabilistic Data Association With Radar and Target-Provided Measurements

This chapter is based on the publication:

[68] A. G. Hem and E. F. Brekke. Variations of joint integrated data association with radar and target-provided measurements. *Journal of Advances in Information Fusion*, 17(2):97–115, Dec. 2022

Changes from the original publication involve shortening of the introduction section, and a restructuring of the results section reflecting the insight from [43].

3.1 Introduction

There are two established approaches to the fusion of sensor signals: track-to-track fusion and track-to-measurement fusion [7]. Here track-to-measurement fusion is examined, and a model suitable for incorporating target-provided measurements, and a tracking algorithm utilizing this model, is presented. The tracking algorithm presented here differs from previous work in some significant ways. We use a hybrid state framework based on Chapter 2, which can include motion and visibility models in addition to target IDs. Furthermore, building upon Chapter 2, we derive the tracking algorithm as a special case of the Poisson Multi-Bernoulli Mixture (PMBM) filter originally proposed in [138]. An important technical detail to enable this is to model the birth model as a marked Poisson point process (PPP), where the target

IDs take the role of the marks. The resulting algorithm can be seen as a generalized version of Joint Integrated Probabilistic Data Association [105].

This leads to the contributions of this chapter. It derives a framework that includes target-provided measurements based on a PMBM formulation of the JIPDA. The resulting target tracker includes both a visibility state and multiple kinematic models. Furthermore, the chapter details a sequential way of handling the incoming target-provided measurements, a method more similar to the one described in [44], and a method similar to how radar measurements are processed. Lastly, we present some implementation-specific considerations to make when handling target-provided measurements in a tracker.

The chapter is organized as follows. We detail the problem formulation in Section 3.3. In Section 3.4, we explain the structure of the hybrid state that facilitates the inclusion of target-provided information. We present the mathematical expressions needed for calculations in Section 3.5. In Section 3.6, three different methods for handling the incoming measurements are detailed. Section 3.7 presents the implementation choices, together with considerations to make to accommodate the target-provided measurements. Lastly, Section 3.8 presents the results. We compare the performance of the different measurement handling methods and how they compare to using only radar and the method from [44].

3.2 Background

This work builds upon the multi-target tracking method presented in Chapter 2 and can be considered an extension of the framework described there. The tracking algorithm, denoted as Visibility interacting multiple models joint integrated probabilistic data association (VIMMJIPDA), combines interacting multiple models (IMM) and a visibility state with the well-established JIPDA framework. The tracking method was derived with a basis in the Poisson multi-Bernoulli mixture (PMBM) filter [138].

Darko Musicki and Rob Evans introduced the JIPDA in [105], where the concept of visibility is mentioned and indicates whether the tracked target is visible to the sensor or not. Later, for example, [141] has explored visibility in connection with the problem of estimating target detectability. The JIPDA is an extension of the Joint Probabilistic Data Association (JPDA) method developed by Yaakov Bar-Shalom [41], which again is an extension of Bar-Shalom's probabilistic data association (PDA) method [9]. These methods are well established in the target tracking community and have been used for a range of different purposes, such as collision avoidance for marine vessels [123], autonomous navigation [33], and air traffic control [93]. Henk A. P. Blom and Yaakov Bar-Shalom introduced the IMM method [15], and it has been used for several decades in,for example, air traffic control. Furthermore, Musicki and Suvorova presented an IMM-JIPDA algorithm

in [108].

The PMBM filter and subsequent tracking algorithms [57] utilize the PMBM density, which is the union between a Poisson point process (PPP) and a multi-Bernoulli mixture. The PPP represents unknown targets, that is, undetected targets hypothesized to exist, and the multi-Bernoulli mixture represents already detected targets. Links between PMBM and JIPDA have been established in [138] (single kinematic model, loopy belief propagation as an alternative to hypothesis enumeration) and in Chapter 2 (multiple kinematic models, standard hypothesis enumeration and mixture reduction).

Here, a multi-sensor network for maritime surveillance is described, utilizing several sensors, including radar and AIS. More recently, research has been conducted into the track-to-track association of radar- and AIS-tracks [121].

3.3 Problem formulation

The modeling assumptions used throughout this chapter are presented below.

M1

The unknown target intensity $u(\mathbf{y})$ describes the not yet discovered targets present in the surveillance area. We model the unknown targets as a Marked PPP, which is equivalent to a PPP on the Cartesian product of the space $\mathbf{R}^{n_x}$ and the discrete spaces the discrete hybrid states can take values from [125, p. 205]. In its general form, this process is

$$b(\mathbf{y}) = f(\mathrm{v})f(\tau|\mathrm{v})p(s|\mathrm{v},\tau)f_\gamma(\mathbf{x}|s,\mathrm{v},\tau). \tag{3.1}$$

where $f_\gamma(\mathbf{x}|s,\mathrm{v},\tau)$ is an intensity function on the the space $\mathbf{R}^{n_x}$, and $f(\cdot)$ are distributions over the discrete states. Rather than using the birth intensity directly, we use Proposition 1 from Chapter 2 to get the converged unknown target intensity

$$u(\mathbf{y}) = U o_u^{\mathrm{v}} \xi_u^\tau \mu_u^s f_u(\mathbf{x}). \tag{3.2}$$

Here, U is the overall unknown target and in this case the rate of new targets, o_u^{v} is the probability of visibility state v, ξ_u^τ is the probability of ID τ, μ_u^τ is the probability of the kinematic mode s, and $f_u(\mathbf{x})$ is the distribution of the kinematic state. The subscript u indicates that the individual expressions are part of the unknown target intensity. Equation (3.2) does not contain the initial values of new targets, as it is a function of the birth intensity and the transition probability matrices. However, for simplicity, the unknown target values are tuned directly and can be viewed as initial values.

Remark 3.1. This method of modeling the target IDs through a Marked PPP implies that two targets can have the same ID. The probability of two targets having the same ID in a surveillance area with relatively few targets is minuscule, but it is nevertheless a possibility [31]. We also note how the modeling of actual, observable IDs here deviates from theoretically assigned IDs. The labels in labeled Random finite sets (RFSs), introduced in [132], are unobservable and analogous to the identifying tags in [50], which ensure the uniqueness of the elements of a RFS. The IDs described here, however, serve no such purpose and can be assumed non-unique without breaking the underlying mathematical assumptions of RFSs. ◁

M2

We model the survival probability as a function of time since the last update. A constant parameter P_{S_c} denotes the probability of survival after one second. Thus, the survival probability of an interval between times t_{k-1} and t_k, denoted as dt, becomes

$$P_{\mathrm{S}}(\mathrm{dt}) = P_{S_c}^{\mathrm{dt}}. \tag{3.3}$$

M3

The ID numbers τ are assumed to be static, in line with the physical reality of the AIS protocol. The IDs are manually set at the installation of the AIS system. We assume that the ID numbers of the unknown targets are distributed according to

$$\xi_u^\tau = \begin{cases} \xi_u^0 & \text{if } \tau = 0 \\ \dfrac{1-\xi_u^0}{|\mathcal{V}|-1} & \text{if } \tau > 0 \end{cases} \tag{3.4}$$

where ξ_u^0 is some parameter denoting the belief that the target has no ID, and $|\mathcal{V}|$ is the number of all possible ID numbers in addition to 0. Not all targets have an ID, and we represent this non-ID by the value $\tau = 0$. If $\tau = 0$, the target does not transmit measurements.

M4

From time step $k-1$ to k, the evolution of a target is given by

$$f_{\mathbf{y}}(\mathbf{y}_k|\mathbf{y}_{k-1}) = f_{\mathbf{x}}^{s\tau}(\mathbf{x}_k|\mathbf{x}_{k-1})\pi^{s_{k-1}s_k} w^{\mathrm{v}_{k-1}\mathrm{v}_k}. \tag{3.5}$$

The π-matrix contains the Markov chain probabilities of changing between different kinematic models. The matrix w contains the Markov chain probabilities of the target switching between the visible state $\mathrm{v} = 1$ and invisible state $\mathrm{v} = 0$. The ID numbers are assumed static and therefore do not change during a prediction.

M5

For radar measurements, the detection probability $P_{\mathrm{D}}(\mathbf{y}_k)$ varies based on the visibility state v and we define it as

$$P_{\mathrm{D}}(\mathbf{y}_k) = \begin{cases} P_{\mathrm{D}} & \text{if v = 1} \\ 0 & \text{if v = 0} \end{cases} \tag{3.6}$$

where P_{D} is a constant describing the probability of a target being detected by the radar at a given time step.

For target-provided measurements, which are assumed to give no missed detections, we have that

$$P_{\mathrm{D}}(\mathbf{y}_k) = \begin{cases} 1 & \text{if a target-provided measurement is received} \\ 0 & \text{otherwise} \end{cases} \tag{3.7}$$

independent of the visibility state. Thus, no conclusions about a target are made from the absence of target-provided measurements. Trying to keep track of when a vessel should transmit measurements is a difficult problem, which, for example, would be subject to intentional randomness from the protocol [16].

M6

Radar clutter measurements are assumed to follow a Poisson process with intensity λ. The target-provided measurements do not contain clutter, the same as if it is following a Poisson process with intensity 0.

M7

The radar measurements are assumed to be synchronized and to arrive simultaneously at a fixed frequency. The synchronicity means that when radar measurements arrive at time step k, the set of radar measurements contains measurements from all detected targets at time step k, in addition to clutter measurements. The radar measurement likelihood is denoted as $f_{\mathbf{z}}^{R}(\mathbf{z}_k|\mathbf{y}_k)$.

M8

The target-provided measurements can arrive whenever and are not synchronized. Thus, a transmitted measurement can be received at any time from any target. We do not assume that targets transmit measurements simultaneously, contrary to what we do for radar measurements. Whenever a target-provided measurement arrives, however, the time of arrival is assumed to be known. The measurement likelihood for the target-provided measurements is

$$f_{\mathbf{z}}^{A}(\mathbf{z}_k|\mathbf{y}_k) = f_{\mathbf{p}}(\mathbf{p}_k|\mathbf{y}_k) f_{\tau}(\tau^{\mathbf{z}_k}|\tau) \tag{3.8}$$

where $\mathbf{z}_k$ is the whole measurement and $\mathbf{p}_k$ only contains the kinematic data of the measurement. Furthermore

$$f_{\tau}(\tau^{\mathbf{z}_k}|\tau) = \begin{cases} P_C & \text{if } \tau_k = \tau_k^{\mathbf{z}_k} \\ \dfrac{1-P_C}{|\mathcal{V}|-1} & \text{if } \tau_k \neq \tau_k^{\mathbf{z}_k} \text{ and } \tau > 0 \\ 0 & \text{if } \tau = 0 \end{cases} \tag{3.9}$$

where P_C is a fixed parameter describing the confidence in the ID number not being corrupted, denoted as the confidence probability. The reasoning behind the above equation comes from the observation that the likelihood of a transmitted measurement coming from a target without an ID is zero. Furthermore, the chance of a transmitted ID being erroneous makes it a possibility, albeit small, that any ID can be the correct one.

3.4 Hybrid states and the PMBM

As formulated in [8, p. 441], a hybrid state is a state where the state space contains both discrete and continuous states or uncertainties. This structure is useful as the kinematic state will be continuous, while, for example, the choice of kinematic model for the target will be discrete.

A PMBM filter represents the posterior multi-target density for discovered targets as a weighted sum of multi-Bernoulli densities. These involve weights for each of the multi-Bernoullis, and kinematic densities and existence probabilities for each of the Bernoullis. The PMB filter, which essentially is the same as a JIPDA, approximates the sum of multi-Bernoullis by a single multi-Bernoulli at the end of each estimation cycle.

Using the equations from [138], one can get general expressions for the weight, existence, and states irrespective of the sensor type assuming the sensors generate measurements adhering to the assumptions made in Assumption 2 in [138]. The assumptions hold for both target-provided and radar measurements. The inclusion of IDs in the target-provided measurements is contained in the measurement likelihood function, and they do not breach any independence assumptions. The goal of this section is to extract expressions for the probabilistic properties of the individual hybrid state elements.

From [138], we have that the weight w, existence probability r and distribution $f(\mathbf{y})$ of a single Bernoulli in general can be written as

$$w = g(\mathbf{y}) + h[1] \tag{3.10}$$

$$r = \frac{h[1]}{g(\mathbf{y}) + h[1]} \tag{3.11}$$

$$f(\mathbf{y}) = \frac{h(\mathbf{y})}{h[1]}. \tag{3.12}$$

for some functions g and h of the state $\mathbf{y}$. The notation $[\cdot]$ indicates a linear functional, defined as

$$g[h] = \int g(\mathbf{x})h(\mathbf{x})d\mathbf{x}. \tag{3.13}$$

These are useful tools for compactly writing normalization constants and likelihoods. For later use, it is convenient to find general expressions for the individual states in the hybrid state $\mathbf{y}$. Using the approximation from [24, Remark 6] that the visibility is independent on the other states, we can write $h(\mathbf{y}) = h(\mathrm{v})h(\tau)h(s|\tau)h(\mathbf{x}|\tau, s)$. We get the individual states by using the rule of conditional probability. Starting with the kinematic state $\mathbf{x}$, it can be acquired by

$$\begin{aligned} f^t(\mathbf{x}|s,\tau,\mathrm{v}) &= \frac{f(\mathbf{x},s,\tau,\mathrm{v})}{\int f(\tilde{\mathbf{x}},s,\tau,\mathrm{v})\mathrm{d}\tilde{\mathbf{x}}} \\ &= \frac{\frac{h(\mathbf{x},s,\tau,\mathrm{v})}{h[1]}}{\frac{\int h(\tilde{\mathbf{x}},s,\tau,\mathrm{v})\mathrm{d}\tilde{\mathbf{x}}}{h[1]}} \\ &= \frac{h(\mathbf{x},s,\tau,\mathrm{v})}{\int h(\mathbf{x},s,\tau,\mathrm{v})\mathrm{d}\mathbf{x}} \\ &= \frac{h(\mathrm{v})h(\mathbf{x},s,\tau)}{h(\mathrm{v})\int h(\tilde{\mathbf{x}},s,\tau,\mathrm{v})\mathrm{d}\tilde{\mathbf{x}}} \\ &= \frac{h(\mathbf{x},s,\tau)}{h(s,\tau)} \end{aligned} \tag{3.14}$$

where we have omitted the time indices for brevity. The $\tilde{(\cdot)}$ notation is used for latent variables which disappear by marginalization. Furthermore, the absence of the visibility state v in the final expression means that $f^t(\mathbf{x}|s,\tau,\mathrm{v}) = f^t(\mathbf{x}|s,\tau)$. Similarly, the mode probabilities are

$$f^t(s|\tau) = \mu^{t\tau s} = \frac{h(s,\tau)}{h(\tau)}, \tag{3.15}$$

the ID probabilities are

$$f^t(\tau) = \xi^{t\tau} = \frac{h(\tau)}{h[1]}. \tag{3.16}$$

and the visibility probabilities are

$$f^t(\mathrm{v}) = o^{t\mathrm{v}} = \frac{h(\mathrm{v})}{h[1]}. \tag{3.17}$$

Note that $\sum_{\tilde{\tau}} \sum_{\tilde{s}} \int h(\tilde{\mathbf{x}}, \tilde{s}, \tilde{\tau}) \mathrm{d}\tilde{\mathbf{x}} = h[1]$, which essentially acts as a normalization constant. Independencies between the states will make it possible to reduce the needed amount of marginalization, as they will appear both in the numerator and the denominator. The independencies will depend on the model choices and are written here according to the assumptions in Section 3.3.

3.5 Including target-provided measurements in the VIMMJIPDA

Table 3.1: Expressions for creating, updating, and predicting the Bernoulli components.

New targets:	g	$=$	λ
	$h[1]$	$=$	$u[P_D(\tilde{\mathbf{y}}) f_{\mathbf{z}}(\mathbf{z}\|\tilde{\mathbf{y}})]$
	$h(\mathbf{y})$	$=$	$u(\mathbf{y}) P_D(\mathbf{y}) f_{\mathbf{z}}(\mathbf{z}\|\mathbf{y})$
Missed detection:	g	$=$	$1 - r^t_{k\|k-1}$
	$h[1]$	$=$	$r^t_{k\|k-1} f[1 - P_D(\tilde{\mathbf{y}})]$
	$h(\mathbf{y})$	$=$	$r^t_{k\|k-1} f_{k\|k-1}(\mathbf{y})(1 - P_D(\mathbf{y}))$
Detection:	g	$=$	0
	$h[1]$	$=$	$r^t_{k\|k-1} f[P_D(\tilde{\mathbf{y}}) f_{\mathbf{z}}(\mathbf{z}\|\tilde{\mathbf{y}})]$
	$h(\mathbf{y})$	$=$	$r^t_{k\|k-1} f^t_{k\|k-1}(\mathbf{y}) P_D(\mathbf{y}) f_{\mathbf{z}}(\mathbf{z}\|\mathbf{y})$
Prediction	g	$=$	$1 - r^t_{k-1} f[P_S(\tilde{\mathbf{y}})]$
	$h[1]$	$=$	$r^t_{k-1} f[P_S(\tilde{\mathbf{y}})]$
	$h(\mathbf{y})$	$=$	$r^t_{k-1} \int f^t_{k\|k-1}(\mathbf{y}\|\tilde{\mathbf{y}}) P_S(\tilde{\mathbf{y}}) f_{k-1}(\tilde{\mathbf{y}}) \mathrm{d}\tilde{\mathbf{y}}$

In the VIMMJIPDA, the unknown target intensity $u(\mathbf{y})$ is assumed stationary and is left unchanged during the prediction and updating of the estimates. We make the same assumption here. This assumption means that only the Bernoulli components have to be considered, and is further simplified by following the JIPDA method

of performing mixture reduction. That is, we merge all Bernoullis originating in the same measurement into a single Bernoulli after each update. Thus, we can omit the weights of the association hypotheses of previous time steps can due to marginalization. Table 3.1 shows the expressions for updating and predicting the Bernoulli components from [138]. These are adapted to simplify insertion in (3.10)-(3.12) and (3.14)-(3.17). Furthermore, they are simplified to reflect the stationary unknown target intensity and the marginalization over the weights during mixture reduction. As the measurement model assumptions made in [138] hold with regards to both radar and target-provided measurements, both $f_{\mathbf{z}}^{R}(\mathbf{z}|\mathbf{y})$ and $f_{\mathbf{z}}^{A}(\mathbf{z}|\mathbf{y})$ can be considered special cases of the more general $f_{\mathbf{z}}(\mathbf{z}|\mathbf{y})$ in the table. The expressions for predicting and updating the Bernoulli estimates based on the potential information acquired by the sensor updates follow.

3.5.1 Prior

For a single track, which in the context of this chapter is analogous to a Bernoulli, we write the hybrid state prior distribution as

$$f_{k-1}^{t}(\mathbf{y}) = f_{k-1}^{t}(\mathbf{x}|\tau, s)\xi_{k-1}^{t\tau}\mu_{k-1}^{t\tau s}o_{k-1}^{t\mathrm{v}} \tag{3.18}$$

while the prior existence probability is r_{k-1}^{t}. As mentioned above, we merge all the hypotheses of the previous time step, giving $w_{k-1}^{t} = 1$. The prior is a joint distribution over the continuous kinematic state and the discrete potential IDs, kinematic modes, and visibility states. In the following propositions, only the probability of the target being in the visible state is presented, that is, o^{t1}, which we denote as η^{t}. The prior is decomposed into several states conditioned on the different discrete states. An example of the structure of a prior with two possible IDs and two possible kinematic modes is shown in Figure 3.1. The expressions in the square boxes are not calculated themselves but can be constructed from the other expressions.

3.5.2 Prediction

All tracks are predicted from the previous time step $k-1$ to the current time step k. The predicted probabilities and densities are denoted by the subscript $(\cdot)_{k|k-1}$.

Proposition 3.1. The prediction for the existence probability r^{t}, the visibility probability η^{t}, the ID probabilities $\xi^{t\tau}$, the mode probabilities $\mu^{t\tau s}$ and the kinematic density $f^{t}(\mathbf{x}|\tau, s)$ are done as

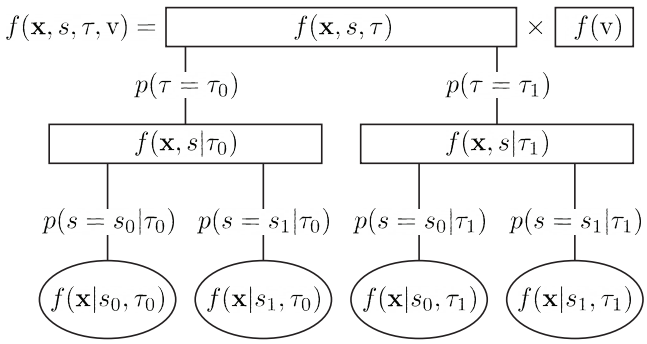

Figure 3.1: The structure of the distribution of a hybrid state with two kinematic modes and two possible IDs.

$$r^t_{k|k-1} = r^t_{k-1} P_{\mathrm{S}}(\mathrm{dt}) \tag{3.19}$$

$$\eta^t_{k|k-1} = (1 - \eta^t_{k-1}) w^{01} + \eta^t_{k-1} w^{11} \tag{3.20}$$

$$\xi^{t\tau}_{k|k-1} = \xi^{t\tau}_{k-1} \tag{3.21}$$

$$\mu^{t\tau s}_{k|k-1} = \sum_{\tilde{s}} \mu^{t\tau\tilde{s}}_{k-1} \pi^{\tilde{s}s}(\mathrm{dt}) \tag{3.22}$$

$$f^t_{k|k-1}(\mathbf{x}|\tau, s) = \int f_{\mathbf{y}}(\mathbf{x}|\tau, s, \tilde{\mathbf{x}}) f^t_{k-1}(\tilde{\mathbf{x}}|\tau, s) \mathrm{d}\tilde{\mathbf{x}} \tag{3.23}$$

where

$$f^t_{k-1}(\tilde{\mathbf{x}}|\tau, s) = \sum_{\tilde{s}} \frac{\mu^{t\tau\tilde{s}}_{k-1} \pi^{\tilde{s}s} f^t_{k-1}(\tilde{\mathbf{x}}|\tau, \tilde{s})}{\sum_{\tilde{s}} \mu^{t\tau\tilde{s}}_{k-1} \pi^{\tilde{s}s}(\mathrm{dt})} \tag{3.24}$$

.

Proof. The proof builds upon Chapter 2, but is modified to also account for the inclusion of the IDs in the state vector. It should be noted that the survival probability is only dependent on the times of the measurements' arrival, which are independent of the state. Because the IDs are assumed static the transition model for the IDs becomes a Kronecker delta $\delta_{\tau\tilde{\tau}}$. It is defined as

$$\delta_{\tau\tilde{\tau}} = \begin{cases} 1 & \text{if } \tau = \tilde{\tau} \\ 0 & \text{if } \tau \neq \tilde{\tau} \end{cases}. \tag{3.25}$$

First we write out $h(\mathbf{y})$ from Table 3.1:

$$
\begin{aligned}
h(\mathbf{y}) &= r_{k-1}^t \int f_{k|k-1}^t(\mathbf{y}|\tilde{\mathbf{y}}) P_{\mathrm{S}}(\tilde{\mathbf{y}}) f_{k-1}(\tilde{\mathbf{y}}) \mathrm{d}\tilde{\mathbf{y}} \\
&= r_{k-1}^t P_{\mathrm{S}}(\mathrm{dt}) \Big(\sum_{\tilde{\mathrm{v}}} f(\tilde{\mathrm{v}}) f(\mathrm{v}|\tilde{\mathrm{v}}) \Big) \sum_{\tilde{\tau}} f_{k-1}(\tilde{\tau}) \delta_{\tau\tilde{\tau}} \times \\
&\sum_{\tilde{s}} f_{k-1}(\tilde{s}|\tilde{\tau}) f_{k|k-1}^t(s|\tilde{s}) \int f_{k|k-1}^t(\mathbf{x}|s,\tau,\tilde{\mathbf{x}}) f_{k-1}(\tilde{\mathbf{x}}|\tilde{s},\tilde{\tau}) \mathrm{d}\tilde{\mathbf{x}} \\
&= r_{k-1}^t P_{\mathrm{S}}(\mathrm{dt}) \Big(\sum_{\tilde{\mathrm{v}}} f(\tilde{\mathrm{v}}) f(\mathrm{v}|\tilde{\mathrm{v}}) \Big) f_{k-1}(\tau) \times \\
&\times \sum_{\tilde{s}} f_{k-1}(\tilde{s}|\tau) f_{k|k-1}^t(s|\tilde{s}) \int f_{k|k-1}^t(\mathbf{x}|s,\tau,\tilde{\mathbf{x}}) \times f_{k-1}(\tilde{\mathbf{x}}|\tilde{s},\tau) \mathrm{d}\tilde{\mathbf{x}} \\
&= r_{k-1}^t P_{\mathrm{S}}(\mathrm{dt}) \Big(\sum_{\tilde{\mathrm{v}}} o_{k-1}^{t\tilde{\mathrm{v}}} w^{\tilde{\mathrm{v}}\mathrm{v}} \Big) \xi_{k-1}^{t\tau} \sum_{\tilde{s}} \mu_{k-1}^{\tau s} \pi^{\tilde{s}s}(\mathrm{dt}) \times \\
&\int f_{k|k-1}^{t\tau s}(\mathbf{x}|s,\tau,\tilde{\mathbf{x}}) f_{k-1}(\tilde{\mathbf{x}}|\tilde{s},\tau) \mathrm{d}\tilde{\mathbf{x}}
\end{aligned} \tag{3.26}
$$

which uses the fact that only the conditioning on the most recent variable is relevant. Marginalizing this, one gets

$$
\begin{aligned}
h(s,\tau) &= r_{k-1}^t \int \sum_{\mathrm{v}} h(\mathbf{x},s,\tau,\mathrm{v}) \mathrm{d}\mathbf{x} \\
&= r_{k-1}^t P_{\mathrm{S}}(\mathrm{dt}) \xi_{k-1}^{t\tau} \sum_{\tilde{s}} \mu_{k-1}^{t\tau\tilde{s}} \pi^{\tilde{s}s}(\mathrm{dt}),
\end{aligned} \tag{3.27}
$$

$$
h(\tau) = r_{k-1}^t \sum_{s} h(s,\tau) = P_{\mathrm{S}}(\mathrm{dt}) \xi_{k-1}^{t\tau}, \tag{3.28}
$$

$$
\begin{aligned}
h(\mathrm{v}) &= r_{k-1}^t \int \sum_{\tau} \sum_{s} h(\mathbf{x},s,\tau,\mathrm{v}) \mathrm{d}\mathbf{x} \\
&= r_{k-1}^t P_{\mathrm{S}}(\mathrm{dt}) \Big(\sum_{\tilde{\mathrm{v}}} o_{k-1}^{t\tilde{\mathrm{v}}} w^{\tilde{\mathrm{v}}\mathrm{v}} \Big)
\end{aligned} \tag{3.29}
$$

$$
h[1] = r_{k-1}^t \sum_{\tau} h(\tau) = r_{k-1}^t P_{\mathrm{S}}(\mathrm{dt}). \tag{3.30}
$$

Inserting this in (3.14)-(3.17) provides the expressions for the hybrid states. Note that the expression for the visibility probability $\eta_{k|k-1}^t$ follows from the fact that $o_{k-1}^{t0} = 1 - o_{k-1}^{t1} = 1 - \eta_{k-1}^t$. The expression for the existence probability $r_{k|k-1}^t$ is found by inserting $g(\mathbf{y}) = r_{k-1}^t P_{\mathrm{S}}(\mathrm{dt})$ from Table 3.1 and $h[1]$ into (3.11). □

3.5.3 Posterior

The individual posterior distributions, conditioned on either a detection or a missed detection, are calculated after the prediction. The four possibilities for a track when new measurements arrive are

- The previously unknown track is detected for the first time.
- The previously detected track is detected again.
- The previously detected track is not detected.
- The previously unknown track is not detected.

Any tracks covered by the fourth alternative will be represented by the unknown target density, and do not need to be considered specifically. The posterior distributions for the three first possibilities are presented in the following propositions.

Proposition 3.2. Initialization of a new track on a measurement indexed by j is done as

$$w_k^{tj} = \begin{cases} \lambda + cUP_{\mathrm{D}}\eta^0 & \text{for radar} \\ cU\sum_{\tilde{\tau}} \xi_u^{\tilde{\tau}} f_\tau(\tau^j|\tilde{\tau}) & \text{for target-provided} \end{cases} \tag{3.31}$$

$$r_k^{tj} = \begin{cases} \dfrac{UP_{\mathrm{D}}\eta^0}{\lambda + UP_{\mathrm{D}}\eta^0} & \text{for radar} \\ 1 & \text{for target-provided} \end{cases} \tag{3.32}$$

$$\eta_k^{tj} = \begin{cases} 1 & \text{for radar} \\ \eta_u & \text{for target-provided} \end{cases} \tag{3.33}$$

$$\xi_k^{t\tau j} = \begin{cases} \xi_u^\tau & \text{for radar} \\ f_\tau(\tau^{\mathbf{z}}|\tau) & \text{for target-provided} \end{cases} \tag{3.34}$$

$$\mu_k^{t\tau sj} = \mu_u^s \tag{3.35}$$

$$f_k^{tj}(\mathbf{x}|s,\tau) = f_{\mathbf{z}}(\mathbf{z}|\mathbf{x},s,\tau) f_u(\mathbf{x})/c \tag{3.36}$$

where $c = \int f_{\mathbf{z}}(\mathbf{z}|\mathbf{x},s,\tau) f_u(\mathbf{x}) \mathrm{d}\mathbf{x}$ is a constant.

Proof. Firstly, for radar measurements, we have that

$$h(\mathbf{y}) = UP_{\mathrm{D}}(\mathrm{v}) o_u^{\mathrm{v}} \xi_u^\tau \mu_u^{\tau s} f_u(\mathbf{x}) f_{\mathbf{z}}(\mathbf{z}|\mathbf{x},s,\tau) \tag{3.37}$$

which follows from (3.2) and Table 3.1. Furthermore

$$h(s,\tau,\mathrm{v}) = cUP_{\mathrm{D}}(\mathrm{v})o_u^{\mathrm{v}}\xi_u^{\tau}\mu_u^{\tau s} \tag{3.38}$$
$$h(\tau,\mathrm{v}) = cUP_{\mathrm{D}}(\mathrm{v})o_u^{\mathrm{v}}\xi_u^{\tau} \tag{3.39}$$
$$h(\mathrm{v}) = cUP_{\mathrm{D}}(\mathrm{v})o_u^{\mathrm{v}} \tag{3.40}$$
$$h[1] = cUP_{\mathrm{D}}\eta^0 \tag{3.41}$$

where c is a constant resulting from the marginalization over $\mathbf{x}$. For target-provided measurements we have that

$$f_{\mathbf{z}}(\mathbf{z}|\mathbf{x},s,\tau) = f_{\mathbf{p}}(\mathbf{p}|\mathbf{x},s,\tau)f_{\tau}(\tau^{\mathbf{z}}|\tau) \tag{3.42}$$

which means that

$$h(\mathbf{y}) = Uo_u^{\mathrm{v}}\xi_u^{\tau}\mu_u^{\tau s}f_{\tau}(\tau^{\mathbf{z}}|\tau)f_u(\mathbf{x})f_{\mathbf{z}}(\mathbf{p}|\mathbf{x},s,\tau). \tag{3.43}$$

The probability of detection is omitted here, as it is defined as 1 whenever a target-provided measurement has been received. Furthermore

$$h(s,\tau,\mathrm{v}) = cUo_u^{\mathrm{v}}\xi_u^{\tau}\mu_u^{\tau s}f_{\tau}(\tau^{\mathbf{z}}|\tau) \tag{3.44}$$
$$h(\tau) = cU\xi_u^{\tau}f_{\tau}(\tau^{\mathbf{z}}|\tau) \tag{3.45}$$
$$h(\mathrm{v}) = cUo_u^{\mathrm{v}}\textstyle\sum_{\tilde{\tau}}\xi_u^{\tilde{\tau}}f_{\tau}(\tau^{\mathbf{z}}|\tilde{\tau}) \tag{3.46}$$
$$h[1] = cU\textstyle\sum_{\tilde{\tau}}\xi_u^{\tilde{\tau}}f_{\tau}(\tau^{\mathbf{z}}|\tilde{\tau}) \tag{3.47}$$

where c again is a constant.

Inserting these expressions in (3.14)-(3.17) gives (3.33)-(3.36), that isthe distributions of the individual hybrid states of a new target. Furthermore, we have from Table 3.1 that g is the clutter density, which is λ for radar measurements, and 0 for target-provided measurements. We insert g and $h[1]$ in (3.10) and (3.11) to get (3.31) and (3.32). The expression for the ID probability in the event of initialization on a transmitted measurement requires some further explanation. Keeping in mind the prior distribution for the IDs (3.4), we have that

$$
\begin{aligned}
\xi_k^{t\tau j} &= \frac{h(\tau)}{h[1]} \\
&= \frac{\xi_u^{\tau} f_\tau(\tau^{\mathbf{z}}|\tau)}{\sum_{\tilde{\tau}} \xi_u^{\tilde{\tau}} f_\tau(\tau^{\mathbf{z}}|\tilde{\tau})} \\
&= \begin{cases} \dfrac{f_\tau(\tau^{\mathbf{z}}|\tau)(1-\xi_u^0)/|\mathcal{V}-1|}{\sum_{\tilde{\tau}} f_\tau(\tau^{\mathbf{z}}|\tilde{\tau})(1-\xi_u^0)/|\mathcal{V}-1|} & \text{if } \tau > 0 \\ 0 & \text{if } \tau = 0 \end{cases} \\
&= \begin{cases} \dfrac{f_\tau(\tau^{\mathbf{z}}|\tau)}{\sum_{\tilde{\tau}} f_\tau(\tau^{\mathbf{z}}|\tilde{\tau})} & \text{if } \tau > 0 \\ 0 & \text{if } \tau = 0 \end{cases} \\
&= \begin{cases} f_\tau(\tau^{\mathbf{z}}|\tau) & \text{if } \tau > 0 \\ 0 & \text{if } \tau = 0 \end{cases} = f_\tau(\tau^{\mathbf{z}}|\tau)
\end{aligned} \tag{3.48}
$$

where we have used that $\sum_{\tilde{\tau}} f_\tau(\tau^{\mathbf{z}}|\tilde{\tau}) = 1$. If a different prior distribution than (3.4) is used for the IDs, it can be accommodated by replacing the final expression with the one in the second line of the above expression. □

Proposition 3.3. Updating based on a missed detection is done as

$$
w_k^{t0} = \begin{cases} 1 - r_{k|k-1}^t \eta_{k|k-1}^t P_{\mathrm{D}} & \text{for radar} \\ 1 & \text{for target-provided} \end{cases} \tag{3.49}
$$

$$
r_k^{t0} = \begin{cases} \dfrac{r_{k|k-1}^t (1 - \eta_{k|k-1}^t P_{\mathrm{D}})}{1 - r_{k|k-1}^t \eta_{k|k-1}^t P_{\mathrm{D}}} & \text{for radar} \\ r_{k|k-1}^t & \text{for target-provided} \end{cases} \tag{3.50}
$$

$$
\eta_k^{t0} = \begin{cases} \dfrac{(1 - P_{\mathrm{D}})\eta_{k|k-1}^t}{1 - P_{\mathrm{D}} \eta_{k|k-1}^t} & \text{for radar} \\ \eta_{k|k-1}^t & \text{for target-provided} \end{cases} \tag{3.51}
$$

$$
\xi_k^{t\tau 0} = f_{k|k-1}^t(\tau) \tag{3.52}
$$

$$
\mu_k^{t\tau s 0} = f_{k|k-1}^t(s|\tau) \tag{3.53}
$$

$$
f_k^{t0}(\mathbf{x}|\tau, s) = f_{k|k-1}^t(\mathbf{x}|\tau, s) \tag{3.54}
$$

Remark 3.2. The inclusion of target-provided measurement types in the case of a missed detection is somewhat artificial. The expressions are the same as for the prediction, as the absence of target-provided measurements gives no additional

information to the tracking algorithm. This follows from the definition of the detection probability in Section 3.3, that isthat $P_{\mathrm{D}} = 0$ for target-provided measurements when they have not been received. For later use, the expressions are nevertheless written out here. ◁

Proof. We have that

$$h(\mathbf{y}) = r^t_{k|k-1}(1 - P_{\mathrm{D}}(\mathrm{v}))o^{t\mathrm{v}}_{k|k-1}\xi^{t\tau}_{k|k-1}\mu^{t\tau s}_{k|k-1}f_{k|k-1}(\mathbf{x}|s,\tau) \tag{3.55}$$

where the corresponding expression from Table 3.1 has been written out. Similarly as to what was done previously, we find through marginalization that:

$$\begin{aligned} h(s,\tau,\mathrm{v}) &= r^t_{k|k-1}(1 - P_{\mathrm{D}}(\mathrm{v}))o^{t\mathrm{v}}_{k|k-1}\xi^{t\tau}_{k|k-1}\mu^{t\tau s}_{k|k-1} \\ h(\tau,\mathrm{v}) &= r^t_{k|k-1}(1 - P_{\mathrm{D}}(\mathrm{v}))o^{t\mathrm{v}}_{k|k-1}\xi^{t\tau}_{k|k-1} \\ h(\mathrm{v}) &= r^t_{k|k-1}(1 - P_{\mathrm{D}}(\mathrm{v}))o^{t\mathrm{v}}_{k|k-1} \end{aligned} \tag{3.56}$$

Again, the different detection probabilities have to be taken into account when summing over the visibility states, giving

$$\begin{aligned} h[1] &= r^t_{k|k-1}((1 - P_{\mathrm{D}})\eta^t_{k|k-1} + (1 - \eta^t_{k|k-1})) \\ &= r^t_{k|k-1}(1 - P_{\mathrm{D}}\eta^t_{k|k-1}) \end{aligned} \tag{3.57}$$

for radar updates, and $h[1] = 1$ for AIS updates. Inserting this in (3.14)-(3.17) gives the wanted expressions for the hybrid states. Furthermore, we get from Table 3.1 that g is given by $1 - r^t_{k|k-1}$, which together with $h[1]$ gives us (3.49) and (3.50) by using (3.10) and (3.11). □

Proposition 3.4. Updating based on a detection is done as

$$w^{tj}_k = \begin{cases} P_{\mathrm{D}}\eta^t_{k|k-1}r^t_{k|k-1}\sum_{\tilde{\tau}}\xi^{t\tilde{\tau}}_{k|k-1}\sum_{\tilde{s}}\mu^{t\tilde{\tau}\tilde{s}}_{k|k-1}l^{t\tilde{\tau}\tilde{s}j} & \text{for radar} \\ r^t_{k|k-1}\sum_{\tilde{\tau}}\xi^{t\tilde{\tau}}_{k|k-1}\sum_{\tilde{s}}\mu^{t\tilde{\tau}\tilde{s}}_{k|k-1}l^{t\tilde{\tau}\tilde{s}j} & \text{for target-provided} \end{cases} \tag{3.58}$$

$$r^{tj}_k = 1 \tag{3.59}$$

$$\eta_k^{tj} = \begin{cases} 1 & \text{for radar} \\ \eta_{k|k-1}^t & \text{for target-provided} \end{cases} \tag{3.60}$$

$$\xi_k^{t\tau j} = \frac{\xi_{k|k-1}^{t\tau} \sum_{\tilde{s}} l^{t\tau\tilde{s}j}}{\sum_{\tilde{\tau}} \xi_{k|k-1}^{t\tilde{\tau}} \sum_{\tilde{s}} l^{t\tilde{\tau}\tilde{s}j}} \tag{3.61}$$

$$\mu_k^{t\tau sj} = \frac{\mu_{k|k-1}^{t\tau s} l^{t\tau sj}}{\sum_{\tilde{s}} \mu_{k|k-1}^{t\tau\tilde{s}} l^{t\tau\tilde{s}j}} \tag{3.62}$$

$$f_k^{tj}(\mathbf{x}|\tau, s) = \frac{f_{\mathbf{z}}(\mathbf{z}|\mathbf{x}, \tau, s) f_{k|k-1}^t(\mathbf{x}|\tau, s)}{l^{t\tau sj}} \tag{3.63}$$

where

$$l^{t\tau sj} = f_\tau(\tau^j|\tau) \int f_{\mathbf{z}}(\mathbf{z}_k^j|\tilde{\mathbf{x}}) f_{k|k-1}^{t\tau s}(\tilde{\mathbf{x}}) \mathrm{d}\tilde{\mathbf{x}}. \tag{3.64}$$

for target-provided measurements, and

$$l^{t\tau sj} = \int f_{\mathbf{z}}(\mathbf{z}_k^j|\tilde{\mathbf{x}}) f_{k|k-1}^{t\tau s}(\tilde{\mathbf{x}}) \mathrm{d}\tilde{\mathbf{x}}. \tag{3.65}$$

for radar measurements.

Proof. Writing out the expression for a detection in Table 3.1, we have that

$$h(\mathbf{y}) = r_{k|k-1}^t P_{\mathrm{D}}(\mathrm{v}) o_{k|k-1}^{t\mathrm{v}} \xi_{k|k-1}^{t\tau} \mu_{k|k-1}^{t\tau s} f_{k|k-1}^{t\tau s}(\mathbf{x}) f_{\mathbf{z}}(\mathbf{z}|\mathbf{x}, s, \tau) \tag{3.66}$$

which we marginalize to obtain

$$\begin{aligned} h(s, \tau, \mathrm{v}) &= r_{k|k-1}^t P_{\mathrm{D}}(\mathrm{v}) o_{k|k-1}^{t\mathrm{v}} \xi_{k|k-1}^{t\tau} \mu_{k|k-1}^{t\tau s} l^{t\tau sj} \\ h(\tau, \mathrm{v}) &= r_{k|k-1}^t P_{\mathrm{D}}(\mathrm{v}) o_{k|k-1}^{t\mathrm{v}} \xi_{k|k-1}^{t\tau} \sum_s \mu_{k|k-1}^{t\tau s} l^{t\tau sj} \\ h(\mathrm{v}) &= r_{k|k-1}^t P_{\mathrm{D}}(\mathrm{v}) o_{k|k-1}^{t\mathrm{v}} \sum_\tau \xi_{k|k-1}^{t\tau} \sum_s \mu_{k|k-1}^{t\tau s} l^{t\tau sj}. \end{aligned} \tag{3.67}$$

For radar, we have that $P_{\mathrm{D}}(\mathrm{v} = 1) = P_{\mathrm{D}}$ and 0 otherwise, and for AIS $P_{\mathrm{D}}(\mathrm{v}) = P_{\mathrm{D}} = 1$ if a measurement has been received. Using this we get

$$h[1] = P_{\mathrm{D}} \eta_{k|k-1}^t r_{k|k-1}^t \sum_\tau \xi_{k|k-1}^{t\tau} \sum_s \mu_{k|k-1}^{t\tau s} l^{t\tau sj} \tag{3.68}$$

for radar updates, and

$$h[1] = r^t_{k|k-1} \sum_{\tau} \xi^{t\tau}_{k|k-1} \sum_{s} \mu^{t\tau s}_{k|k-1} l^{t\tau sj} \tag{3.69}$$

for AIS updates. The expressions for the hybrid states result from inserting this in (3.14)-(3.17). We see from Table 3.1 that $g = 0$, and using this, together with $h[1]$ we get (3.58) and (3.59) from (3.10) and (3.11). □

3.5.4 Mixture reduction

The mixture reduction is done similarly to what is done in the JIPDA. That is, all the association hypotheses for each track are merged. An association hypothesis $\mathbf{a}_k$ from the set of all possible association hypotheses $\mathcal{A}_k$ contains individual track-to-measurement associations a^t. The probabilities for the individual association hypotheses are

$$\Pr(\mathbf{a}_k) \propto \prod_{t \ s.t \ a^t=0} w_k^{ta^t} \prod_{t \ s.t \ a^t>0} w_k^{ta^t}/\lambda \tag{3.70}$$

where λ is the Poisson intensity for the false alarms, and the fact that

$$\sum_{\mathbf{a}_k \in \mathcal{A}_k} \Pr(\mathbf{a}_k) = 1 \tag{3.71}$$

is used to normalize the probabilities. This in turn provides the marginal probabilities for the associations as

$$p_k^{tj} = \sum_{\mathbf{a}_k \ s.t. \ a^t=j} \Pr(\mathbf{a}_k). \tag{3.72}$$

The mixture reduction remains the same irrespective of the type of measurement, as all differences are handled during the calculation of the individual posterior distributions.

Proposition 3.5. We have that

$$r_k^t = \sum_{j=0}^{m_k} r_k^{tj} p^{tj} \tag{3.73}$$

$$\eta_k^t = \sum_{j=0}^{m_k} \underbrace{\frac{1}{r_k^t} r_k^{tj} p_k^{tj}}_{\beta_k^{tj}} \eta_k^{tj} \tag{3.74}$$

$$\xi_k^{t\tau} = \sum_{j=0}^{m_k} \underbrace{\frac{1}{r_k^t} r_k^{tj} p_k^{tj}}_{\beta_k^{tj}} \xi_k^{t\tau j} \tag{3.75}$$

$$\mu_k^{t\tau s} = \sum_{j=0}^{m_k} \underbrace{\frac{1}{\xi_k^{t\tau} r_k^t} \xi_k^{t\tau j} r_k^{tj} p_k^{tj}}_{\beta_k^{t\tau j}} \mu_k^{t\tau sj} \tag{3.76}$$

$$f_k^{t\tau s}(\mathbf{x}) = \sum_{j=0}^{m} \underbrace{\frac{\mu_k^{t\tau sj} \xi_k^{t\tau j} r_k^{tj} p_k^{tj}}{\mu_k^{t\tau s} \xi_k^{t\tau} r_k^t}}_{\beta^{t\tau sj}} f_{k_i}^{t\tau sj}(\mathbf{x}) \tag{3.77}$$

where

$$\beta_k^{tj} = \frac{r_k^{tj} p_k^{tj}}{r_k^t} = \begin{cases} \dfrac{p_k^{tj}}{r_k^t}, & j > 0 \\ \dfrac{r_k^0 p_k^{t0}}{r_k^t}, & j = 0 \end{cases} \tag{3.78}$$

$$\beta_k^{t\tau j} = \frac{\xi_k^{t\tau j} r_k^{tj} p_k^{tj}}{\xi_k^{t\tau} r_k^t} = \beta_k^{tj} \frac{\xi_k^{t\tau j}}{\xi_k^{t\tau}} \tag{3.79}$$

$$\beta_k^{t\tau sj} = \frac{\mu_k^{t\tau sj} \xi_k^{t\tau j} r_k^{tj} p_k^{tj}}{\mu_k^{t\tau s} \xi_k^{t\tau} r_k^t} = \beta_k^{t\tau j} \frac{\mu_k^{t\tau sj}}{\mu_k^{t\tau s}} \tag{3.80}$$

Using the individual $f_k^{t\tau sj}(\mathbf{x})$, the combined state $f_k^{t\tau s}(\mathbf{x})$ can be approximated by use of moment matching techniques.

Proof. The Multi-Bernoulli Mixture (MBM) containing the posterior track estimates, weights, and existence probabilities can be approximated as a Multi-Bernoulli. A thorough proof of this, and more context regarding the MBM, can be found in [138]. Drawing from the aforementioned proof, in combination with the proof in [24, Appendix D], we have that the posterior distribution over $\mathbf{y}$ can be approximated as

$$f_k^t(\mathbf{y}) \approx \sum_{j=1}^{m_k} \beta_k^{tj} f_k^{tj}(\mathbf{y}) \tag{3.81}$$

$$\tag{3.82}$$

where

$$\beta_k^{tj} = \frac{r_k^{tj} p_k^{tj}}{r_k^t} \tag{3.83}$$

and

$$f_k^{tj}(\mathbf{y}) = o_k^{\mathrm{tv}j} \xi_k^{t\tau j} \mu_k^{t\tau sj} f_k^{ts\tau j}(\mathbf{x}). \tag{3.84}$$

Using this, together with the approximation that the visibility is independent of the other states, we can write

$$\begin{aligned}
&\sum_{j=1}^{m_k} \beta_k^{tj} f_k^{tj}(\mathbf{y}) \approx \sum_{j=1}^{m_k} \beta_k^{tj} \xi_k^{t\tau j} \mu_k^{t\tau sj} f_k^{ts\tau j}(\mathbf{x}) \sum_{j=1}^{m_k} \beta_k^{tj} o_k^{\mathrm{tv}j} \\
=&\frac{\sum_{j=1}^{m_k} \beta_k^{tj} \xi_k^{t\tau j} \mu_k^{t\tau sj} f_k^{ts\tau j}(\mathbf{x})}{\sum_{j=1}^{m_k} \beta_k^{tj} \xi_k^{t\tau j} \mu_k^{t\tau sj}} \frac{\sum_{j=1}^{m_k} \beta_k^{tj} \xi_k^{t\tau j} \mu_k^{t\tau sj}}{\sum_{j=1}^{m_k} \beta_k^{tj} \xi_k^{t\tau j}} \sum_{j=1}^{m_k} \beta_k^{tj} \xi_k^{t\tau j} \sum_{j=1}^{m_k} \beta_k^{tj} o_k^{\mathrm{tv}j} \\
=&\sum_{j=1}^{m_k} \frac{\beta_k^{tj} \xi_k^{t\tau j} \mu_k^{t\tau sj}}{\sum_{j=1}^{m_k} \beta_k^{tj} \xi_k^{t\tau j} \mu_k^{t\tau sj}} f_k^{ts\tau j}(\mathbf{x}) \sum_{j=1}^{m_k} \frac{\beta_k^{tj} \xi_k^{t\tau j}}{\sum_{j=1}^{m_k} \beta_k^{tj} \xi_k^{t\tau j}} \mu_k^{t\tau sj} \sum_{j=1}^{m_k} \beta_k^{tj} \xi_k^{t\tau j} \sum_{j=1}^{m_k} \beta_k^{tj} o_k^{\mathrm{tv}j} \\
=&\underbrace{\sum_{j=1}^{m_k} \beta_k^{t\tau sj} f_k^{ts\tau j}(\mathbf{x})}_{f_k^{ts\tau}(\mathbf{x})} \underbrace{\sum_{j=1}^{m_k} \beta_k^{t\tau j} \mu_k^{t\tau sj}}_{\mu_k^{t\tau s}} \underbrace{\sum_{j=1}^{m_k} \beta_k^{tj} \xi_k^{t\tau j}}_{\xi_k^{t\tau}} \underbrace{\sum_{j=1}^{m_k} \beta_k^{tj} o_k^{\mathrm{tv}j}}_{o_k^{\mathrm{tv}}}
\end{aligned} \tag{3.85}$$

Keeping in mind that $r_k^{tj} = 1 \; \forall \; j > 0$ and that $o_k^{t1j} = \eta_k^{tj} = 1 \; \forall \; j > 0$ we get the wanted expressions. Lastly, we get the expression for the existence probability r_k^t directly from Chapter 2.

□

3.6 Target-provided measurement handling

The method shown in the previous section does not specify how the target-provided measurements are grouped before being sent to the tracker. In this section, we present three different ways of considering the target-provided measurements.

3.6.1 Method A: Sequential measurement processing

The first method for handling the incoming target-provided measurements is to process them, and perform the data association, as they arrive. This would mean that the predicting and updating of tracks is performed for each target-provided measurement, which can arrive at any time between radar measurement batches. This approach demands no further extensions to what is described above. The method is shown in Algorithm 1.

Algorithm 1 Method A: Sequential measurement processing

Require: target-provided measurements $Z_A = \{\mathbf{z}_A^1, \ldots, \mathbf{z}_A^m\}$, radar measurements $Z_R = \{\mathbf{z}_R^1, \ldots, \mathbf{z}_R^m\}$, tracks from previous time step $X = \{\mathbf{x}^1, \ldots, \mathbf{x}^n\}$
for target-provided measurement $\mathbf{z}_A^j \in Z_A$ **do**
$X \leftarrow$ PREDICT(X, t_A^j) ▷ predict tracks to time of $\mathbf{z}_A^j$
$X \leftarrow$ UPDATE(X, $\mathbf{z}_A^j$)
end for
$X \leftarrow$ PREDICT(X, t_R) ▷ predict tracks to time of Z_R
$X \leftarrow$ UPDATE(X, Z_R)

3.6.2 Method B: Precise batch measurement processing

The second method performs the data association for the target-provided measurements at the times when radar measurements arrive. The method considers all the target-provided measurements that have arrived between the previous and current time steps as a batch of measurements. This method is conceptually similar to what is done in [44] and [59]. The method is shown in Algorithm 2. The target-provided measurements with the same ID are clustered together, and the data association is performed based on these clusters. The clustering means that the measurement likelihood has to be calculated for each cluster rather than for each measurement. The measurement likelihood for I_m measurements with the same ID is

$$f_{\mathbf{z}}(\mathbf{z}|\mathbf{x}) = f_{\mathbf{z}}(\mathbf{z}^1, \ldots, \mathbf{z}^{I_m}|\mathbf{x}) = \prod_{i=1}^{I_m} f_{\mathbf{z}}(\mathbf{z}^i|\mathbf{z}^{i-1}, \ldots, \mathbf{z}^1, \mathbf{x}) \tag{3.86}$$

where

$$f_{\mathbf{z}}(\mathbf{z}^i|\mathbf{z}^{i-1}, \ldots, \mathbf{z}^1, \mathbf{x}) = \int f_{\mathbf{z}}(\mathbf{z}^i|\mathbf{x}^i) f_{\mathbf{x}}(\mathbf{x}^i|\mathbf{z}^{i-1}, \ldots, \mathbf{z}^1, \mathbf{x}) \mathrm{d}\mathbf{x}^i. \tag{3.87}$$

This has to be calculated for each measurement that has arrived between the radar updates. The measurements are sorted according to their time stamp, with $\mathbf{z}^{I_m}$ being

the most recent measurement. This expression effectively replaces the integral in (3.64). The individual kinematic states are calculated as

$$f_k^{t\tau sj}(\mathbf{x}|\mathbf{z}^i, \mathbf{z}^{i-1}, \ldots, \mathbf{z}^1, \mathbf{x}) = \frac{f_{\mathbf{z}}(\mathbf{z}^i|\mathbf{x}^i) f_{\mathbf{x}}(\mathbf{x}^i|\mathbf{z}^{i-1}, \ldots, \mathbf{z}^1, \mathbf{x})}{\int f_{\mathbf{z}}(\mathbf{z}^i|\mathbf{x}^i) f_{\mathbf{x}}(\mathbf{x}^i|\mathbf{z}^{i-1}, \ldots, \mathbf{z}^1, \mathbf{x}) \mathrm{d}\mathbf{x}^i} \tag{3.88}$$

This expression can be calculated using, for example, a Kalman filter. A thorough explanation of this recursive measurement likelihood calculation can be found in the supplementary material of [44]. With these expressions established, the other calculations and expressions are identical to Method A.

Algorithm 2 Method B: Precise batch measurement processing

Require: target-provided measurement clusters $Z_A = \{\mathbf{z}_A^1, \ldots, \mathbf{z}_A^m\}$, radar measurements $Z_R = \{\mathbf{z}_R^1, \ldots, \mathbf{z}_R^m\}$, tracks from previous time step $X = \{\mathbf{x}^1, \ldots, \mathbf{x}^n\}$

```
for track x^t ∈ X do
    for target-provided measurement cluster z_A^j ∈ Z_A do
        x^{t,j} ← COPY(x^t)
        for target-provided measurement z^i ∈ z_A^j do
            x^{t,j} ← PREDICT(x^t, t_A^{j,i})
            x^{t,j} ← UPDATE(x^j, t_A^{j,i})
        end for
        l^{t,j} ← MEASUREMENTLIKELIHOOD(x^{t,j}, z_A^j)
        x^{t,j} ← PREDICT(x^{t,j}, t_R)
    end for
    X_new^{t,j} ← x^{t,j}
end for
X ← MIXTUREREDUCTION(X_new, l)
X ← UPDATE(X, Z_R)
```

3.6.3 Method C: Batch measurement processing with added noise

In Section 3.3 it is assumed that the radar measurements of a single measurement batch are synchronized, that is, they all arrive at the same time. We do not make the same assumption for the target-provided measurements. However, making this assumption would allow us to simplify the handling of the measurements and remove some of the computational complexity of the above methods. Such an approach would be well suited when the radar frequency is high, as the timing errors would be small. Algorithm 3 describes the approach. Furthermore, only the most

recent measurement is considered when a target has transmitted more than one measurement between radar updates. In addition, this method should be used with a higher measurement noise level to account for the synchronization errors.

Algorithm 3 Method C: Batch measurement processing with added noise

Require: target-provided measurements $Z_A = \{\mathbf{z}_A^1, \ldots, \mathbf{z}_A^m\}$, radar measurements $Z_R = \{\mathbf{z}_R^1, \ldots, \mathbf{z}_R^m\}$, tracks from previous time step $X = \{\mathbf{x}^1, \ldots, \mathbf{x}^n\}$

$X \leftarrow$ PREDICT(X, t_R) ▷ predict tracks to time of Z_R, Z_A

$X \leftarrow$ UPDATE(X, Z_A)

$X \leftarrow$ UPDATE(X, Z_R)

Remark 3.3. When grouping the same-ID target-provided measurements, one has to keep in mind the assumption of only one measurement arising from each target. If a target transmits two target-provided measurements between radar updates, and one of the measurements has a corrupted ID number, this would breach the assumption. The most obvious way to amend this is to discard target-provided measurements whenever there are more measurements than tracks present. This will, however, interfere with initializing new tracks on the target-provided measurements. It should also be noted that if the radar frequency is higher than the target-provided measurement transmission frequency, a cluster will always only contain a single measurement. This would avoid the aforementioned problem, and simplify calculations. ◁

Remark 3.4. When using (3.86), the discrete hybrid states will take their most likely value as a mean over the information from the measurements in the cluster. This is as opposed to obtaining the most likely value at the most recent target-provided measurement. This could theoretically impact the estimation of the discrete states. For exampleif two measurements in a cluster indicate two different kinematic models this disparity will not be captured when using the batch processing methods. ◁

3.7 Implementation

3.7.1 Utilization of Gaussian-linearity

To make the implementation tractable, we model the individual kinematic states and the measurement likelihoods as Gaussian distributions. This allows us to use an Extended Kalman Filter when predicting and updating the kinematic estimates. The measurement likelihoods are defined as

$$f_{\mathbf{z}}^R(\mathbf{z}_k|\mathbf{y}_k) = \mathcal{N}(\mathbf{z}_k|\mathbf{H}_R\mathbf{x}, \mathbf{R}_R) \tag{3.89}$$

for radar measurements, and as

$$f_{\mathbf{p}}(\mathbf{p}_k|\mathbf{y}_k) = \mathcal{N}(\mathbf{p}_k|\mathbf{H}_A\mathbf{x}, \mathbf{R}_A), \tag{3.90}$$

for the positional part of the AIS measurements. Furthermore, the kinematic transition density $f_{\mathbf{x}}^{s\tau}(\mathbf{x}_k|\mathbf{x}_{k-1})$ is assumed to be in the form of a Gaussian

$$f_{\mathbf{x}}^{s\tau}(\mathbf{x}_k|\mathbf{x}_{k-1}) = \mathcal{N}(\mathbf{x}_k|\mathbf{f}^{(s)}(\mathbf{x}_{k-1}), \mathbf{Q}^{(s)}). \tag{3.91}$$

The transition model is linearized when needed to enable EKF prediction and Gaussian moment matching for mixture reduction.

The kinematic unknown target density from (3.2) is defined as

$$f_u(\mathbf{x}) = 1_\Omega(\mathbf{H}^{(s)}\mathbf{x})\mathcal{N}(\mathbf{H}^{*(s)}\mathbf{x}; \mathbf{0}, \mathbf{P}_v). \tag{3.92}$$

where $1_\Omega(\cdot)$ is an indicator function which is 1 when the unknown target is within the surveillance area, and $\mathbf{H}^{*(s)}$ is the permutation matrix corresponding to the non-positional states of the state vector $\mathbf{x}$. Using this we have that

$$f_{\mathbf{z}}(\mathbf{z}|\mathbf{x}, s, \tau)f_u(\mathbf{x}) \;=\; 1_\Omega(\mathbf{H}^{(s)}\mathbf{x})\mathcal{N}(\mathbf{z}_k^{a^t}|\mathbf{H}^{(s)}\mathbf{x}, \mathbf{R}^s)\mathcal{N}(\mathbf{H}^{*(s)}\mathbf{x}|\mathbf{0}, \mathbf{P}_v^{(s)}) \tag{3.93}$$

In the case of a large enough surveillance area Ω, and under the assumption of Gaussian-linearity, this can be approximated as $\mathcal{N}(\mathbf{x}|\hat{\mathbf{x}}_0^s, \mathbf{P}_0^s)$. Furthermore, this means that the constant c in Proposition 3.2 becomes

$$c = \int f_{\mathbf{z}}(\mathbf{z}|\mathbf{x}, s, \tau)f_u(\mathbf{x})\mathrm{d}\mathbf{x} \approx \int \mathcal{N}(\mathbf{x}|\hat{\mathbf{x}}_0^s, \mathbf{P}_0^s)\mathrm{d}\mathbf{x} = 1. \tag{3.94}$$

A more thorough proof regarding the unknown target density can be found in Appendix C of Chapter 2.

3.7.2 Gating

Because the target-provided measurements can arrive at any time, the number of times we have to perform gating increases considerably. The main computational cost of this is the number of predictions. Thus, we should consider this when creating the gating procedure.

Several different gating methods are presented in [136]. The first method relies on gating for each kinematic model, and it uses all measurements that have been gated by any of the models. A different method is a centralized gating procedure which makes an approximation across all models using a single gate. We use a somewhat more refined method, the Two-step Model Probability Weighted Gating (TS-MPWG) method. TS-MPWG was also presented in [136]. The first step in the method is a centralized gating procedure

$$f_{k|k-1}^{t}(\mathbf{x}) = \sum_{\tilde{\tau}} \xi_{k|k-1}^{t\tilde{\tau}} \sum_{\tilde{s}} \mu_{k|k-1}^{t\tilde{\tau}\tilde{s}} f_{k|k-1}^{t\tilde{\tau}\tilde{s}}(\mathbf{x}) \tag{3.95}$$

where $f_{k|k-1}^{t}(\mathbf{x}) = \mathcal{N}(\mathbf{x}|\hat{\mathbf{x}}_{k|k-1}, \hat{\mathbf{P}}_{k|k-1})$ provides the gate center $\hat{\mathbf{x}}_{k|k-1}$ and the predicted covariance $\hat{\mathbf{P}}_{k|k-1}$. Furthermore, the innovation covariance becomes

$$\mathbf{S} = \mathbf{H}\hat{\mathbf{P}}_{k|k-1}\mathbf{H}^{\top} + \mathbf{R}_k. \tag{3.96}$$

If no measurements are gated during the first step, the next step is initiated. Here, the gate is determined by the largest possible model error and should encompass any measurements generated by the target even if the chosen kinematic model is wrong. Thus, the TS-MPWG method can exploit the more computationally effective nature of the central gating method while compensating for eventual model errors. Adapting the expressions in [136] to this model, the gate in the second step is determined in by the maximal difference between $\hat{\mathbf{x}}_{k|k-1}$ and the individual $\hat{\mathbf{x}}_{k|k-1}^{t\tau s}$. This error is

$$K_{\max} = \arg\max_{\tau,s} \|\mathbf{H}\hat{\mathbf{x}}_{k|k-1} - \mathbf{H}\hat{\mathbf{x}}_{k|k-1}^{t\tau s}\|^2. \tag{3.97}$$

Using this, we calculate the gate volume as

$$\mathbf{S}_d = \mathbf{S} + \mathbf{K}_{\max} \tag{3.98}$$

where

$$\mathbf{K}_{\max} = \operatorname{diag}[\overbrace{K_{\max}, \ldots, K_{\max}}^{n}] \tag{3.99}$$

for a measurement space of dimension n.

Furthermore, it would be beneficial to have the possibility of gating target-provided measurements between two radar time steps without having to predict the state of all tracks. We can achieve this by utilizing one of the methods described in [142]. The method involves expanding the gate size according to a fixed presumed maximum velocity. That is, rather than predicting the track from time t_{k-1} to t_k, the gate accounts for movement in all directions at a very high speed. This method gives very large validation gates, and we only use it as a preliminary step before using the TS-MPWG method. Here, the radius of the gate is decided by

$$r_k = 2r_{k_0} + (\mathrm{t}_k - \mathrm{t}_{k-1})v_{\max} \tag{3.100}$$

where v_{max} is a parameter representing the largest possible speed for a target, and

$$r_{k_0} = \sqrt{g \times \mathrm{eig}(\mathbf{R})_{\max}}. \tag{3.101}$$

Here, g is the gate size, and $\mathrm{eig}(\mathbf{R})_{\max}$ is the largest eigenvalue of the measurement covariance matrix.

3.7.3 Initialization and termination

Due to target-provided measurements never being clutter measurements, care should be taken when choosing the initialization scheme. In JIPDA tracking algorithms, new tracks are usually only initialized on so-called free measurements, that is, measurements that have not been gated by any tracks at the current time step. When using this scheme, a target-provided measurement belonging to an uninitialized target, which falls within the validation gate of a previously initialized target, would most likely assign the measurement to the previously initialized target. However, a scheme that initiates tracks on all measurements will avoid this problem.

Initializing a new track on every measurement is computationally expensive and requires measures to mitigate computational complexity. For this purpose, we classify the tracks as newborn, adolescent, and ordinary. Newborn tracks are tracks that have been initialized at the current time step, adolescent tracks are tracks that were initialized at the previous time step, and ordinary tracks are all other tracks. The adolescent tracks are not allowed to compete for measurements in the same way as the ordinary tracks. The restriction comes into play when an adolescent track i and an ordinary track t have gated measurement j at the current time step, and they have both gated the same measurement at the previous time step. Then, the adolescent track i is only allowed to compete for the measurement if it has a larger weight relative to the measurement than the other track

$$\max_{t,j} w_k^{tj} < T_B w_k^{ij}. \tag{3.102}$$

where T_B is a threshold parameter. Otherwise, the adolescent track is not allowed to compete for measurement j, which is enforced by setting $w_k^{ij} = 0$.

Termination is done as described in [143]. First, any tracks with an existence probability under a predetermined threshold T_d are removed. Furthermore, any two tracks deemed to be identical are identified by the use of the hypothesis test in [7, p. 447]. The most recently initialized of these are then terminated. Lastly, any tracks that have not been associated with a measurement for N_T radar intervals are terminated.

3.7.4 Kinematic models

The implementation uses two different kinematic models: the Constant Velocity (CV) model and the Coordinated Turn (CT) model. Due to the varying prediction intervals, we use the discretized continuous formulation of the models. The CV model has the kinematic state $\mathbf{x} = [x, y, v_x, v_y]^\top$ where v denotes the velocity, and the state evolves according to $\mathbf{x}_k = \mathbf{F}^{(s)}(\mathrm{dt})\mathbf{x}_{k-1} + \mathbf{v}_k$, $\mathbf{v}_k \sim \mathcal{N}(\mathbf{0}, \mathbf{Q}^{(s)})$ where

$$\mathbf{F}^{(s)} = \begin{bmatrix} \mathbf{I}_2 & \mathrm{dt}\mathbf{I}_2 \\ \mathbf{0} & \mathbf{I}_2 \end{bmatrix}, \quad \mathbf{Q}^{(s)} = \begin{bmatrix} {}^{(\mathrm{dt})^3}/_{3}\mathbf{I}_2 & {}^{(\mathrm{dt})^2}/_{2}\mathbf{I}_2 \\ {}^{(\mathrm{dt})^2}/_{2}\mathbf{I}_2 & \mathrm{dt}\mathbf{I}_2 \end{bmatrix} q. \tag{3.103}$$

Here, $\mathbf{I}$ is the identity matrix, dt is the prediction interval, and q is the process noise intensity [8, p. 270] of the process noise. The CT model has an additional state ω, which is the turn rate. It evolves as $\mathbf{x}_k = \mathbf{F}^{(s)}(\mathbf{x}_{k-1})\mathbf{x}_{k-1} + \mathbf{v}_k$, $\mathbf{v}_k \sim \mathcal{N}(\mathbf{0}, \mathbf{Q}^{(s)})$ where

$$\mathbf{F}^{(s)}(\mathbf{x}) = \begin{bmatrix} 1 & 0 & \frac{\sin \mathrm{dt}\omega}{\omega} & \frac{-1+\cos \mathrm{dt}\omega}{\omega} & 0 \\ 0 & 1 & \frac{1-\cos \mathrm{dt}\omega}{\omega} & \frac{\sin \mathrm{dt}\omega}{\omega} & 0 \\ 0 & 0 & \cos \mathrm{dt}\omega & -\sin \mathrm{dt}\omega & 0 \\ 0 & 0 & \sin \mathrm{dt}\omega & \cos \mathrm{dt}\omega & 0 \\ 0 & 0 & 0 & 0 & 1 \end{bmatrix} \tag{3.104}$$

and

$$\mathbf{Q}^{(s)} = \begin{bmatrix} \mathbf{Q}^{(1)} & \mathbf{0} \\ \mathbf{0} & \mathrm{dt}q_\omega \end{bmatrix} \tag{3.105}$$

where $\mathbf{Q}^{(1)}$ is a CV model covariance matrix, and q_ω is the intensity of the turn rate process noise. In the implementation, the CT model is linearized as in [8, §11.7.2].

Remark 3.5. In most IMM applications, the transition matrix is constant. Thus, an aspect that has to be considered when the measurements do not arrive at a fixed frequency, is how to design the time-varying transition matrix $\pi(\mathrm{dt})$. A solution is to use the theory of Continuous Markov Chains to get an approximation for $\pi(\mathrm{dt})$ from the time-independent transition matrix π. As described in [58], this can be done by use of a generator matrix G. The generator matrix is closely related to the time-independent transition matrix π and is formulated as

a) no transition takes place in the time interval dt with probability $1+g_{ii}\mathrm{dt}+o(\mathrm{dt})$

b) a transition takes place in the time interval dt with probability $g_{ij}\mathrm{dt} + o(\mathrm{dt})$.

where g_{ij} are the individual elements of G, and $o(\mathrm{dt})$ indicates some small additional term which is ignored. This approximation is reasonable for relatively small dt. Thus, the generator matrix G for M number of states can be written as

$$G = \begin{bmatrix} \pi^{11} - 1 & \dots & \pi^{1M} \\ \vdots & \ddots & \vdots \\ \pi^{M1} & \dots & \pi^{MM} - 1 \end{bmatrix} \tag{3.106}$$

where π^{ij} are the individual elements of π. Furthermore we have from [58] that

$$\pi^{ij}(\mathrm{dt}) \approx g_{ij}\mathrm{dt} \text{ if } i \neq j \text{ and } \pi^{ii}(\mathrm{dt}) \approx 1 + g_{ii}\mathrm{dt}. \tag{3.107}$$

Using this we get

$$\pi(\mathrm{dt}) \approx \begin{bmatrix} 1 + (\pi^{11} - 1)\mathrm{dt} & \dots & \pi^{1M}\mathrm{dt} \\ \vdots & \ddots & \vdots \\ \pi^{M1}\mathrm{dt} & \dots & 1 + (\pi^{MM} - 1)\mathrm{dt} \end{bmatrix} \tag{3.108}$$

. ◁

3.7.5 Measurement models

Radar measurements

The radar measurements only contain positional data, and the measurements can be written as

$$\mathbf{z}_k = \mathbf{H}\mathbf{x}_k + \mathbf{w}_k, \;\; \mathbf{w}_k \sim \mathcal{N}(\mathbf{0}, \mathbf{R}_R) \tag{3.109}$$

The noise matrix has both a Cartesian and polar element, to account both for errors in range and bearing, and clustering errors. The measurement noise matrix for the radar measurement becomes

$$\mathbf{R}_R = \mathbf{R}_c + \mathbf{R}_p \tag{3.110}$$

Here $\mathbf{R}_c$ is the Cartesian noise component, while $\mathbf{R}_p$ is the polar noise component converted to Cartesian coordinates. The conversion is done by using the unbiased conversion equations from [95].

Target-provided measurements

The target-provided measurements can contain both positional and velocity data. The kinematic part of the measurements can be written as

$$\mathbf{p}_k = \mathbf{H}\mathbf{x}_k + \mathbf{H}^{\text{vel}}\mathbf{x}_k + \mathbf{w}_k, \;\; \mathbf{w}_k \sim \mathcal{N}(\mathbf{0}, \mathbf{R}_A) \tag{3.111}$$

where $\mathbf{H}$ and $\mathbf{H}^{\text{vel}}$ are the position and velocity measurement matrices, respectively. The position is usually derived from GPS information, while the velocity is derived either from a combination of speed and heading data [16]. Due to the nature of the data, we approximate the positional errors as Cartesian noise, while we approximate the velocity errors as polar noise. The measurement noise matrix for the AIS measurement becomes

$$\mathbf{R}_A = \mathbf{H}\mathbf{R}_{c,A} + \mathbf{H}^{\text{vel}}\mathbf{R}_{p,A} \tag{3.112}$$

where $\mathbf{R}_{c,A}$ is the Cartesian noise component, while $\mathbf{R}_{p,A}$ is the polar noise component converted to Cartesian coordinates, again by using [95].

3.8 Results

3.8.1 Simulation environment

We created the simulated data in line with the assumptions in Section 3.3. The ownship is situated at the origin and is stationary. The surveillance area is circular with a radius of 500m. We track five targets, all appearing at the edge of the area. Three of the targets appear at time $\text{t} = 0\ \text{s}$, while the last two appear at time $\text{t} = 10\ \text{s}$. The data consists of true target positions, radar, and AIS measurements. The movement of the targets follows a CV model with process noise intensity $q = 0.1^2\ \text{m}^2/\text{s}^3$, with occasional maneuvers according to a CT model. Furthermore, all targets are guided towards the center of the surveillance area until they are within 50m of it. The measurements are created according to the measurement models in Section 3.7.5.

The tracking parameters were tuned to achieve good performance on experimental data and are similar to the ones in Chapter 2. We list the parameters in Table 3.2. These are also the parameters used for creating the simulated data. The AIS measurement noise was also chosen according to the experimental data and would correspond to the measurements providing high location accuracy. Furthermore, in practical applications, the precision of the AIS location data can be dynamically adjusted according to a position accuracy flag in the AIS protocol [112].

To evaluate the results we used five different performance measures : the Optimal subpattern assignment (OSPA) metric [122], the track localization error (TLE), track fragmentation rate (TFR), track false alarm rate (TFAR) and track probability of detection (TPD). The last four evaluation methods are described in [109]. The OSPA metric provides an overall performance assessment, while the other measures provide information about specific aspects of the methods.

Table 3.2: Tracking system parameters

Quantity	Symbol	Unit	Value
Radar sample interval	T	[s]	2.5
Model 1 process noise intensity	$q_{a,1}$	[$\mathrm{m}^2/\mathrm{s}^3$]	0.1^2
Model 2 process noise intensity	$q_{a,2}$	[$\mathrm{m}^2/\mathrm{s}^3$]	1.5^2
Turn rate process noise intensity	q_ω	[$\mathrm{rad}^2/\mathrm{s}^3$]	0.02^2
Cartesian noise std. radar	σ_{c_R}	[m]	6.6
Cartesian noise std. AIS	σ_{c_A}	[m]	3.0
Polar range std.	σ_r	[m]	8.0
Polar bearing std.	σ_θ	[°]	1.0
Detection probability	P_D	[%]	92
Survival probability	P_S	[%]	99.9
Non-corrupted ID probability	P_C	[%]	99
Initial visibility probability	η_u	[%]	90
Visibility Markov probability	$\omega^{\tilde{\mathrm{vv}}}$	[%]	$\left[\begin{smallmatrix} 90 & 10 \\ 52 & 48 \end{smallmatrix}\right]$
Gate size	g	[−]	3.5
Clutter intensity	λ	[$1/\mathrm{m}^2$]	5×10^{-7}
Unknown target rate	U	[$1/\mathrm{m}^2$]	5×10^{-8}
Initial velocity std.	σ_v	[m/s]	10
Initial model probability	μ_u^s	[%]	$\left[\begin{smallmatrix} 80 & 10 & 10 \end{smallmatrix}\right]$
Unknown target no ID probability	ξ_u^0	[−]	0.5
Existence confirmation threshold	T_c	[%]	99.9
Existence termination threshold	T_d	[%]	1
IMM transition probability	$\pi^{\tilde{ss}}$	[%]	$\left[\begin{smallmatrix} 99 & .5 & .5 \\ .5 & 99 & .5 \\ .5 & .5 & 99 \end{smallmatrix}\right]$

We tested four different methods: the three methods described in Section 3.6, and a method using only the radar measurements.. The code implementing Method A from Section 3.6 is available at [69].

3.8.2 Simulated data

We tested the methods on 100 simulated data sets over a range of different detection probabilities. The results are seen in Figure 3.2, Figure 3.3, and Figure 3.4. Not surprisingly, the pure radar tracking method performs worse than the AIS-aided tracking methods from Section 3.6 when the P_D is low. The difference becomes smaller as P_D approaches 1, but is still significant.

Furthermore, we see that the three methods from Section 3.6 perform similarly. As expected, the batch processing method using added noise gives slightly less

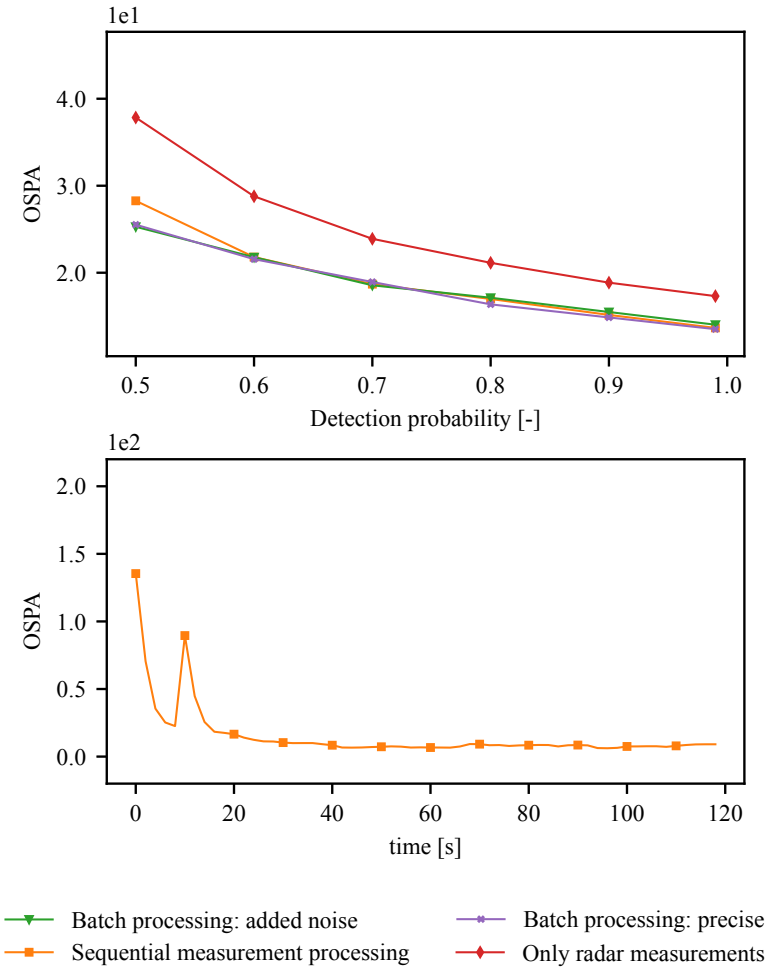

Figure 3.2: Comparison of the different methods, using the OSPA metric. The above figure shows the average OSPA values of each method for different detection probabilities. The below figure shows the average OSPA value for each time step, with $P_{\mathrm{D}} = 0.9$. Here, we only consider the sequential measurement processing method. Both figures contain results from the same 100 scenarios. The OSPA values are calculated using $p = 2$ and $c = 200$. The purpose of the two parameters is described in [122].

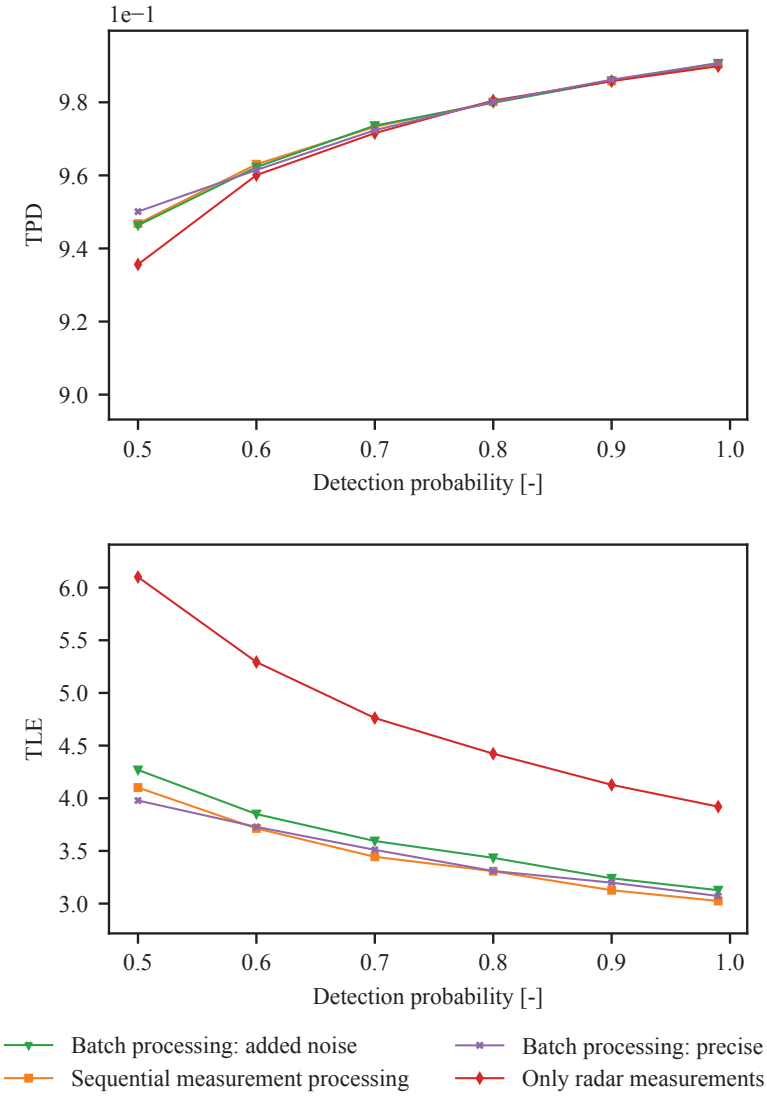

Figure 3.3: TPD and TLE for the four different methods, for different detection probabilities. The values were calculated by running the methods on the same 100 scenarios as above.

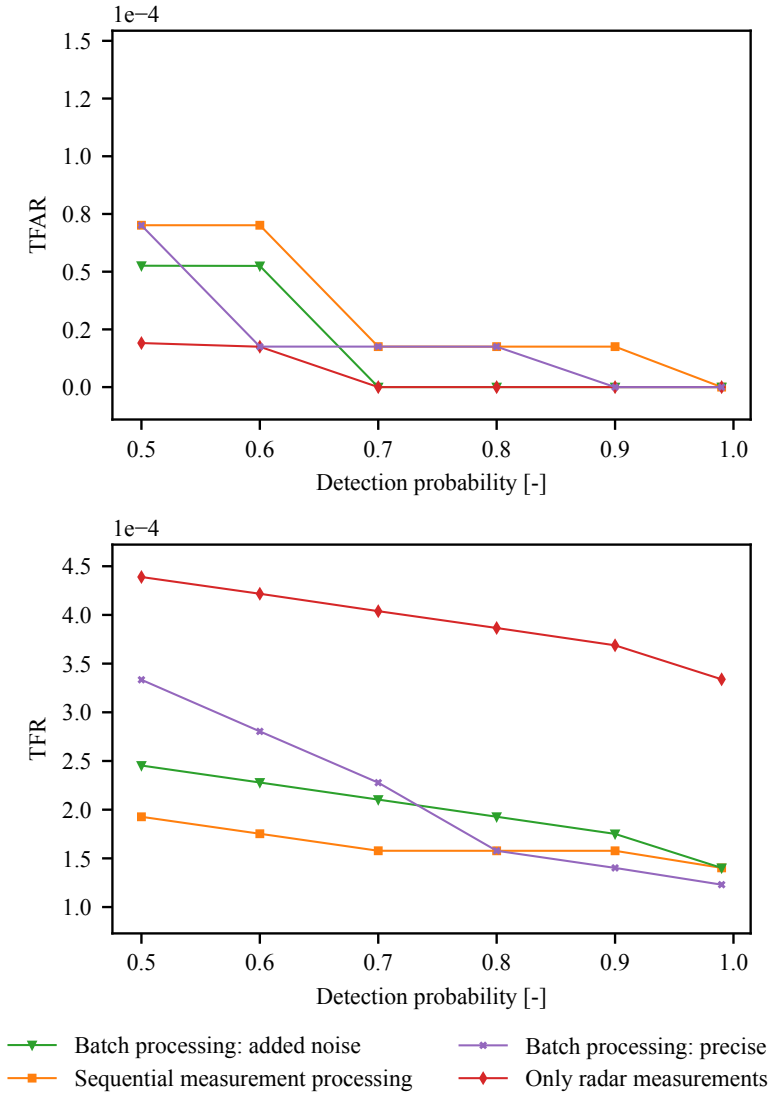

Figure 3.4: TFAR and TFR for the four different methods, for different detection probabilities. The values were calculated by running the methods on the same 100 scenarios as above.

precise estimates. While we see some differences between the methods for TFR and TFAR, the errors are of an overall small magnitude. However, the pure radar tracker is more prone to track fragmentation than the other methods.

The computational complexity of the methods also warrants a comparison. The pure radar tracker is the least computationally demanding, as all the other methods add functionality in addition to performing the calculations of the pure radar tracker. The precise batch processing method is the most demanding of the target-provided measurement handling methods. This is because it requires predictions and updates of each track for each measurement. The least demanding of the three is the batch processing method with added noise, as it does not need to perform more predictions than the pure radar methods. Nevertheless, the three methods generally do not introduce a prohibiting amount of complexity and can all be implemented using a Kalman filter.

3.8.3 Experimental data

In addition to the simulated data, the sequential measurement handling method and the pure radar tracker were tested on experimental data collected as part of the Autosea project at NTNU [26]. The data set is the same set used in Chapter 2. We consider two scenarios, which include three different ships using AIS, of which two provide frequent measurements. The transmission frequency for the two ships is higher than what is mandated by IMO [112], but the data set is nevertheless helpful for demonstrating the functionality and usefulness of the tracking method. Due to the AIS data previously being used as ground truth for the AIS-equipped vessels, the AIS data has been interpolated to increase the number of measurements. This interpolation was undone prior to using the data, that is, we removed any artificially added measurements.

Figure 3.5 shows the results from the first scenario. The scenario contains three fast-moving and maneuvering targets and a single slow-moving target. The slow-moving target is a large vessel with an AIS transmitter, while the three fast-moving targets are small, rigid inflatable boats (RIBs). Only one of the RIBs has an AIS transmitter, and it only transmits a single AIS measurement. The large vessel, however, provides high-quality AIS measurements. As can be seen, both the sequential measurement handling method and the pure radar method can track the scenario well. The two methods have flexibility in their use of IMM, and they can thus use different kinematic models for the RIBs and the large ship. When combining target-provided measurements with IMM the tracker is also better able to select the correct kinematic model for each target. Furthermore, the sequential measurement handling method can use the AIS measurements when tracking the large vessel, improving upon the track from the pure radar method. It also correctly associates the single AIS measurement transmitted by the RIB.

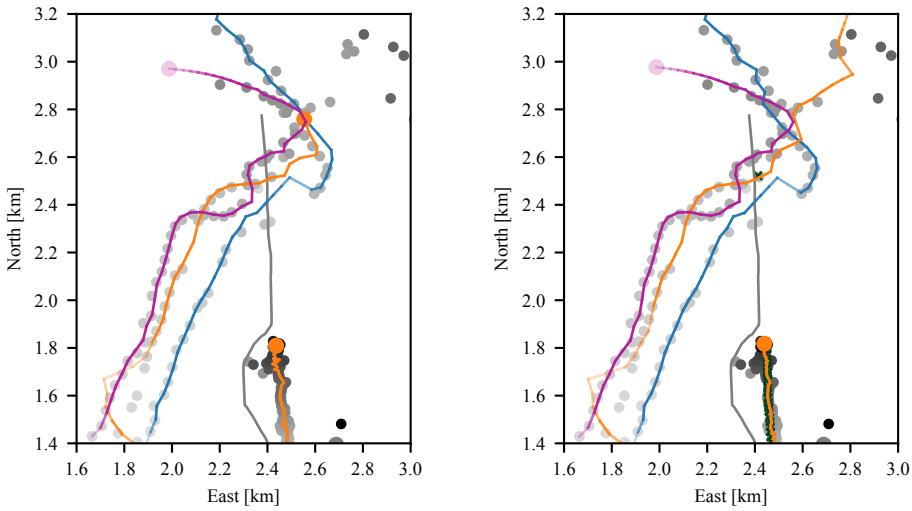

(a) Results when tracking the scenario using only radar.

(b) Results when tracking the scenario using Method A: Sequential measurement processing.

Figure 3.5: A scenario showing four targets. The ownship is the gray line, moving southwards, while the targets all move northwards. The gray dots are radar measurements, and the green crosses are AIS measurements. The measurements become more transparent as time passes, that is, the darker ones have arrived closer to the end of the scenario. The transparency of the tracks is decided by the existence probability, with the more transparent having a lower probability of existence. The target originating furthest to the right is a large vessel with an AIS transmitter, while the three other targets are small, fast-moving rigid inflatable boats (RIBs). Of the RIBs, only the orange has an AIS transmitter, which transmits a single measurement during the scenario. The RIBs make several maneuvers before moving beyond the radar range.

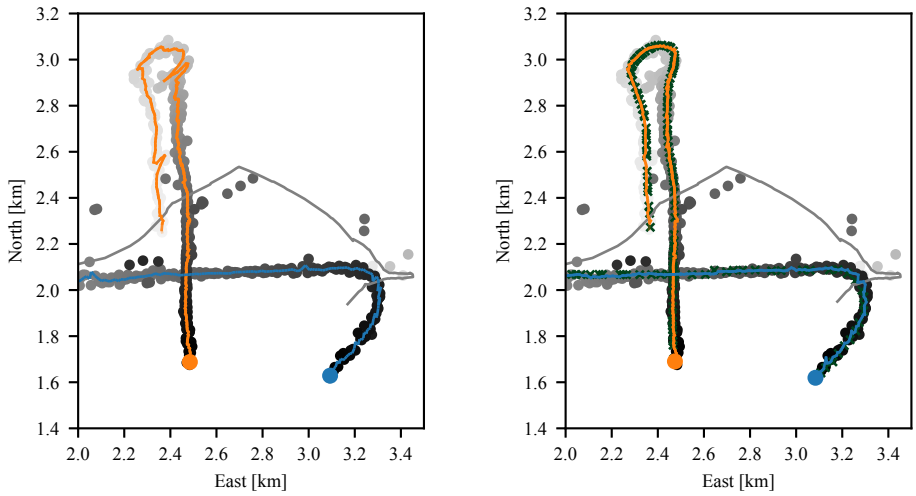

(a) Results when tracking the scenario using only radar.

(b) Results when tracking the scenario using Method A: Sequential measurement processing.

Figure 3.6: A scenario showing two large vessels with AIS transmitters (with tracks shown as blue and orange lines), in addition to an ownship (gray line). We depict the measurements and tracks as in Figure 3.5. Initially, the orange target moves north, while the blue target moves east. After some time, the orange target makes a u-turn, while the blue target makes a turn towards south-west. The ownship moves in a clockwise motion. The orange and blue dots represent the track positions at the end of the scenario.

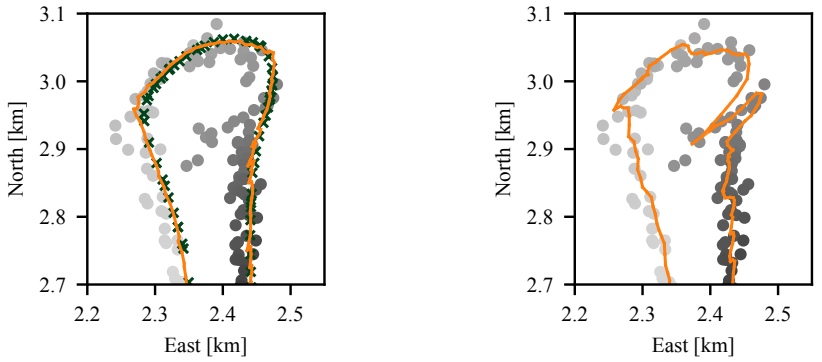

(a) A target making a clockwise turn, while being tracked using AIS and radar.

(b) A target making a clockwise turn, while being tracked using only radar.

Figure 3.7: A closer look at the northernmost turn for the orange track in the scenario in Figure 3.6. A single large vessel makes a clockwise turn, resulting in significant amounts of radar clutter.

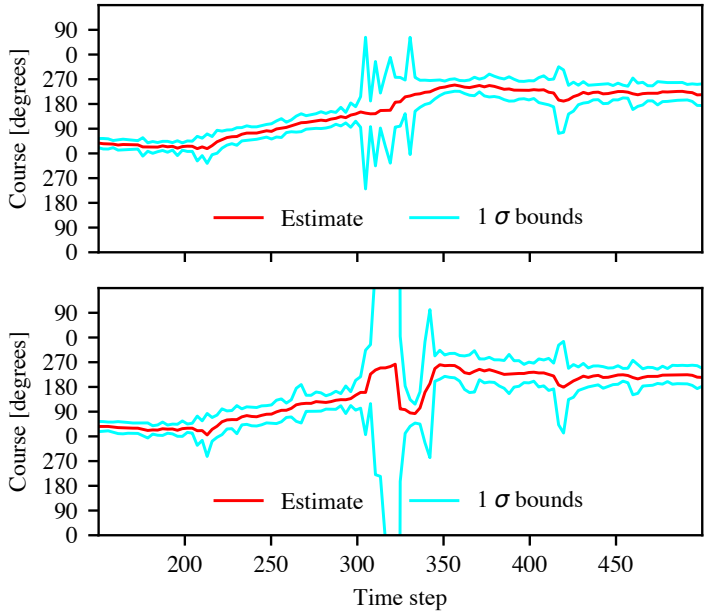

Figure 3.8: Course estimate for the turn depicted in Figure 3.7 using both radar and AIS (top) and using only radar (bottom).

The second scenario can be seen in Figure 3.6. The plots show the two vessels with frequent AIS transmissions and the ownship. Figure 3.7 displays a close-up of the northernmost turn, with and without AIS measurements. The second scenario highlights some advantages of utilizing the AIS measurements when available. The main event occurs during the turn depicted in Figure 3.7, where the radar measurements are poor due to the large vessel making a maneuver and generating numerous clutter measurements. A similar effect also occurs on the straight leading up to the turn. Both these effects cause the purely radar-guided tracking method to veer off track, while the sequential measurement handling method can utilize the AIS measurements to avoid this. Figure 3.8 shows the estimated course of the target during the turn, in addition to the standard deviation of the estimates. The poor radar measurements make the course estimates unreliable when not also utilizing the AIS measurements. When using the AIS measurements standard deviation of the course estimates during the turn is significant, but they are still considerably smaller than when the tracker uses only radar measurements. Furthermore, the track avoids sudden course changes. In this scenario, the inclusion of AIS measurements causes no unwanted consequences, opening the possibility of utilizing all the potential enhancements information given by the messages can bring.

3.8.4 Comparison with the particle based method of Gaglione et al.

In [43], a thorough comparison between the method from [44] and the sequential measurement processing method in Algorithm 1 is presented. The method from [44] uses a particle filter and loopy belief propagation and is implementation-wise very different from the one described in this chapter. We denote the method from [44] as the Belief propagation, particle filter (BP-PF) method.

Figure 3.9 show the mean generalized OSPA (GOSPA) [113] values for the VIMMJIPDA with and without AIS input, and the BP-PF for different detection probabilities. The results are based on the scenario described in Section 3.8.1, across 100 simulations. The BP-PF implementation uses a single CV model with process noise intensity $q = 0.8^2 \ \mathrm{m}^2/\mathrm{s}^3$, and the same parameters as in Table 3.2 where applicable. The GOSPA metric is a generalization of the OSPA metric and penalizes localization errors, missed detections, and false alarms. The additional penalties for missed detections and false alarms explains the differences in the values between Figure 3.9 and Figure 3.2. The results show that the two methods that use AIS input performs similarily, and that both outperform the pure radar tracker. This is, as in Figure 3.2, most evident when the detection probability is low, as this is where the proportional advantage of using AIS together with radar, in comparison to using only radar, is largest.

In Figure 3.10, a comparison is made between the BP-PF method and the VIMMJIPDA with AIS on a different data set. The scenario in question is the

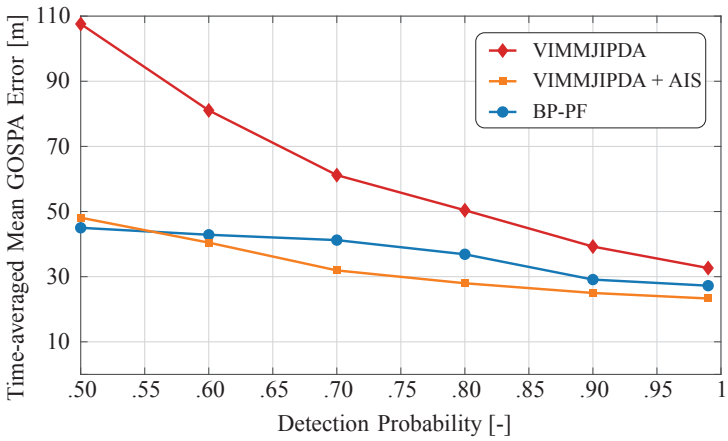

Figure 3.9: Comparison of the different methods for the same scenario as in Figure 3.2 across 100 simulations, using the GOSPA metric. From [43].

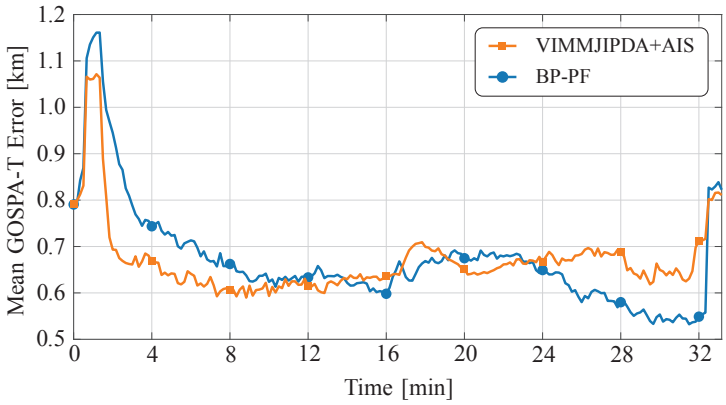

Figure 3.10: Performance of the two trackers during the scenario from [44] across 100 simulations, calculated using the trajectory GOSPA metric. From [43].

one used in [44], which involves eight targets. The targets start their trajectories 4 kilometers from the center of the area, before moving towards and eventually meeting at the center. When there, all targets turn 60 degrees before continuing for 4 additional kilometers. The radar used to track the targets is placed 15 kilometers from the area center, which results is large amounts of measurement noise. Both trackers use a single CV model with $q = 0.15^2 \ \mathrm{m}^2/\mathrm{s}^3$, and model the measurement noise with a range standard deviation of $250 \ \mathrm{m}$ and bearing standard deviation of $2.56°$. Furthermore, the clutter density is $1.7 \times 10^{-9} \ 1/\mathrm{m}^2$ and the unknown target density for the VIMMJIPDA with AIS is set to $10^{-10} \ 1/\mathrm{m}^2$. The trajectory-GOSPA metric [46] is used to evaluate the results. It differs from the regular GOSPA-metric in that it also penalizes track switches, and calculates the error based on trajectories rather than single target estimates. It is computed as an average of 100 simulation runs, and plotted for each time step in the scenario. The results show that the methods perform similarly overall, with the BP-PF method performing slightly worse in the early stages of the scenario, and slightly better in the later stages. The behavior of the VIMMJIPDA with AIS can be explained by the difficult situation at the mid point of the scenario, where the targets are closely spaced. There, it may struggle to keep track of the targets, leading to termination of tracks, and subsequently a struggle to reinitialze them. The BP-PF method, with its use of belief propagation, may be better able to compute correct association probabilities when the number of possible associations are very large. A comparison can be drawn to the results in [138], and the performance of the Track-oriented marginal MeMBer/Poisson filter. The filter is very similar to the VIMMJIPDA, albeit without multiple models and visibility modeling, and exhibits similar behavior on a scenario that also includes closely spaced targets at its mid point.

Lastly, the BP-PF method is tested on the same experimental data as in Figure 3.5, see Figure 3.11. The results are similar to that of the AIS capable tracker in Figure 3.5b, but the use of only a single kinematic model with high process noise results in some additional noise in the estimates for Gunnerus. This is consistent with the analysis made regarding the use of a single model in the VIMMJIPDA, see Section 2.5.2. Furthermore, a false alarm is created in the upper right corner of the area.

3.9 Chapter summary

In this chapter, we have presented a framework for including target-provided measurements in a JIPDA-based tracking algorithm. We use AIS measurements as an example of such measurements. It is seen that the inclusion of such measurements can help a pure radar tracking method and improve performance greatly when the radar measurements are of low quality. In addition to the pure performance

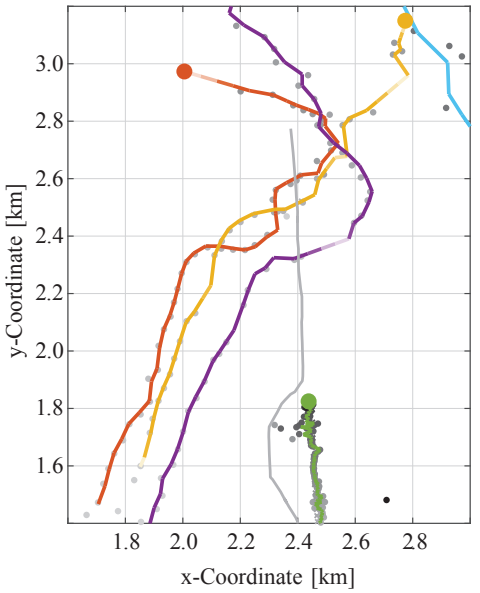

Figure 3.11: The same scenario as in Figure 3.5 when using the BP-PF method, from [43].

improvements, target-provided measurements can facilitate the identification of targets, which can be useful for, e.g., a collision avoidance system. Furthermore, we have presented and compared three different methods of handling the target-provided measurements: One method where the tracker processes the target-provided measurements when they arrive and two methods where the tracker processes them at the time of the radar update.

Acknowledgment

The radar data were recorded by Erik Wilthil, Andreas Flåten, Bjørn-Olav Eriksen and Giorgio D. K. M. Kufoalor, with assistance from Maritime Robotics and Kongsberg.

Chapter 4

Poisson Multi-Bernoulli Mixture Filtering With Fusion of Target-Provided and Exteroceptive Measurements

This chapter is based on the publication:

[65] 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. To appear in *IEEE Access*. doi: 10.1109/ACCESS.2024.3389824

Changes from the submitted publication involve shortening and editing of the introduction and background sections.

4.1 Introduction

In the last decades, many new multi-target tracking methods have been developed. The one used in this chapter is the Poisson multi-Bernoulli mixture (PMBM) filter [138]. The PMBM filter is based on Random Finite Sets (RFS), specifically a potentially detected target is modelled as a Bernoulli RFS and the set of undetected targets is modelled as a Poisson point process (PPP). The PMBM filter has been shown to be among the state-of-the-art methods in target tracking [22, 57, 138], and has also been used for extended object tracking [6, 52, 145]. Furthermore, the interacting multiple model (IMM) filter [15] can be used to provide more flexibility in the modeling of the target behavior. The IMM filter uses a set of motion models,

and switches between these depending on the current behavior of the target. The method has been used together with many tracking methods, including the PMBM filter [91].

This chapter combines many of the recent innovations regarding the PMBM filter and the use of target-provided information in target tracking. We present a method for including such information in the IMM-PMBM filter, using the model assumptions from Chapter 3. Furthermore, we show how to use target-provided information together with an extended object tracking method, a PMBM with a Gaussian process extent model (GP-PMBM) [6]. The extended object tracker and the point measurement tracker are evaluated using both simulated and experimental data. Additionally, the point measurement tracker is used in a closed-loop collision avoidance experiment that demonstrates its feasibility for use in real-world applications, which is further detailed in Chapter 5.

The chapter is organized as follows. Section 4.2 provides a brief explanation of relevant target tracking methods and concepts, together with an overview of previous work on use of target-provided data. In Section 4.3 we present the general model used to describe the targets and measurements. A general PMBM filter for use with target-provided measurements is described in Section 4.4, and it is specified for the point target and the extended object cases in Section 4.5. The test setups and results for the point target case are presented in Section 4.6, and for the extended object case in Section 4.7. Lastly, we conclude the chapter in Section 4.8.

4.2 Background

For the Probabilistic data association (PDA) filter and its derived methods, each estimate is first found as a mixture of the estimate conditioned on different measurements. The mixture is then combined into a single estimate. If the estimates are Gaussian, this is done by use of moment matching. The mixture reduction avoids most problems with computational complexity, as the number of estimates will be relatively limited. However, temporal information is lost in the process, and as such the approach is called single-scan target tracking because the mixture reduction discards information from previous time steps.

4.2.1 Multi-scan target tracking

The PMBM filter is a multi-scan method. In broad terms, what distinguishes multi-scan methods from single-scan methods is that the latter performs a mixture reduction of the estimates after each time step, whereas the former does not. Instead, each measurement-conditioned estimate is propagated to the next time step. Multiple hypothesis tracking (MHT) [115] is an early and widely used multi-scan method.

MHT constructs a tree structure of possible associations between the measurements and the estimates. Each time new measurements are received, the tree is updated by adding new branches representing new track hypotheses, with a path from the root to a leaf representing a single track hypothesis. MHT weights the single measurement-conditioned estimates according to the likelihood of the association between the track and the measurement, which allows the method to find the most likely global hypothesis. A global hypothesis is a valid combination of all possible associations between tracks and measurements. The PMBM filter works in a conceptually similar way, but it initializes the new potential targets from the PPP representing the undetected targets, whereas the different potential target hypotheses are represented as a multi-Bernoulli mixture. This structure ensures conjugacy in the Bayesian recursion and allows for more mathematically rigorous procedures when initializing the new potential tracks. Both MHT and the PMBM filter produce an exponentially increasing number of hypotheses. They mitigate this by removing the most unlikely hypotheses, a process called pruning. The exponentially growing tree of hypotheses does, however, demand more care in managing computational complexity.

4.2.2 Extended object tracking

Most early works on multi-target tracking assumes that objects only generate a single measurement, the so-called point target assumption. Relaxing this assumption to allow targets to generate a varying number of measurements leads to the problem of extended object tracking [51]. The manner in which an extended object generates measurements is most commonly modelled as an inhomogeneous PPP. The Poisson rate governs the expected number of measurements, and a specific spatial distribution indicates how these measurements are spatially distributed across the target [49]. This spatial distribution allows us to estimate the shape and size of an object, which is referred to as object extent. Initial approaches used elliptical shapes as priors for the spatial distribution, this is commonly referred to as the random matrix model [83]. The extent of the object is then modelled by a symmetric and positive definite $d \times d$ matrix called the shape matrix, where d is the dimension of the object. The elements of this matrix are then estimated according to the spatial distribution of the measurements. This model is very popular because it is a linear model. Another approach instead models the extent as a generic star-convex shape by parametrizing the shape contour. This enables the modelling of more complex shapes, but it should be noted that the estimation problem then becomes non-linear [11]. This method is well suited to modeling contour generated measurements, such as those generated by LiDARs. The shape contour is parametrized by a radius function which can be estimated by a variety of techniques, the seminal paper used Fourier series, but today the most common method is to use Gaussian processes [134]. The Gaussian process method also allows the use of specific symmetry properties of the tracked objects

when estimating their extent.

4.2.3 Multiple extended object tracking

When tracking multiple extended objects, the data association becomes harder because each target can generate an unknown number of measurements. The first theoretical framework for a multiple extended object tracker was derived from the probability hypothesis density (PHD) filter by using the above-mentioned PPP model [99]. This filter was then implemented using the random matrix model, and an inverse Wishart distribution was used for estimation of the shape matrix [54]. This approach was augmented by estimating the Poisson rate governing the expected number of generated measurements for each target using a gamma distribution [53]. Combining them resulted in the gamma Gaussian inverse Wishart (GGIW) model [96]. Later developments explored extended object formulations of other types of filters, such as the PMBM filter [52]. Initially it was presented with the GGIW model, but the same formulation has been used to implement a PMBM filter using the Gaussian process model augmented with a gamma distribution [6]. This filter has also been demonstrated on maritime LiDAR data.

4.2.4 Relation to previous work

Much of the mathematical framework used in this chapter is based on the work in Chapter 3. This is possible because the tracker in Chapter 3, even though it ultimately is a JIPDA-type tracker, can be seen as a special case of the PMBM filter. It is an extension of the IMM-JIPDA tracker with visibility modeling presented in Chapter 2, where the steps needed to go from the PMBM filter to a JIPDA-type filter are thoroughly presented. The link between the two filters can be explained by using the track-oriented marginal multiple target multi-Bernoulli/Poisson (TOMB/P) filter as an intermediate step. Williams notes in [138, Sec. IV-A] that the TOMB/P filter results from forcing the individual track hypotheses in a global PMBM hypothesis to be independent. This approximation results in tracks formed by the marginal track-to-measurement association probabilities, as is done in the JPDA and JIPDA filters. By assuming that new targets are born according to a stationary birth density, and by neglecting the influence of unknown targets when calculating the association probabilities, the TOMB/P filter becomes identical to a JIPDA filter. Despite these differences, the formulations in Chapter 3 and Chapter 2 are similar enough to make the transition from an IMM-JIPDA filter which utilizes target-provided information to an IMM-PMBM filter relatively straightforward. Note that a visibility state is present in both Chapter 2 and Chapter 3, which models the possibility of a target being occluded. This state is omitted in this chapter but is possible to include without much effort.

4.3 Model

The model used to describe the targets and measurements is similar to the one presented in Chapter 3, which itself is based in the model underlying the PMBM filter. The model describes how the targets are represented, how they are created, and how they evolve over time. Furthermore, it describes how the measurements and the information they contain relate to the targets. Thus, the model forms the framework we use to later describe how to estimate the target states based on the measurements.

4.3.1 The Poisson multi-Bernoulli mixture

In general, the PMBM filter models the targets as the union of undetected targets and detected targets. The undetected targets are represented as a PPP, whereas the hypothesized detected targets are modeled as a multi-Bernoulli mixture (MBM). The combination of a PPP and an MBM ensures a conjugate prior in the context of recursive Bayesian estimation. We write the multi-target density as

$$f(X) = \sum_{Y \uplus W = X} f^{ppp}(Y) f^{mbm}(W) \tag{4.1}$$

where X is the set of all targets, Y is the set of undetected targets, W is the set of detected targets, $\uplus$ denotes the disjoint union, $f^{ppp}(\cdot)$ is the PPP, and $f^{mbm}(\cdot)$ is the MBM. The PPP is defined as

$$f^{ppp}(X) = \exp(-\int \mu(\tilde{\mathbf{x}}) \mathrm{d}\tilde{\mathbf{x}}) \prod_{\mathbf{x} \in X} \mu(\mathbf{x}) \tag{4.2}$$

where $\mu(\tilde{\mathbf{x}})$ is the intensity function and the notation $\tilde{\cdot}$ indicates marginalization over a variable. Boldface notation is used for vectors. Furthermore, the MBM is defined as

$$f^{mbm}(X) \propto \sum_{j} \sum_{X_1 \uplus \ldots \uplus X_n = X} \prod_{i=1}^{n} w^{j,i} f^{j,i}(X_i). \tag{4.3}$$

The first sum accounts for all global hypotheses, and the second for all hypothesized targets within the global hypothesis. $w^{j,i}$ is the weight and distribution of potentially detected target i in global hypothesis j. The distribution $f^{j,i}(X_i)$, a Bernoulli RFS, is defined as

$$f^{j,i}(X^i) = \begin{cases} 1 - r^{j,i} & \text{if } X^i = \emptyset \\ r^{j,i} f^{j,i}(\mathbf{x}) & \text{if } X^i = \{\mathbf{x}\} \\ 0 & \text{otherwise} \end{cases} \tag{4.4}$$

where $r^{j,i}$ is the existence probability, and $f^{j,i}(\mathbf{x})$ is the state density.

4.3.2 The hybrid state space

The full state $\mathbf{y}$ of a target can contain both discrete and continuous states. Such a combination is often denoted as a hybrid state [8, p. 411]. The continuous part of the state is denoted as $\mathbf{x}$, and typically contains information such as position, velocity, or target extent. The discrete states can contain information such as target ID, what kinematic model the target is following, or if the target is occluded. In this section and the next, the exact information contained in the states is not important, but rather how they relate to each other, and how they evolve.

Nevertheless, we select two discrete states to illustrate the concepts, which are also useful later, namely the ID τ and the kinematic model s. The distribution of the hybrid state can be written as

$$f(\mathbf{x}, \tau, s) = f(\mathbf{x}|\tau, s)f(s|\tau)f(\tau). \tag{4.5}$$

Furthermore, the actual ID of a target is assumed to not change over time, whereas the kinematic model a target moves according to can change. For any discrete state with the same properties as either the ID or the kinematic model, the following can also be used to describe how to incorporate that state into the model.

4.3.3 New targets

We assume that new targets are born according to a PPP with intensity $b(\mathbf{y})$. A birth intensity with N_b components is defined as

$$b(\mathbf{y}) = \sum_{i=1}^{N_b} w^{b,i} f^{b,i}(\tau) f^{b,i}(s|\tau) f^{b,i}(\mathbf{x}|s, \tau). \tag{4.6}$$

Here, $f^{b,i}(\tau)$ is the distribution of the IDs τ, $f^{b,i}(s|\tau)$ is the distribution of the kinematic models s, and $f^{b,i}(\mathbf{x}|s, \tau)$ is the distribution of the state $\mathbf{x}$. As the hybrid state includes both discrete and continuous states, the PPP is what is denoted as a marked PPP, equivalent to a PPP on the Cartesian product of the continuous space of the kinematic density and the discrete spaces of the other states. Under the assumption that the distributions of all states are independent of other elements in the marked PPP, it nevertheless inherits the properties of a PPP, and can be used interchangeably in the PMBM. This means that a restriction of unique IDs across targets can not be enforced, as it would break the independence assumption underlying the PPP.

We assume that the IDs of unknown targets are independent and identically distributed (i.i.d.) with an initial prior distribution

$$f^{b,i}(\tau) = \begin{cases} \xi^{b,0} & \text{if } \tau = 0 \\ \dfrac{1 - \xi^{b,0}}{|\mathcal{V}| - 1} & \text{if } \tau > 0 \end{cases}. \tag{4.7}$$

Here, $\xi^{b,0}$ denotes the initial probability of a specific ID, whereas $|\mathcal{V}|$ is the number of possible IDs and $\tau \in \mathcal{V}$. The first case accounts for the probability of a target not having an ID, whereas the second case accounts for the probability of a target having an ID. For the latter case, the probability is uniformly distributed among all possible IDs. Furthermore, the kinematic models s are assumed to be i.i.d. with an initial distribution μ_s^0.

4.3.4 Target evolution

For the different variables, we use the subscript k to indicate that we are considering their value at the current time step, and with $k-1$ we indicate the previous time step.

Each target is assumed to survive a duration of T with probability

$$P_S(T) = P_{S_c}^{\mathrm{dt}} \tag{4.8}$$

where S_c is the survival probability per time unit, usually seconds.

The target IDs do not change over time, whereas the kinematic model can change between time steps as part of a Markov chain. The transition matrix π contains the Markov chain probabilities of a change occurring, whereas the target ID transition is modeled as a Kronecker delta δ. The kinematic transition density depends on the kinematic model, and we can write the prediction of a target as

$$f_{\mathbf{y}}(\mathbf{y}_k|\mathbf{y}_{k-1}) = f_{\mathbf{x}}(\mathbf{x}_k|\mathbf{x}_{k-1}, \tau_k, s_k)\pi^{s_{k-1}s_k}\delta^{\tau_{k-1}\tau_k}. \tag{4.9}$$

This transition is that of a target with ID τ_{k-1} and kinematic model s_{k-1} at the previous time step, and target ID τ_k and kinematic model s_k at the current time step.

4.3.5 Measurements

The detection probability of a target by way of exteroceptive sensors is assumed constant in both time and space and is denoted as $P_D(\mathbf{y}) = P_D$. For the target-provided measurements, the detection probability is designed to more precisely reflect their physical reality. This is done by setting the detection probability to 1 if a target-provided measurement is received, and 0 otherwise. We express this as

$$P_D(\mathbf{y}) = \begin{cases} 1 & \text{if a measurement is received} \\ 0 & \text{otherwise} \end{cases}. \tag{4.10}$$

Only the exteroceptive sensors are assumed to provide false alarms, referred to as clutter, and the clutter is modeled as a PPP with intensity $\lambda(Z)$. The intensity can depend on measurement position but is assumed to not change over time. For the target-provided measurements, the absence of clutter is modeled by using a PPP with intensity 0.

We assume that the exteroceptive measurements are synchronized, and all detections in each individual scan come from time step k. This is, for most sensors, an approximation. If needed, the time disparity between detections in a scan can be accounted for by relatively simple means Chapter 7. The likelihood for a set of exteroceptive measurements is $f_{\mathbf{z}}^{ex}(Z_k|\mathbf{y}_k)$, where Z_k is a set of measurements. For point target tracking, the set is either empty, or contains a single measurement. For extended object tracking, the set can contain several measurements. The information provided by the exteroceptive measurements is assumed to only contain the position of the detections, or potentially also the speed of the reflecting surface.

The target-provided measurements are not assumed to be synchronized. They can arrive at any time, and furthermore, they arrive at different times for different targets. The likelihood is denoted as $f_{\mathbf{z}}^{tp}(Z_k|\mathbf{y}_k)$. We assume that a set of target-provided measurements will contain at most a single measurement. As opposed to the case for the exteroceptive measurements $\mathbf{z}_k$ can now also contain additional information, such as ID and target dimensions. We assume that, when conditioned on the target state, the information in the measurements is independent of other information contained in the measurement. Furthermore, we keep in mind that the detection probability is 1 if a measurement is received. This means that the likelihood of a set with a measurement containing, for example, kinematic information $\mathbf{p}$, ID τ, length z^L, and width z^W can be decomposed as

$$f_{\mathbf{z}}^{tp}(Z_k|\mathbf{y}_k) = f_{\mathbf{p}}^{tp}(\mathbf{p}_k|\mathbf{y}_k) f_{\tau}^{tp}(\tau_k|\mathbf{y}_k) f_{L}^{tp}(z^L|\mathbf{y}_k) f_{W}^{tp}(z^W|\mathbf{y}_k) \tag{4.11}$$

whereas a set without a measurement has a likelihood of 1 because the detection probability is 0 when no measurement is received.

4.4 Method

From the previous time step $k-1$, we assume that the Poisson component representing an unknown target is given by

$$u_{k-1}(\mathbf{y}_{k-1}) = \sum_{i=1}^{N} w_{k-1}^{u,i} f_{k-1}^{u,i}(\tau_{k-1}) f_{k-1}^{u,i}(s_{k-1}) f_{k-1}^{u,i}(\mathbf{x}_{k-1}|s_{k-1}, \tau_{k-1}). \tag{4.12}$$

which is a sum of mixture components, where $w_{k-1}^{u,i}$ is the weight of component i. A potentially detected target i in a global hypothesis j at the previous time step is represented by a Bernoulli with existence probability $r_{k-1}^{j,i}$, weight $w_{k-1}^{j,i}$ and state density

$$f_{k-1}^{j,i}(\mathbf{y}_{k-1}) = f_{k-1}^{j,i}(\mathbf{x}_{k-1}|\tau_{k-1}, s_{k-1}) \times f_{k-1}^{j,i}(s_{k-1}|\tau_{k-1}) f_{k-1}^{j,i}(\tau_{k-1}). \tag{4.13}$$

For ease of notation, we henceforth write $f_{k-1}^{\cdot,i}(\tau_{k-1})$ as $\xi_{k-1}^{\cdot,i\tau}$ and $f_{k-1}^{\cdot,i}(s_{k-1}|\tau_{k-1})$ as $\mu_{k-1}^{\cdot,is\tau}$.

4.4.1 Prediction

We find the expressions for the prediction by use of the expressions from Section 3.5 and [47, Sec. V]. For the Poisson component, we get that the predicted intensity is

$$u_{k|k-1}(\mathbf{y}_k) = b_k(\mathbf{y}_k) + P_S(\mathrm{d}t) \sum_{i=1}^{N} w^{u,i} \mu_{k|k-1}^{u,is\tau} \xi_{k|k-1}^{u,i\tau} f_{k|k-1}^{u,i}(\mathbf{x}_k|s_k, \tau_k) \quad (4.14)$$

whereas the predicted Bernoulli components are

$$f_{k|k-1}^{j,i}(\mathbf{y}_k) = \mu_{k|k-1}^{j,is\tau} \xi_{k|k-1}^{j,i\tau} f_{k|k-1}^{j,i}(\mathbf{x}_k|s_k, \tau_k). \quad (4.15)$$

Furthermore, we predict the existence probability and discrete states in the hybrid state as

$$r_{k|k-1}^{\cdot,i} = r_{k-1}^{\cdot,i} P_S(\mathrm{d}t) \quad (4.16)$$

$$\xi_{k|k-1}^{\cdot,i\tau} = \xi_{k-1}^{\cdot,i\tau} \quad (4.17)$$

$$\mu_{k|k-1}^{\cdot,is\tau} = \sum_{\tilde{s}} \mu_{k-1}^{\cdot,i\tilde{s}\tau} \pi^{\tilde{s}s}(\mathrm{d}t). \quad (4.18)$$

where $\cdot$ denotes that the expressions are valid for both unknown and potentially detected targets. The kinematic density is predicted as

$$f_{k|k-1}^{\cdot,i}(\mathbf{x}_k|\tau_k, s_k) = \int f_{\mathbf{x}}(\mathbf{x}_k|\tau_k, s_k, \tilde{\mathbf{x}}) f_{k-1}^{\cdot,i}(\tilde{\mathbf{x}}|\tau_k, s_k) \mathrm{d}\tilde{\mathbf{x}} \quad (4.19)$$

where

$$f_{k-1}^{\cdot,i}(\tilde{\mathbf{x}}|\tau_k, s_k) = \sum_{\tilde{s}} \frac{\mu_{k-1}^{\cdot,i\tau\tilde{s}} \pi^{\tilde{s}s}(\mathrm{d}t) f_{k-1}^{\cdot,i}(\tilde{\mathbf{x}}|\tau_k, s_k, \tilde{s})}{\sum_{\tilde{s}} \mu_{k-1}^{\cdot,i\tau\tilde{s}} \pi^{\tilde{s}s}(\mathrm{d}t)}. \quad (4.20)$$

4.4.2 Update

The posterior is also found by use of Section 3.5 in combination with [47, Sec. V]. Three different types of updates must be considered:

- Update of undetected targets.
- Update of new potentially detected targets.
- Update of previously potentially detected targets.

The update of undetected targets only involves updating the weight. For convenience, the unknown target intensity is rewritten as

$$u_k(\mathbf{y}_k) = \sum_{i=1}^{N_u} w_{k|k-1}^{u,i} f_{k|k-1}^{u,i}(\mathbf{y}_k). \tag{4.21}$$

We update the weights by multiplying them with the probability of a missed detection

$$w_k^{u,i} = w_{k|k-1}^{u,i}(1 - P_D). \tag{4.22}$$

Update of new potentially detected targets

When a potential new target is detected, a new Bernoulli is initialized. We need to find the Bernoulli's existence probability $r_k^{j,i}(Z_k)$, state density $f_k^{j,i}(\mathbf{y}_k|Z_k)$, and weight $w^{j,i}(Z_k)$. This is done by updating the unknown target intensity with some non-empty subset Z_k of all the received detections. We have from [47, Sec. V] that

$$r_k^{j,i}(Z_k) = \frac{e(Z_k)}{\rho(Z_k)} \tag{4.23}$$

$$f_k^{j,i}(\mathbf{y}_k|Z_k) = \sum_{i_u=1}^{N_u} w^{i_u}(Z_k) f_k^{u,i_u}(\mathbf{y}_k|Z_k) \tag{4.24}$$

where

$$\begin{aligned} e(Z_k) &= \sum_{i_u=1}^{N_u} w^{u,i_u} f_{\mathbf{z}}^{u,i_u}(Z_k|\mathbf{y}_k) \\ \rho(Z_k) &= e(Z_k) + \lambda(Z_k) \\ w^{i_u}(Z_k) &\propto w_{k|k-1}^{u,i_u} f_{\mathbf{x}}^{u,i_u}(Z_k|\mathbf{y}_k). \end{aligned} \tag{4.25}$$

The weight $w_k^{j,i}$ of the new Bernoulli in a global hypothesis j has value $\rho(Z_k)$ if the global hypothesis includes the new target, and otherwise the weight is set to 1 with the existence probability set to 0. Furthermore, we need to find expressions for $f_k^{u,i_u}(\mathbf{y}_k|Z_k)$ and $f_{\mathbf{x}}^{u,i_u}(Z_k|\mathbf{y}_k)$. We provide these on a general form that holds for exteroceptive measurements and target-provided information in both point target and extended object tracking, and from Section 3.5 we get that

$$f_k^{\cdot,i}(\mathbf{y}_k|Z_k) = f_{k|k-1}^{\cdot,i}(\mathbf{x}_k|\tau_k, s_k, Z_k)\mu_k^{\cdot,is\tau}\xi_k^{\cdot,i\tau} \tag{4.26}$$

where

$$f_k^{\cdot,i}(\mathbf{x}_k|\tau_k, s_k, Z_k) = \frac{f_{\mathbf{z}}^{\cdot,i}(Z_k|\tau_k, s_k, \mathbf{x}_k) f_{k|k-1}^{\cdot,i}(\mathbf{x}_k|\tau_k, s_k)}{l^{\cdot,i\tau s}} \tag{4.27}$$

and

$$\mu_k^{\cdot,is\tau} = \frac{\mu_{k|k-1}^{\cdot,i\tau s} l^{\cdot,i\tau s}}{\sum_{\tilde{s}} \mu_{k|k-1}^{\cdot,i\tau\tilde{s}} l^{\cdot,i\tau\tilde{s}}} \tag{4.28}$$

$$\xi_k^{\cdot,i\tau} = \frac{\xi_{k|k-1}^{\cdot,i\tau} \sum_{\tilde{s}} \mu_{k|k-1}^{\cdot,i\tau\tilde{s}} l^{\cdot,i\tau\tilde{s}}}{\sum_{\tilde{\tau}} \xi_{k|k-1}^{\cdot,i\tilde{\tau}} \sum_{\tilde{s}} \mu_{k|k-1}^{\cdot,i\tilde{\tau}\tilde{s}} l^{\cdot,i\tilde{\tau}\tilde{s}}}. \tag{4.29}$$

Furthermore, we have that

$$l^{\cdot,i\tau s} = \int f_{\mathbf{z}}^{\cdot,i}(Z_k|\tau_k, s_k, \tilde{\mathbf{x}}) f_{k|k-1}^{\cdot,i}(\tilde{\mathbf{x}}|\tau_k, s_k) \mathrm{d}\tilde{\mathbf{x}}. \tag{4.30}$$

Update of previously potentially detected targets

A potentially detected target can either be updated based on a detection or a missed detection. First, we define the combined likelihood of a measurement set across all discrete states as

$$L^{j,i} = \sum_{\tilde{\tau}} \xi_{k|k-1}^{j,i\tilde{\tau}} \sum_{\tilde{s}} \mu_{k|k-1}^{j,i\tilde{\tau}\tilde{s}} l^{j,i\tilde{\tau}\tilde{s}}. \tag{4.31}$$

Note that a missed detection would mean that the likelihood (4.30) is that of an empty set of measurements. From Section 3.5, we have that for the missed detection case

$$w_k^{j,i} = \begin{cases} w_{k|k-1}^{j,i}(1 - r_{k|k-1}^{j,i} + r_{k|k-1}^{j,i} L^{j,i}) \\ \quad \text{for exteroceptive} \\ w_{k|k-1}^{j,i} \quad \text{for target-provided} \end{cases} \tag{4.32}$$

$$r_k^{j,i} = \begin{cases} \dfrac{r_{k|k-1}^{j,i} L^{j,i}}{1 - r_{k|k-1}^{j,i} + r_{k|k-1}^{j,i} L^{j,i}} & \text{for exteroceptive} \\ r_{k|k-1}^{j,i} & \text{for target-provided} \end{cases} \tag{4.33}$$

whereas the state density remains unchanged from the prediction, a consequence of the model choice of state-independent detection probability. We distinguish between the two measurement types to highlight how the absence of target-provided information does not impact the weight and existence probability, as it does for the exteroceptive measurements. This is due to how the detection probability (4.10) is defined. Updating the weights based on a detection is done as

$$w_k^{j,i} = w_{k|k-1}^{j,i}(r_{k|k-1}^{j,i} L^{j,i}). \tag{4.34}$$

The updated existence probability $r_k^{j,i}$ is 1, and the state density is updated by use of (4.26)-(4.29).

4.4.3 Global hypotheses

With the updated Bernoulli components we now need to form global hypotheses and calculate their weights. This is done as described in [47, Sec. C.3], from which we give a summary. For each previous global hypothesis j, we must assign each new measurement to either an existing track or a new track. One such set of assignments amounts to a new global hypothesis. Its weight w_k^j is calculated as the product of the weights of the individual Bernoulli components in the hypothesis. The number of global hypotheses can quickly become untenable, so rather than considering all possible assignments, we only consider the most likely assignments. These can, for example, be found by use of Murty's method [104] or stochastic optimization [55].

4.5 Application to point target tracking and extended object tracking

We want closed form recursions for the AIS-IMM-PMBM filter and the AIS-GP-PMBM filter. For that purpose, we specify the state spaces, kinematic models, and measurement models for the two filters. We show how the resulting expressions relate to those in Section 4.4, and that they allow us to perform the calculations by use of Kalman filtering. Furthermore, we show how to estimate the states we want to output to the surrounding system.

4.5.1 The AIS-IMM-PMBM filter for point target tracking

For the AIS-IMM-PMBM point target tracker, the hybrid state is

$$\mathbf{y} = \begin{bmatrix} \mathbf{x} \ \tau \ s \end{bmatrix}^\top \tag{4.35}$$

in which the kinematic state $\mathbf{x}$ is

$$\mathbf{x} = \begin{bmatrix} x \ v_x \ y \ v_y \ \omega \end{bmatrix}^\top. \tag{4.36}$$

Here, x and y is the position, v_x and v_y are the velocities, and ω is the angular velocity. We have omitted the time index, the global hypothesis index j, and the track index i for brevity. We model the estimated kinematic states, the kinematic state transition, and the measurements as Gaussians distributions. That is, we have that the state density (4.13) is given by

$$f(\mathbf{y}) = \mathcal{N}(\mathbf{x}; \hat{\mathbf{x}}^{s\tau}, \mathbf{P}^{s\tau}) \mu^{s\tau} \xi^{\tau} \tag{4.37}$$

where $\hat{\mathbf{x}}^{s\tau}$ and $\mathbf{P}^{s\tau}$ are the estimated mean and covariance conditioned on kinematic model s and ID τ, $\mu^{s\tau}$ is the probability of kinematic model s conditioned on ID τ, and ξ^{τ} is the probability of ID τ.

We use two constant velocity (CV) models and one coordinated turn (CT) model to model the movement of the point targets. The CV models model linear, straight-line motion, whereas the CT model in addition models the possibility of a target turning. In general, the state evolves according to

$$\mathbf{x}_k = \mathbf{F}^s(\mathbf{x}_{k-1})\mathbf{x}_{k-1} + \mathbf{v}_k, \quad \mathbf{v}_k \sim \mathcal{N}(\mathbf{0}, \mathbf{Q}^s). \tag{4.38}$$

The CV models are defined as

$$\mathbf{F}^{CV} = \begin{bmatrix} 1 & \mathrm{dt} \\ 0 & 1 \end{bmatrix} \otimes \mathbf{I}_2, \tag{4.39}$$

$$\mathbf{Q}^{CV} = \begin{bmatrix} {}^{(\mathrm{dt})^3}\!/_3 & {}^{(\mathrm{dt})^2}\!/_2 \\ {}^{(\mathrm{dt})^2}\!/_2 & \mathrm{dt} \end{bmatrix} \otimes \begin{bmatrix} q_a & 0 \\ 0 & q_a \end{bmatrix}. \tag{4.40}$$

and the CT model is defined as

$$\mathbf{F}^{CT}(\mathbf{x}_k) = \begin{bmatrix} 1 & 0 & {}^{\sin \mathrm{dt}\omega}\!/_\omega & {}^{-1+\cos \mathrm{dt}\omega}\!/_\omega & 0 \\ 0 & 1 & {}^{1-\cos \mathrm{dt}\omega}\!/_\omega & {}^{\sin \mathrm{dt}\omega}\!/_\omega & 0 \\ 0 & 0 & \cos \mathrm{dt}\omega & -\sin \mathrm{dt}\omega & 0 \\ 0 & 0 & \sin \mathrm{dt}\omega & \cos \mathrm{dt}\omega & 0 \\ 0 & 0 & 0 & 0 & 1 \end{bmatrix}, \tag{4.41}$$

$$\mathbf{Q}^{CT} = \begin{bmatrix} \mathbf{Q}^{CV} & \mathbf{0} \\ \mathbf{0} & \mathrm{dt}q_\omega \end{bmatrix}. \tag{4.42}$$

This is a non-linear model, so for use in an extended Kalman filter it is linearized as it is done in [8, Ch. 11.72]. The kinematic state transition density is given by

$$f_{\mathbf{x}}(\mathbf{x}_k|\mathbf{x}_{k-1}, \tau_k, s_k) = \mathcal{N}(\mathbf{x}_k; \mathbf{F}^{s_k\tau}(\mathbf{x}_{k-1}), \mathbf{Q}^{s_k\tau}). \tag{4.43}$$

Here, $\mathbf{F}^{s_k\tau}(\mathbf{x}_{k-1})$ is the state transition matrix for kinematic model s_k conditioned on ID τ, and $\mathbf{Q}^{s_k,\tau}$ is the process noise covariance matrix. By using (4.37) and (4.43) in (4.19), we get the predicted kinematic states. The Gaussian mixture in (4.20) can be approximated by use of moment matching.

Update with radar measurements

The radar measurements are modeled as

$$\mathbf{z} = \mathbf{H}_R\mathbf{x} + \mathbf{w}, \;\; \mathbf{w} \sim \mathcal{N}(\mathbf{0}, \mathbf{R}_R) \tag{4.44}$$

where $\mathbf{H}_R$ is the measurement matrix, $\mathbf{R}_R$ is the measurement noise matrix, and $\mathbf{w}$ is the measurement noise. We model the measurement noise as a combination of Cartesian and polar noise

$$\mathbf{R}_R = \mathbf{R}_c + \mathbf{R}_p \tag{4.45}$$

where $\mathbf{R}_C$ is the Cartesian measurement noise and $\mathbf{R}_p$ is the polar measurement noise, which is converted to Cartesian coordinates with the method from [95]. The Cartesian noise is meant to account for errors from clustering and sensor noise, and the polar noise is meant to account for errors in the range and bearing.

We assume that a set Z of radar measurements contains either one or zero measurements. The measurement likelihood for a single radar measurement is given by

$$f_{\mathbf{z}}(\mathbf{z}|\mathbf{x}, \tau, s) = \mathcal{N}(\mathbf{z}; \mathbf{H}_R\mathbf{x}, \mathbf{R}_R) \tag{4.46}$$

and the likelihood for a set of radar measurements is given by

$$f_{\mathbf{z}}(Z|\mathbf{x}, \tau, s) = \begin{cases} P_D f_{\mathbf{z}}(\mathbf{z}|\mathbf{x}, \tau, s) & Z = \{\mathbf{z}\} \\ 1 - P_D & Z = \emptyset \end{cases} \tag{4.47}$$

which is used in (4.27) and (4.30) to get a closed-form solution. Note that the case with the empty measurement set is only relevant for previous potentially detected targets. Furthermore, the radar clutter measurements are modeled as a PPP with constant intensity $\lambda(Z) = N_c/\pi R^2$. Here N_c is the expected number of clutter measurements and R is the radius of the surveillance area.

Update with AIS measurements

The position and velocity part of the AIS messages are defined as

$$\mathbf{z} = \mathbf{H}_A\mathbf{x} + \mathbf{w}, \;\; \mathbf{w} \sim \mathcal{N}(\mathbf{0}, \mathbf{R}_A) \tag{4.48}$$

where $\mathbf{H}_A$ is the measurement matrix and $\mathbf{R}_A$ is the measurement noise matrix. As for the radar measurements, the noise matrix is a combination of Cartesian and polar noise, but here the Cartesian noise models the error in the position and the polar noise models the error in the velocity. This is because the positional errors are independent on the distance, and derived from GPS data, whereas the velocity is transmitted as a speed and course. The measurement noise matrix is given by

$$\mathbf{R}_A = \mathbf{H}\mathbf{R}_{c,A} + \mathbf{H}^{\text{vel}}\mathbf{R}_{p,A} \tag{4.49}$$

where $\mathbf{H}$ and $\mathbf{H}^{\text{vel}}$ are measurement matrices for the position and velocity parts of the state, respectively. We use the same method as above when we convert the polar noise matrix to Cartesian coordinates.

For the position- and velocity-information in the AIS messages, the measurement likelihood $f^{tp}_{\mathbf{p}}(\mathbf{p}|\mathbf{y})$ in (4.11) is given by

$$f^{tp}_{\mathbf{p}}(\mathbf{p}|\mathbf{y}) = \mathcal{N}(\mathbf{z}; \mathbf{H}_A\mathbf{x}, \mathbf{R}_A). \tag{4.50}$$

Furthermore, we utilize the ID information in the AIS messages. We account for the small possibility of the ID provided by the measurement is incorrect, relative to the actual ID of the target, by use of a probability P_C representing our confidence in the ID provided by the measurement being correct. The likelihood of the ID only depends on the ID of the target. We formulate $f_\tau^{tp}(\tau^{\mathbf{z}}|\mathbf{y}_k)$ as

$$f_\tau(\tau^{\mathbf{z}}|\tau) = \begin{cases} P_C & \text{if } \tau = \tau^{\mathbf{z}} \\ \dfrac{1-P_C}{|\mathcal{V}|-1} & \text{if } \tau \neq \tau^{\mathbf{z}} \text{ and } \tau > 0 \\ 0 & \text{if } \tau = 0 \end{cases} \tag{4.51}$$

where $\tau^{\mathbf{z}}$ is the ID provided by the measurement. A zero-valued ID τ is used to represent a target that does not transmit any information, and as such has no observable ID. We combine (4.50) and (4.51) in (4.11), and disregard the length- and width-related terms in the latter equation. This provides a closed form solution to the integral in (4.30), and furthermore allows us to calculate the updated kinematic states in (4.27) by use of the Kalman filter equations.

4.5.2 The AIS-GP-PMBM filter for extended object tracking

The hybrid state in the AIS-GP-PMBM filter is

$$\mathbf{y} = \begin{bmatrix} \mathbf{x} \; \tau \; \alpha \; \beta \end{bmatrix}^\top \tag{4.52}$$

and includes the ID τ and two gamma distribution shape parameters α and β. The gamma distribution is used to estimate the expected number of detections from a target. Furthermore, $\mathbf{x}$ is given by

$$\mathbf{x} = \begin{bmatrix} x \; v_x \; y \; v_y \; \phi \; \omega \; \mathbf{x}^f \end{bmatrix}^\top. \tag{4.53}$$

Here, ϕ is the target heading, ω is the angular velocity, and $\mathbf{x}^f$ is a vector which parametrizes the contour of the target extent, specifically it contains the values of a radius function f at equidistant angles. The single state estimates are represented as gamma-Gaussian distributions on the form

$$f(\mathbf{y}) = \mathcal{N}(\mathbf{x}; \hat{\mathbf{x}}^\tau, \mathbf{P}^\tau)\mathcal{G}(\alpha, \beta)\xi^\tau. \tag{4.54}$$

where we again have omitted the time index, the track index i, and the global hypothesis index j for brevity. For the augmented state in the AIS-GP-PMBM filter, we combine the CV model with a process model for the extent. In addition to linear velocity, the CV model also models the heading and angular velocity, and is defined

as

$$\mathbf{F}^{CV} = \begin{bmatrix} 1 & \mathrm{dt} \\ 0 & 1 \end{bmatrix} \otimes \mathbf{I}_3 \tag{4.55}$$

$$\mathbf{Q}^{CV} = \begin{bmatrix} (\mathrm{dt})^3/3 & (\mathrm{dt})^2/2 \\ (\mathrm{dt})^2/2 & \mathrm{dt} \end{bmatrix} \otimes \begin{bmatrix} q_a & 0 & 0 \\ 0 & q_a & 0 \\ 0 & 0 & q_\theta \end{bmatrix}. \tag{4.56}$$

We have that $\mathbf{F} = diag(\mathbf{F}^{CV}, \mathbf{F}^f)$ and $\mathbf{Q} = diag(\mathbf{Q}^{CV}, \mathbf{Q}^f)$ where

$$\begin{aligned} \mathbf{F}^f &= \exp(-\Delta \mathrm{t} \eta_\gamma)\mathbf{I}, \\ \mathbf{Q}^f &= (1 - \exp(-2\Delta \mathrm{t} \eta_\gamma))\mathbf{K}(\Theta^f, \Theta^f). \end{aligned} \tag{4.57}$$

$\mathbf{K}(\Theta^f, \Theta^f)$ is the covariance matrix of the Gaussian process, and Θ^f contains the angles of the points which define the extent. For further details, we refer to [134] for the derivation of the Gaussian process model, or to [6] for a shorter summary.

When predicting and updating the single state estimates, conditioned on their IDs, we separate the gamma and Gaussian parts of the distributions. This allows us to use the Kalman filter equations to calculate the predicted and updated Gaussian parts the same way as in Section 4.5.1, whereas the variables in the gamma distribution are predicted as

$$\alpha_{k|k-1} = \alpha_{k-1}/\eta_\gamma, \quad \beta_{k|k-1} = \beta_{k-1}/\eta_\gamma \tag{4.58}$$

and updated as

$$\alpha_k = \alpha_{k|k-1} + |Z|, \quad \beta_k = \beta_{k|k-1} + 1. \tag{4.59}$$

Here, η_γ is called the forgetting factor and is a parameter that controls how quickly the measurements received in the past should be forgotten by the gamma distribution parameters.

Update with LiDAR measurements

A generic measurement equation for one contour generated measurement with the target contour parametrized by a radial function f can be written as

$$\begin{aligned} \mathbf{z}^l &= \mathbf{x}^c + \mathbf{p}(\theta^l) f(\theta^l) + \mathbf{w}^l \\ \mathbf{p}(\theta^l) &= \begin{bmatrix} \cos\theta^l \\ \sin\theta^l \end{bmatrix} \end{aligned} \tag{4.60}$$

where $\mathbf{z}_k^l$ is measurement l and θ^l is the corresponding angle of the origin of the measurement of the target contour. θ^l can be expressed both in a global frame $\theta^{l(G)}$

and the local target body frame $\theta^{l(B)}$ as

$$\begin{aligned}\theta^{l(B)}(\mathbf{x}^c, \phi) &= \theta^{l(G)}(\mathbf{x}^c) - \phi \\ \theta^{l(G)}(\mathbf{x}^c) &= \angle\left(\mathbf{z}^l - \mathbf{x}^c\right)\end{aligned} \tag{4.61}$$

The value of $f(\theta^l)$ can be found by calculating the value of the matrix $\mathbf{H}^{(f)}$ for the specific angle using Gaussian process regression and multiplying it with the vector parametrizing the extent $\mathbf{x}^f$. Therefore, the measurement equation can be written as

$$\begin{aligned}\mathbf{z}^l &= \mathbf{x}^c + \mathbf{p}^l(\theta^{l(G)}(\mathbf{x}^c))\mathbf{H}^{(f)}\left(\theta^{l(B)}(\mathbf{x}^c, \phi)\right)\mathbf{x}^f + \mathbf{w}^l \\ &= h^l(\mathbf{x}) + \mathbf{w}^l, \qquad \mathbf{w}^l \sim \mathcal{N}(\mathbf{0}, \mathbf{R}^l).\end{aligned} \tag{4.62}$$

The measurement equation is therefore dependent on the state space components $\mathbf{x}^c$, ϕ and $\mathbf{x}^f$, of which the two former are non-linear terms in the measurement equation. It should be noted that this is an implicit equation, due to the dependency of $\theta^{l(G)}$ on $\mathbf{z}^l$. Similarly, the estimated error of the Gaussian process regression can be calculated by calculating a matrix $\mathbf{R}^f$ and by projecting this into 2D we get

$$\mathbf{R}^l = \mathbf{p}^l(\mathbf{x}^c)\mathbf{R}^f\left(\theta^{l(B)}(\mathbf{x}^c, \phi)\right)\mathbf{p}^l(\mathbf{x}^c)^T + \mathbf{R}_C. \tag{4.63}$$

We use an iterated extended Kalman filter to deal with the non-linearities that are introduced by the LiDAR measurement model.

A global association hypothesis in extended object tracking does not assign each measurement to a single potential target, but rather assigns sets of measurements to potential targets. The measurement likelihood for a single measurement can thus be written as

$$f_{\mathbf{z}}(\mathbf{z}^l|\mathbf{x}) = N(\mathbf{H}^l\mathbf{x}, \mathbf{R}^l). \tag{4.64}$$

Each potential target i is associated to a specific measurement cell C and the measurements in such a cell is denoted Z_C. The measurement set likelihood is given by the following inhomogeneous PPP [52]

$$f_{\mathbf{z}}(Z_C|\mathbf{x}) = \exp(-\lambda_m)\lambda_m^{|Z_C|}\prod_{\mathbf{z}^l \in Z_C} f_{\mathbf{z}}(\mathbf{z}^l|\mathbf{x}) \tag{4.65}$$

here λ_m is the Poisson rate governing the expected number of measurements, which is estimated by the gamma distribution, i.e. $\lambda_m \sim \mathcal{G}(\alpha, \beta)$. Given this, the predictive likelihood can be calculated as in (4.30), which results in

$$l^\tau = P_D\frac{\Gamma(\alpha + |Z_C|)\beta^\alpha}{\Gamma(\alpha)(\beta+1)^{(\alpha+|Z_C|)}|Z_C|!}\prod_{\mathbf{z}^l \in Z_C}\mathcal{N}(\mathbf{z}^l; \mathbf{H}^l\hat{\mathbf{x}}^{\tau,l}, \mathbf{S}^{\tau,l}) \tag{4.66}$$

where $\mathbf{S}^{\tau,l}$ is the innovation covariance matrix for measurement l conditioned on ID τ. If the measurement cell is empty, the predictive likelihood is instead given by the effective probability of missed detection defined by

$$l^{\tau} = 1 - P_D + P_D \exp(\lambda_m) \tag{4.67}$$

which represents the fact that we can have a missed detection either due to the probability of detection or the probability that the target is detected but generates zero measurements. Lastly, as for the radar measurements, we assume that the clutter intensity is uniform over the surveillance area. Furthermore, the intensity also reflects the result that a cell with more than one measurement never is a false alarm. This means that

$$\lambda(Z_C) = \begin{cases} N_c/\pi R^2 & |Z_C| = 1 \\ 0 & |Z_C| > 1 \end{cases}. \tag{4.68}$$

Further details regarding the Gaussian process PMBM filter can be found in [6].

Update with AIS measurements

The incoming target-provided measurements are handled in a similar manner as for the point target tracking case, as they uphold the assumption of each target only providing a single measurement. The difference is that we now include the dimensions of the target in the measurement vector. The dimension is modeled by considering length and width separately. The measurement model matrix can be found using Gaussian process regression, in this case by using fixed angles which correspond to the length

$$z^L = \mathbf{H}^{(f)}(0)\mathbf{x}^f + \mathbf{H}^{(f)}(\pi)\mathbf{x}^f + w^L, \qquad w^L \sim \mathcal{N}(0, \sigma^2_{z,L}) \tag{4.69}$$

and width

$$z^W = \mathbf{H}^{(f)}\left(\frac{\pi}{2}\right)\mathbf{x}^f + \mathbf{H}^{(f)}\left(\frac{3\pi}{2}\right)\mathbf{x}^f + w^W, \qquad w^W \sim \mathcal{N}(0, \sigma^2_{z,W}). \tag{4.70}$$

The estimated measurement noise covariance has a component for Gaussian process regression, and is given by

$$\begin{aligned} \sigma^2_L &= \mathbf{R}^{(f)}(0) + \mathbf{R}^{(f)}(\pi) + \sigma^2_{z,L} \\ \sigma^2_W &= \mathbf{R}^{(f)}\left(\frac{\pi}{2}\right) + \mathbf{R}^{(f)}\left(\frac{3\pi}{2}\right) + \sigma^2_{z,W}. \end{aligned} \tag{4.71}$$

Using this, we get that the AIS measurement likelihood (4.11) becomes

$$f_{\mathbf{z}}(Z|\mathbf{y}) = f_\tau(\tau^{\mathbf{z}}|\tau)\mathcal{N}(\mathbf{p};\mathbf{H}\mathbf{x}^c,\mathbf{R}^c)\times \mathcal{N}(z^L;\mathbf{H}^{(f)}(0)\mathbf{x}^f+\mathbf{H}^{(f)}(\pi)\mathbf{x}^f,\sigma_L^2)\times \mathcal{N}(z^W;\mathbf{H}^{(f)}\left(\frac{\pi}{2}\right)\mathbf{x}^f+\mathbf{H}^{(f)}\left(\frac{3\pi}{2}\right)\mathbf{x}^f,\sigma_W^2) \quad (4.72)$$

when $Z=\{\mathbf{z}=[\mathbf{p},\tau,z^L,z^W]\}$. This is used in (4.27) and (4.30). The gamma distribution is only concerned with the exteroceptive measurements and is as such not updated when target-provided measurements are received because no new information is received. It is also ignored when calculating the target-provided measurement likelihoods.

Remark 4.1. Note that we do not use multiple models in the extended object AIS-GP-PMBM. This is equivalent to having a single model with constant probability 1, and the simplified expressions are easily derived from the ones in Sections 4.4.1 and 4.4.2. ◁

4.5.3 State estimation and complexity management

The implementation for the point target tracker largely follows [47], whereas the extended object tracker is based on [52]. The way we perform state estimation and manage the computational complexity is, however, conceptually the same. After each update step we have several global hypotheses, each containing Bernoulli components with potentially detected targets. In [47] three estimation methods are presented, which allows us to choose which of the large number of possible target states we decide are the most likely. For both point targets and extended objects, we use the first of the three methods, which simply decides upon the global hypothesis with the highest weight. Of the Bernoulli components in the chosen global hypothesis, all with an existence probability larger than some threshold T_r are selected as the output. Furthermore, each state estimate is a Gaussian mixture with weights corresponding to the probabilities of the discrete states, such as IDs. A single Gaussian is extracted by use of moment matching.

Even when limiting the amount of new global hypotheses, with Murty's algorithm for point targets and stochastic optimization for extended objects, the number of Bernoulli components can become very large. Thus, to avoid an unmanageable number of possible target states some approximations are made. First, we limit the amount of Bernoulli components created at each time step by way of gating the measurements. This is done by only considering the measurements less than $\sqrt{g}$ standard deviations from a given prediction. Thus, new potential targets are only

created for measurements close enough to an unknown target, and already potentially detected targets are only updated based on the measurements within its gate area.

Furthermore, we do not propagate all Bernoulli components from each time step to the next. We follow the strategy from [47], and only keep the N_{hyp} global hypotheses with highest weights. In addition, we remove all Bernoulli components with existence probability lower than some threshold T_b and those not present in any of the kept global hypotheses. Furthermore, we only keep the Poisson densities with weight higher than T_p. A brief description covering one iteration of the AIS-PMBM is given in Algorithm 4.

Algorithm 4 An iteration of the AIS-PMBM

Input: Previous unknown target densities $\mu(\mathbf{y})$, potentially detected target densities $p^{i,j}(\mathbf{y})$, global hypotheses, and new measurements Z

Output: Updated unknown target densities $\mu(\mathbf{y})$, potentially detected target densities $p^{i,j}(\mathbf{y})$, global hypotheses

1: Predict unknown target densities with (4.14) and perform gating.
2: Initialize new potentially detected targets on the gated measurements with the expressions from Section 4.4.2.
3: Predict potentially detected targets with (4.16)-(4.19) and perform gating.
4: Initialize new Bernoulli components for the previously potentially detected targets with the expressions from Section 4.4.2.
5: Find new global hypotheses based on the previous global hypotheses and the new Bernoulli components as described in Section 4.4.3.
6: Output the Bernoulli components in the best global hypothesis with existence probability higher than T_r.
7: Remove all but the N_{hyp} best global hypotheses, prune Bernoulli components with low existence probability, and prune Poisson densities with low weight.

4.6 Results for point target tracking

To evaluate the performance of the trackers we use the GOSPA metric for trajectories, presented in [46]. The metric works on sets of trajectories, and penalizes track switches in addition to localization errors, false alarms, and missed detections. It also allows us to look at the different error sources in isolation, highlighting the

advantages and shortcomings of different methods. The metric is defined as

$$d_p^{(c,\gamma)}(X,Y) \triangleq \min_{\substack{a_k \in \prod^{X,Y} \\ k=1,\dots,T}} \left(\sum_{k=1}^{T} d_k^{X,Y}(X,Y,a_k)^p) + \sum_{k=1}^{T-1} s^{X,Y}(a_k,a_{k+1})^p \right)^{\frac{1}{p}} \tag{4.73}$$

where the first sum penalizes localization errors and cardinality errors, and the second sum the track switch error. a_k is one of the possible associations at time k between the sets of trajectories X and Y. In our case, the sets of trajectories are the true trajectories and the target states estimated as described in Section 4.5.3. The continuity of the estimated target states is based on their originating measurement. The parameters of the metric are the order p, the cutoff c, and the switch penalty γ. The cutoff decides the point where the distance between two tracks is too large to be considered a feasible association. We have that

$$\begin{aligned} d_k^{X,Y}(X,Y,a_k)^p \triangleq & \sum_{(i,j)\in\theta_k(a_k)} d(\mathbf{x}_k^i, \mathbf{y}_k^j)^p \\ & + \frac{c^p}{2}(|\tau_k(X)| + |\tau_k(Y)| - 2|\theta_k(a_k)|). \end{aligned} \tag{4.74}$$

Here, $d(\cdot,\cdot)$ is the distance function, $\theta_k(\cdot)$ is the set of feasible associations in the kth time step, and $\tau_k(\cdot)$ is the set of elements in a set of trajectories at time step k. The track switch error is defined as

$$s^{X,Y}(a_k,a_{k+1})^p \triangleq \gamma^p \sum_{i=1}^{n_X} s(a_k^i, a_{k+1}^i) \tag{4.75}$$

where $s(\cdot,\cdot)$ is 0 if the association between the trajectories is unchanged, 1 if the association is changed to a different trajectory, and $1/2$ if the association is changed and the trajectory is now unassigned, or was previously unassigned. We use the implementation available at [45] to perform the computations, and only compute errors at the time steps where exteroceptive measurements have been received.

4.6.1 Simulated data

We test seven different trackers and tracker configurations. These are

1. The MHT algorithm with AIS from [94].
2. The PMBM with AIS of Miao et al. [102].
3. The AIS-VIMMJIPDA from Chapter 3.

Table 4.1: Table continues on next page.

Quantity	Symbol Unit	Value (1st, 2nd)
Common parameters		
Survival probability	P_{S} [−]	0.999
Gate size	g [−]	20
Max. number of hypotheses	$N_{\max}$ [−]	400
Area radius	R [m]	1000, 100
Clutter intensity	λ [$1/\mathrm{m}^2$]	$4/(\pi R^2)$
Initial velocity std.	σ_v [$\mathrm{m/s}$]	10
Existence confirmation threshold	T_c [−]	0.99
Radar measurements		
Cartesian noise std.	σ_{c_R} [m]	1.0
Polar range std.	σ_r [m]	2.0
Polar bearing std.	σ_θ [°]	2
Detection probability	P_{D} [−]	0.9
AIS measurements		
Cartesian noise std.	σ_{c_R} [m]	1.0
Polar range std.	σ_r [m]	2.0
Polar bearing std.	σ_θ [°]	2
Confidence probability	P_C [−]	0.999
Unknown target no ID probability	ξ_u^0 [−]	0.9
PMBM parameters		
Poisson pruning threshold	T_{Pp} [−]	1.0×10^{-5}
Bernoulli pruning threshold	T_{Bp} [−]	1.0×10^{-6}
Birth weight	U [$1/\mathrm{m}^2$]	1×10^{-3}
CV model 1 process noise intensity	$q_{a,1}$ [$\mathrm{m}^2/\mathrm{s}^3$]	0.2^2
CV model 2 process noise intensity	$q_{a,2}$ [$\mathrm{m}^2/\mathrm{s}^3$]	2^2
Turn rate process noise intensity	q_ω [$\mathrm{rad}^2/\mathrm{s}^3$]	0.002^2
Single model process noise intensity	$q_{a,1}$ [$\mathrm{m}^2/\mathrm{s}^3$]	0.8^2
Initial model probability	μ_u^s [%]	$\begin{bmatrix} 25 & 25 & 50 \end{bmatrix}$
IMM transition probability	$\pi^{\tilde{s}s}$ [%]	$\begin{bmatrix} 99 & .5 & .5 \\ .5 & 99 & .5 \\ .5 & .5 & 99 \end{bmatrix}$

Table 4.1: Tuning parameters for the point target scenarios. If the parameters differ between the two scenarios, they are listed with the parameter for the first scenario first, and for the second scenario last.

Quantity	Symbol Unit	Value (1st, 2nd)
JIPDA parameters		
Existence termination threshold	T_d [−]	1.0×10^{-5}
CV model 1 process noise intensity	$q_{a,1}$ [$\mathrm{m}^2/\mathrm{s}^3$]	0.1^2
CV model 2 process noise intensity	$q_{a,2}$ [$\mathrm{m}^2/\mathrm{s}^3$]	1.5^2
Turn rate process noise intensity	q_ω [$\mathrm{rad}^2/\mathrm{s}^3$]	0.15^2
Initial existence probability	r_u [-]	0.18
Initial visibility probability	η_u [-]	0.9
Visibility transition probability	$\pi^{\tilde{s}s}$ [%]	$\left[\begin{smallmatrix} 48 & 52 \\ 10 & 90 \end{smallmatrix}\right]$
Initial model probability	μ_u^s [%]	$\left[\begin{smallmatrix} 80 & 10 & 10 \end{smallmatrix}\right]$
MHT parameters		
Number of tailed time steps	N [−]	5
Initialization window size	N_M [−]	5
Measurements needed for initialization	M [−]	2
Birth rate	U [$1/\mathrm{m}^2$]	1×10^{-3}
Cartesian noise std.	σ_{c_R} [m]	4.0
Process noise intensity	q [$\mathrm{m}^2/\mathrm{s}^3$]	1.0^2
Initial velocity std.	σ_v [m/s]	20
Similar track prune threshold	T_{STP} [−]	0

4. The PMBM operating only on the radar measurements.
5. The IMM-PMBM operating only on the radar measurements.
6. The AIS-PMBM.
7. The AIS-IMM-PMBM.

The trackers are tested on two different types of simulated data sets. The first of these is a data set configuration described in [138]. The data sets consider several targets that are all situated approximately at the origin of the area halfway through the simulation. The scenario is created by first choosing the points where the targets are to meet, and then stepping backwards and forwards with some kinematic model to create their trajectories. For the purposes of our simulations, we consider ten targets, whose midpoint positions and velocities are drawn from $\mathcal{N}(\mathbf{0}, 0.25 \times \mathbf{I})$, the same as for Case 2 in [138]. The radar and AIS measurements are generated

according to the measurement models described in Section 4.5, with the frequency of the AIS measurements decided by the requirements set by the protocol [73].

Furthermore, the different PMBM variants are tested on the Ravens data set detailed in [22]. The data set is designed to pose a challenge to the trackers and consists of eight targets and one ownship. During a span of 23 minutes, the targets and the ownship move in formation, making maneuvers underway. For the data to be suitable for testing of fusion between radar and AIS, additional AIS measurements are created for six of the eight targets.

The tuning parameters are shown in Table 4.1 for all the trackers. For the AIS-PMBM from Miao et al., we calculate the AIS detection probability directly from the total number of AIS messages received, together with the number of targets present in the area and their life span. Furthermore, the Cartesian noise component in the AIS measurement noise is increased to 10^2 to account for the time difference between transmission and processing of the measurements. Also note that the MHT algorithm does not implement the measurement noise with a polar component, and thus only uses Cartesian measurement noise. The process noise and initial model probabilities of the PMBM and JIPDA trackers differ, in the case of the initial mode probabilities significantly. During simulations, it became evident that the PMBM trackers needed more process noise than the JIPDA tracker to follow the targets successfully. This can be due to the step where the JIPDA combines the track-to-measurement associations, which has the observed effect of somewhat averaging the individual target movements across closely spaced targets.

Results

Results for the first data sets, with targets that meet in the middle of the surveillance area, are shown in Figure 4.1 and Table 4.2. Overall, the plots show that the most difficult parts of the scenario are at the very start, the middle, and at the very end. Most trackers perform best in the intervening periods. That the trackers struggle at the midpoint is no surprise, as all targets are closely spaced. Furthermore, as the targets are initialized with relatively low velocities and evolve from the midpoint, they usually have higher velocities at the start and end points. This means that they are harder to track correctly, and errors caused by delays in initialization and termination of tracks also contribute to the observed effect.

Next, we look at the performance of the different trackers relative to each other. It is evident that the MHT algorithm struggles more than the other trackers, especially with initializing tracks on all targets. Furthermore, the AIS-VIMMJIPDA performs well in many aspects but struggles in the middle of the scenario. These struggles manifest themselves in both false alarms, and missed detections which are present almost to the end of the scenario. Its IMM capabilities does, however, result in a low localization error, perhaps helped by the fact that the most troublesome targets are not

Table 4.2: GOSPA-T values for the different tracker configurations when tested on the data set configuration from [138], corresponding to the averaged values in Figure 4.2. The best result for each error source is highlighted in bold.

	PMBM	AIS-PMBM	IMM-PMBM	AIS-IMM-PMBM	AIS-PMBM from Miao et al.	AIS-VIMMJIPDA	MHT
Total GOSPA	10.424	9.964	8.809	**8.465**	10.111	13.211	27.013
Localization	9.107	8.718	7.943	**7.520**	8.702	9.341	12.300
Missed detection	1.013	0.739	0.511	**0.482**	1.185	4.649	15.852
Missed detection	1.013	0.739	0.511	**0.482**	1.185	4.649	15.852
False alarm	0.238	0.328	0.262	0.344	**0.206**	0.482	4.961
Track switch	1.055	1.140	0.732	0.755	0.962	**0.670**	1.513

tracked. For the different PMBM variants, the results show that both IMM and use of AIS measurements reduce the GOSPA value. The pure PMBM tracker outperforms the two previously mentioned trackers, albeit with regards to the AIS-VIMMJIPDA tracker by only a slight margin. Furthermore, all its augmentations perform even better. When including the AIS measurements, the general trend across the different trackers is an improvement in all error sources except for the false alarms. Because tracks initialized on AIS measurements get a higher weight than those initialized on radar measurements, this is somewhat expected. On the other hand, the use of AIS noticeably decreases the number of missed detections. Regarding the PMBM variant from [102], its performance lands between that of the regular PMBM and the AIS-PMBM presented here. This can be explained by the fact that the AIS messages are limited by a low P_D, thus giving them reduced influence on the results in comparison to the AIS-PMBM presented here. Increasing the P_D, however, could result in premature terminations when AIS messages are not available. Furthermore, we see that use of IMM improves performance noticeably for all error sources except the false alarms, where we see a slight increase.

For the Ravens data set we only consider the four different PMBM variants. The results are shown in Figure 4.2 and Table 4.3, and show the same trend as we saw in the previous data with more advanced methods performing better overall. There are, however, some interesting idiosyncrasies. The trackers with IMM provide better localization estimates than their counterparts before the midway point but

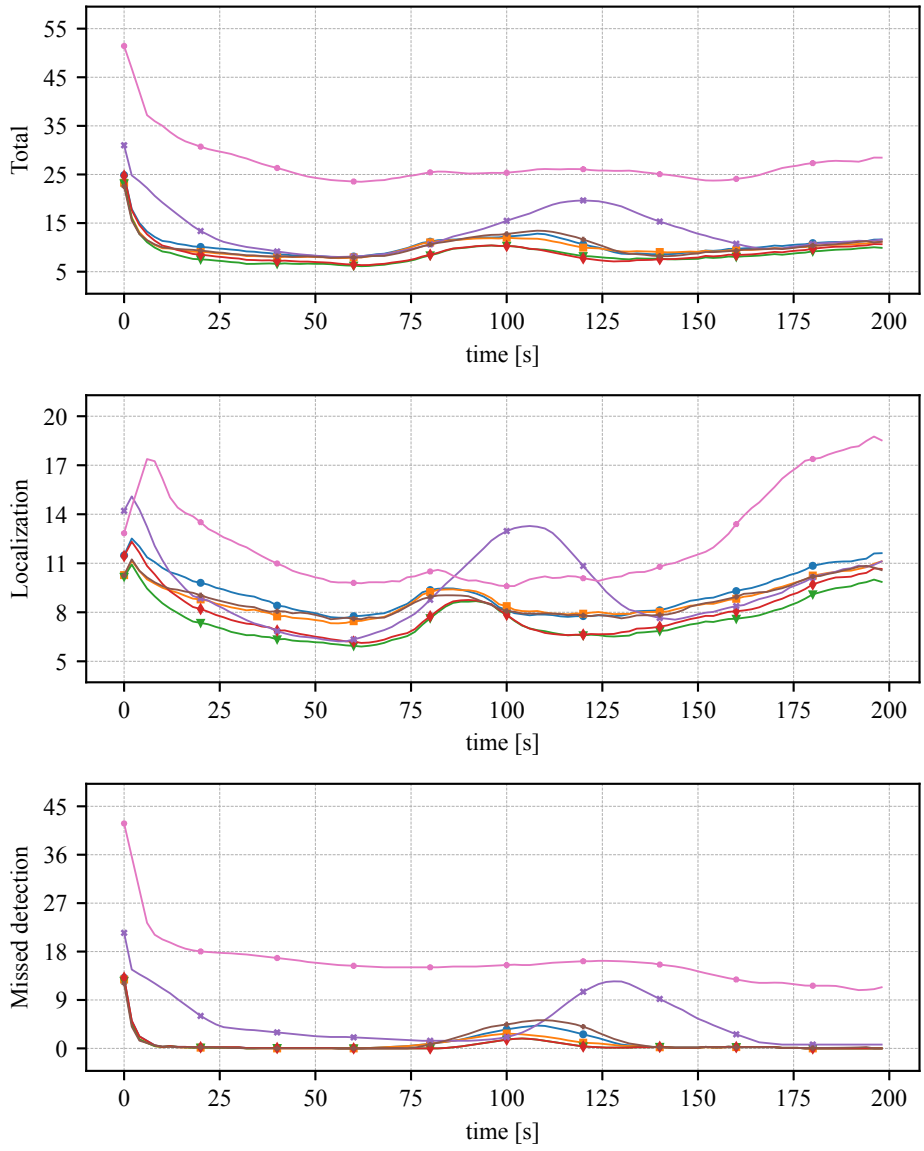

Figure 4.1: Figure with legend continues on next page.

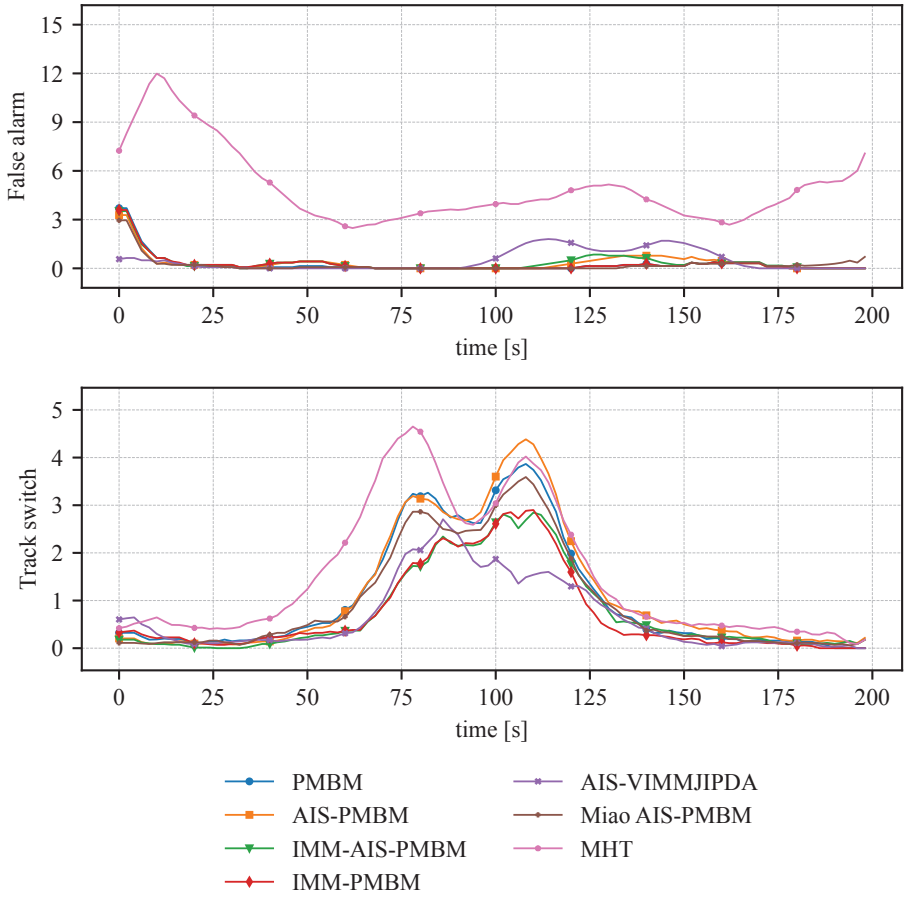

Figure 4.1: GOSPA-T values for the different trackers and configurations, where we see how the different trackers perform for different error sources evolve for the duration of the scenarios. We use the GOSPA parameters $p = 2$, $c = 40$, and $\gamma = 5$. Shown, from the top, is the combined GOSPA-T, the localization error, the missed detection error, the false alarm error, and the track switch penalty. For visual clarity the values are plotted as the moving average over 10 time steps.

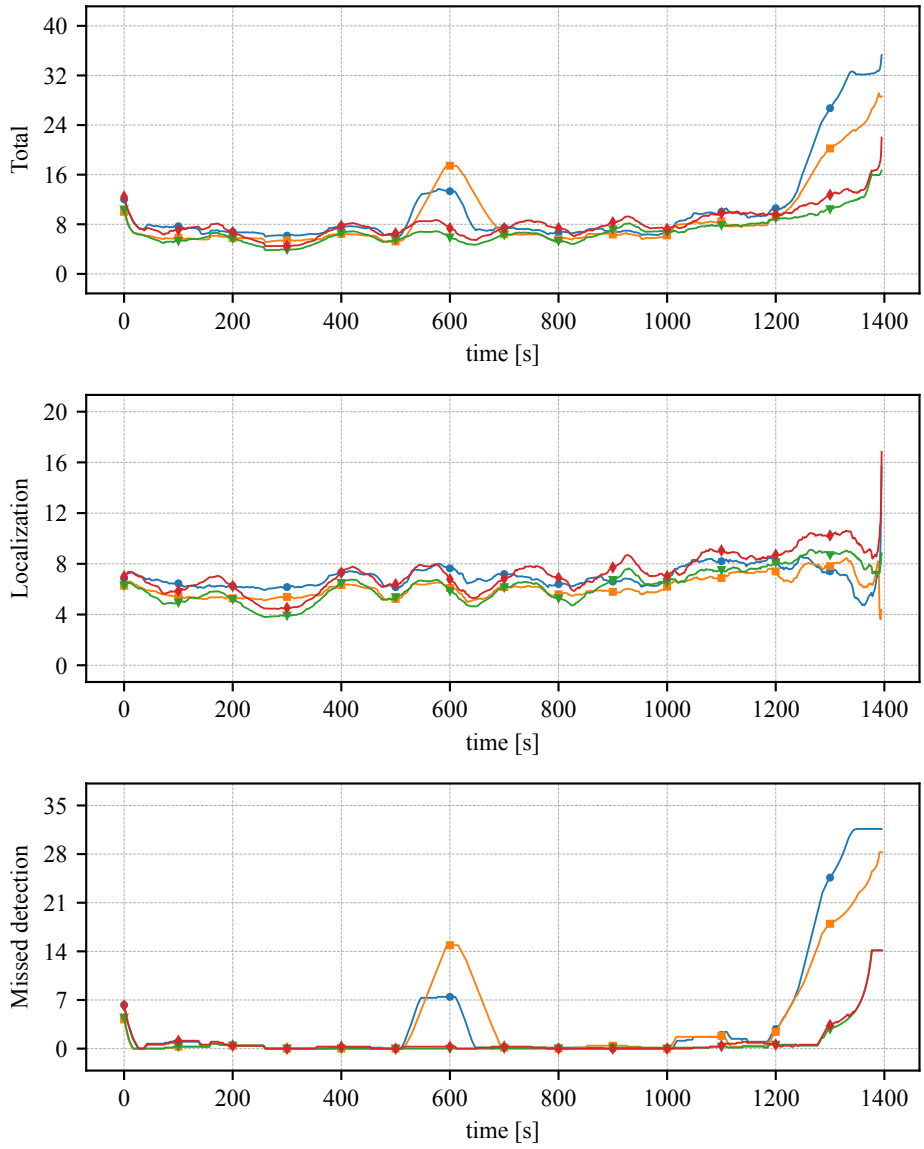

Figure 4.2: Figure with legend continues on next page.

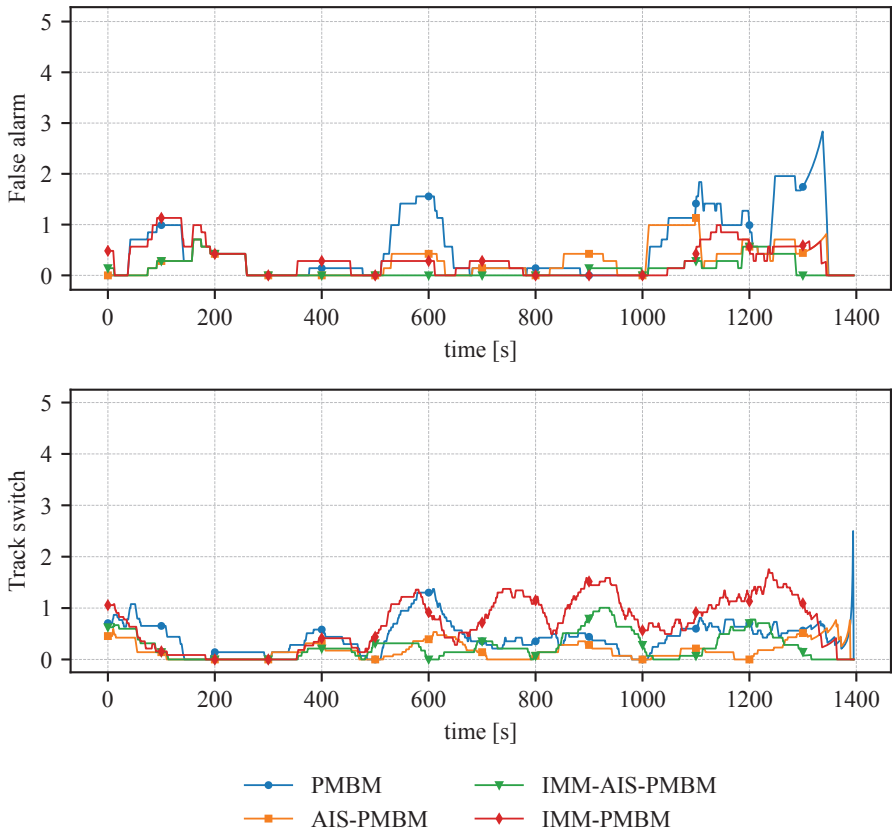

Figure 4.2: GOSPA-T values for the different PMBM tracker configurations. We use the GOSPA parameters $p = 2$, $c = 20$, and $\gamma = 5$. Shown, from the top, is the combined GOSPA-T, the missed detection error, the false alarm error, the localization error, and the track switch error. The values are from the whole duration of the Ravens data set, and for visual clarity plotted values are the moving average over 100 time steps.

Table 4.3: GOSPA-T values for the different PMBM configurations when tested on the Ravens data set, corresponding to the averaged values in Figure 4.2. The best result for each error source is highlighted in bold.

	PMBM	AIS-PMBM	IMM-PMBM	AIS-IMM-PMBM
Total GOSPA	10.149	8.859	8.151	**6.845**
Localization	6.956	**6.145**	7.280	6.274
Missed detection	3.853	3.556	0.861	**0.670**
False alarm	0.555	0.279	0.292	**0.123**
Track switch	0.446	**0.175**	0.687	0.246

worse after. At the very end of the scenario this can be explained with the non-IMM trackers losing track of their targets, and thus the localization error of the more difficult targets is not included. This does not, however, explain why this also occurs before the track losses happen. An explanation can be that we experience a trade-off between precise estimation of individual targets and being able to track all targets. That is, to be able to handle large changes in acceleration the added covariance in the prediction is often too large when the targets do not perform challenging maneuvers. Even though the IMM framework should be able to account for this, it may struggle to estimate the mode probabilities correctly due to the challenging scenario. Nevertheless, the differences between methods are not large, and the IMM methods perform better by the total GOPSA metric. Furthermore, this indicates that looking only at single components of the GOSPA metric will not give a complete picture of the performance of the trackers.

4.6.2 Experimental validation in a closed-loop experiment

The AIS-IMM-PMBM was also used during the testing of an autonomous surface vehicle (ASV) in the Trondheimsfjord in October 2023. A more extensive description of the testing and additional scenarios is available at [70]. In the context of this chapter, the purpose of using the AIS-IMM-PMBM method in the experiment was to demonstrate the feasibility of using an advanced tracker in a larger system with higher demands regarding latency, and to show that it works together with a collision avoidance system. The ASV was equipped with a radar and an AIS receiver. The collision avoidance algorithm, a scenario-based model predictive control method described in [86], was responsible for maneuvering the ASV to avoid collisions with the targets based on input from the AIS-IMM-PMBM.

The scenario involves the ownship, a Mariner from Maritime Robotics, and Juggernaut, a motorboat. The boats are shown in Figure 4.3. The motorboat was operated by a human operator, and the ownship was tasked with avoiding collisions

with the target while maintaining a course towards a waypoint. The scenario was designed to challenge the collision avoidance algorithm, and for it to properly respond to the movements of the target it needed accurate target estimates. What amounts to accurate estimates in this use case is somewhat different from what is usually considered when evaluating target trackers. The course and speed estimates are very important as they are used to predict the future positions of the targets. The position of the target, however, is not as important because the collision avoidance algorithm does not mainly act based on the current target position, but rather on the predicted future positions. An inaccurate course estimate will then have a larger impact than an inaccurate position estimate.

(a) (b)

Figure 4.3: The two ships used in the experimental validation: Mariner (a) and Juggernaut (b). Photos: Maritime Robotics.

Results

The scenario is depicted in Figure 4.4. The Mariner was tasked with moving on a north-eastern course while the motorboat made several maneuvers that demanded action from the Mariner. The target deliberately acted counter to collision avoidance guidelines, so that the collision avoidance algorithm was forced to make evasive maneuvers. Nevertheless, the AIS-IMM-PMBM tracker was able to provide accurate estimates of the target, and the collision avoidance algorithm was able to respond to the maneuvers. The course and speed estimates are stable and with relatively low uncertainty. Jittery and uncertain estimates would demand greater caution on the part of the collision avoidance algorithm, and it could create situations where the ownship path would have to be recalculated at a higher frequency than necessary. Furthermore, any false alarms were avoided, partly helped by the low clutter density.

The data collected during the scenario was also input to an IPDA tracker, see Figure 4.5. This was the tracker used in [85], and whereas it is able to track the target its course and speed estimates are considerably noisier than for the AIS-IMM-PMBM.

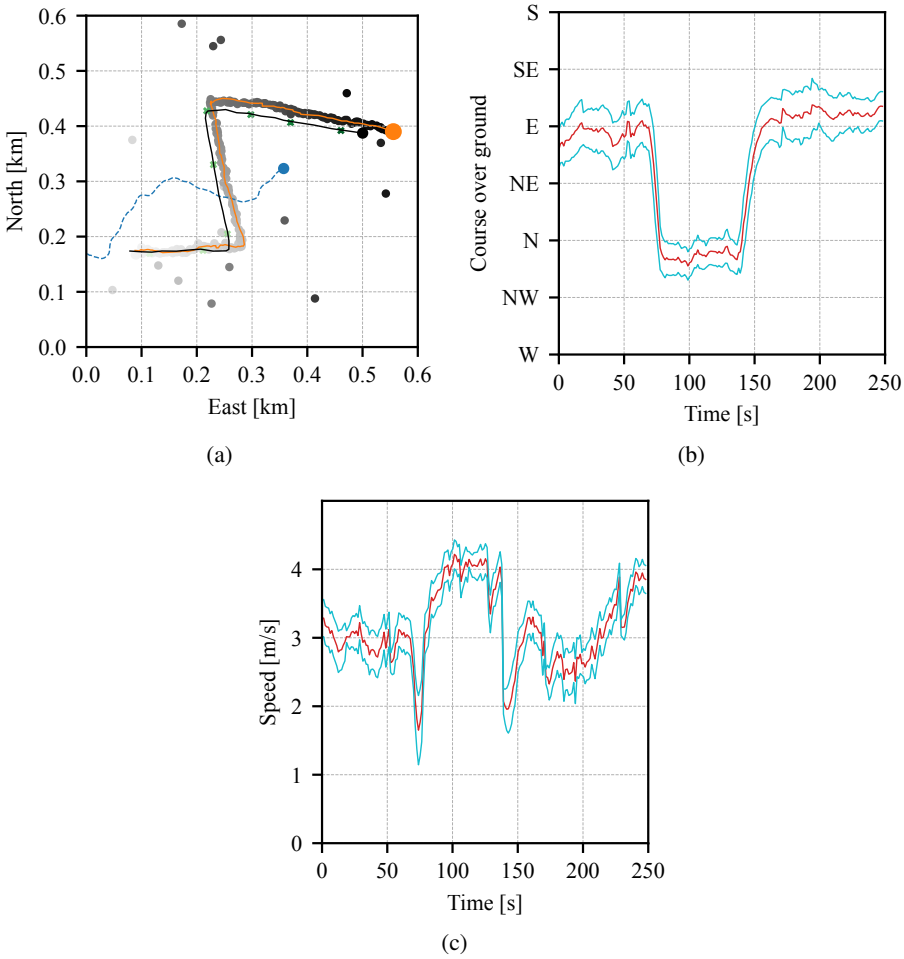

Figure 4.4: A scenario with a target conducting sharp and unpredictable maneuvers, showing how the tracker is able to successfully estimate the target state in a real-world setting. The ownship is shown as the dotted blue line in (a) and tries to move in a northeastern direction while upholding collision avoidance regulations. The target is shown as a solid orange line and tries to make maneuvers which hinder the Mariner in its objective. The radar measurements are shown as black dots which fade to grey as time passes, the AIS measurements are shown as green crosses which also fade as time passes, and the reported GPS position of the target is shown as a solid black line. (b) shows the evolution of the course estimates as a solid red line, together with one standard deviation in each direction as solid cyan lines. The speed estimates are shown in (c) in the same way, with the estimate as a red line and the standard deviations in cyan.

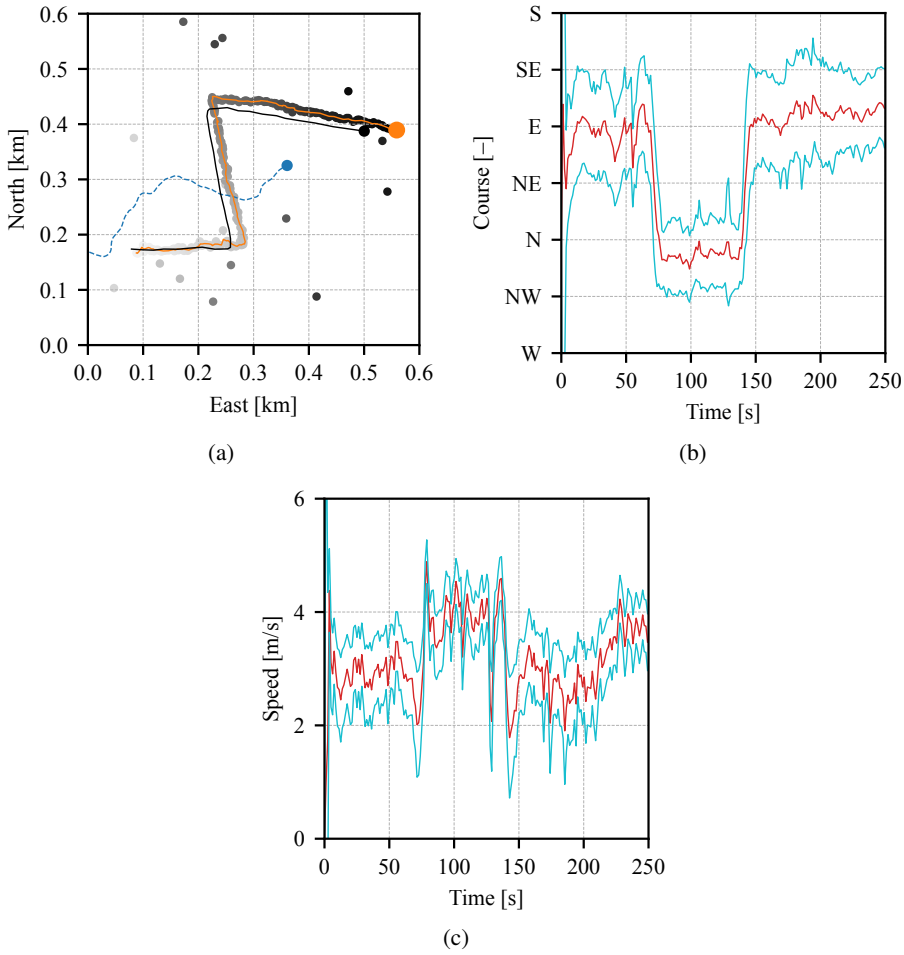

Figure 4.5: The same scenario as in Figure 4.4, but now an IPDA is used for tracking. an overview of the scenario is shown in (a), the estimated course is shown in (b), and the estimated speed is shown in (c). The tracker is able to follow the target, but the course and speed estimates are not as precise.

The standard deviations are also larger: for the IPDA the standard deviations are 0.36 knots for the speed and 34.45 degrees for the course, whereas they are 0.25 knots and 15.63 degrees for the speed and course estimated by the PMBM. Because the IPDA uses only a single kinematic model, and furthermore does not use the AIS messages, this is to be expected. However, as the collision avoidance method was not used together with the IPDA, it is difficult to say how this would impact operations. Previous work on use of PDA in radar-based maritime collision avoidance [36] indicates that additional filtering of the speed and course estimates is needed for successful operation. Furthermore, the IPDA is well suited for such a single-target scenario but lacks the flexibility of a more advanced method regarding multiple targets with different movement characteristics.

4.7 Results for extended object tracking

For the AIS-GP-PMBM we also use GOSPA for trajectories to evaluate the performance with regards to the position and velocities of the target estimates. Additionally, we need a method to evaluate the extent estimates. For this purpose, we use the intersection-over-union (IOU) between the estimated and true extents of the targets. IOU is found by taking the true extent of the target as the area $\mathcal{E}$ and the estimated extent as the area $\hat{\mathcal{E}}$, which we then use to calculate

$$IOU = \frac{\mathcal{E} \cap \hat{\mathcal{E}}}{\mathcal{E} \cup \hat{\mathcal{E}}}. \tag{4.76}$$

To pair the correct estimate with the correct ground truth, we use the same assignment procedure that is used when calculating the GOSPA metric. We consider two scenarios that highlight two key challenges that make estimation based on exteroceptive measurements difficult: occlusion and clutter measurements. To properly handle occluded targets is difficult, and the lack of measurements will in most cases result in reduced estimate quality. The clutter model we use is quite simple and assumes that the clutter measurements are uniformly distributed among the sensor cells. As the experimental data demonstrates, this is not necessarily the case, and especially target-dependent clutter, such as that originating from wakes, is difficult to handle. Both these problems are in some ways mitigated when performing point target tracking, as the sensor detections are clustered before they are input to the tracker. This means that partial occlusions will not necessarily result in loss of measurements, and that target-dependent clutter can be clustered together with target detections. For extended object tracking, however, we do not have these advantages.

The tuning parameters for both scenarios are shown in Table 4.4.

Table 4.4: Tuning parameters for the extended object scenarios. If the parameters differ between the two scenarios, they are listed with the parameter for the first scenario first, and for the second scenario last. Any parameters not listed here are identical to ones found in [6].

Quantity	Symbol Unit	Value
Survival probability	P_{S} $[-]$	0.999
Gate size	g $[-]$	20
Max. number of hypotheses	$N_{\max}$ $[-]$	20
Poisson pruning threshold	T_{Pp} $[-]$	1.0×10^{-5}, 1.0×10^{-2}
Bernoulli pruning threshold	T_{Bp} $[-]$	1.0×10^{-5}, 1.0×10^{-2}
Area radius	R $[\mathrm{m}]$	$100, 60$
Clutter rate	N_c $[-]$	$16, 80$
Initial velocity std.	σ_v $[\mathrm{m/s}]$	3
Initial heading std.	σ_ψ $[^\circ]$	180
Initial ang. vel. std.	σ_ω $[^\circ/\mathrm{s}]$	45
Existence confirmation threshold	T_c $[-]$	0.99
Motion noise intensity	q_a $[\mathrm{m}^2/\mathrm{s}^3]$	0.2^2
Heading noise intensity	q_θ $[^\circ]$	$0.1^2, 0.3^2$
Number of birth components	N_b $[-]$	36
Birth weight	U $[1/\mathrm{m}^2]$	$1/36$
Forgetting factor	η_γ $[-]$	0.99
Gamma distribution shape parameters	α, β $[-]$	$[1000\ 100]$, $[500\ 100]$
Lidar measurements		
Detection probability	P_D $[-]$	0.9
Cartesian noise std.	σ_r $[\mathrm{m}]$	$0.1, 0.5$
AIS measurements		
Cartesian noise std.	σ_{c_R} $[\mathrm{m}]$	1.0
Polar range std.	σ_r $[\mathrm{m}]$	1.0
Polar bearing std.	σ_θ $[^\circ]$	10
Confidence probability	P_C $[-]$	0.999
Unknown target no ID probability	ξ_u^0 $[-]$	0.9
Length and width std.	$\sigma_{l,w}$ $[\mathrm{m}]$	0.5

Table 4.5: Values for the performance metrics in Figure 4.6 and Figure 4.7. The values are averaged over the whole duration of the scenarios, and the best values for each metric are highlighted in bold.

	Simulated data		Experimental data	
	GP-PMBM	AIS-GP-PMBM	GP-PMBM	AIS-GP-PMBM
Total GOSPA	2.564	**1.420**	8.684	**4.789**
Localization	2.154	**1.091**	4.577	**3.693**
IOU	0.600	**0.765**	0.192	**0.263**

4.7.1 Simulated data

The simulated data used to evaluate the EOT methods is the same as that used in [6]. It consists of four targets, with a LiDAR located at the center of a surveillance area with a 100-meter radius. As for the Ravens data set, AIS measurements were created for the targets. Here, all targets transmit AIS messages, which also include their width and length. All the targets are 6.63 meters long and 2.4 meters wide, the scenario lasts for 240 seconds, and the LiDAR measurements generate scans of the area at 1 Hz.

Results

Figure 4.6 and Table 4.5 show the results. The targets enter and depart the surveillance area at different times, as can be seen by short spikes in the GOSPA values. These spikes correspond to missed detections when they arrive and false alarms when they leave, due to latency in the initialization and termination of tracks. Regardless, for most of the scenario the trackers are able to both track and estimate the extents of the targets quite successfully, with low localization error and high IOU. Use of AIS measurements both decreases the localization error and increases the IOU. This is shown throughout the whole scenario but is most evident before the 150 second mark. Here, occlusion effects result in low LiDAR measurement quality, which in turn makes it difficult to estimate the target states. AIS measurements, however, are not impacted by occlusion and allows the tracker to recuperate after the intermediate LiDAR scans. This effect becomes more evident because the AIS messages are transmitted at approximately the same time, resulting in noticeable spikes in the performance metrics also when averaging across all targets.

4.7.2 Experimental data

The experimental data used to test the AIS-GP-PMBM were collected in the Trondheim channel as part of the Autoferry project at NTNU, and was published in [62] as scenario 13. It is also discussed in the context of the GP-PMBM in [6]. The data were collected using a LiDAR mounted on the Milliampere ferry, which was stationary, with a range of approximately 60 meters. Furthermore, the scenario contains two targets, both 7 meters long and 3 meters wide motorboats. They travel across the length of the canal in opposing directions, passing each other approximately when closest to Milliampere. The targets did not transmit AIS messages, but GPS positions are available, which we used to create AIS measurements in the same manner as previously.

Results

The results can be seen in Figure 4.7 and Table 4.5. The overall performance is not as good as for the simulated data set. Because we are no longer dealing with simulated measurements that adhere to our modeling assumptions, this is to be expected. The ground truth also contains a bias which skews the results somewhat, but it is nevertheless useful for comparison purposes. Furthermore, the data set contains a lot of clutter, both from the ship wakes and other sources. The clutter model may struggle especially when encountering wake clutter. When evaluating the measurements in an association hypothesis, wake clutter will often be included and given a high weight. This leads to the extent estimate growing larger than it should, and the kinematic estimate being pulled towards the wake. Closely spaced targets which generate wake clutter can exacerbate performance degradation further by associating the wake clutter from one target with the other, and the tracker will figuratively try pulling the extent estimate apart. Some variation of this effect is seen in this data set, with the tracker struggling at the midpoint of the scenario when the targets are closest to each other.

Initially, both variations of the GP-PMBM are able to track the targets well, albeit with some false alarms for the pure LiDAR tracker. As the targets get closer to each other and the LiDAR, the performance degrades. The AIS-GP-PMBM is able to handle the situation better, helped by the AIS messages that both provide good kinematic information and the dimensions of the ship. In a situation where the LiDAR measurements cause the extent estimate to blow up, the AIS measurements help reduce its size. The pure LiDAR tracker, however, struggles to maintain a good extent estimate. This is seen clearly in the plots of the width and length, and in Figure 4.8. Furthermore, we see that the oversized extent of the non-AIS variant creates an offset in the position estimate, and that the direction of the estimate is flipped. At the displayed time step this is avoided when including AIS information,

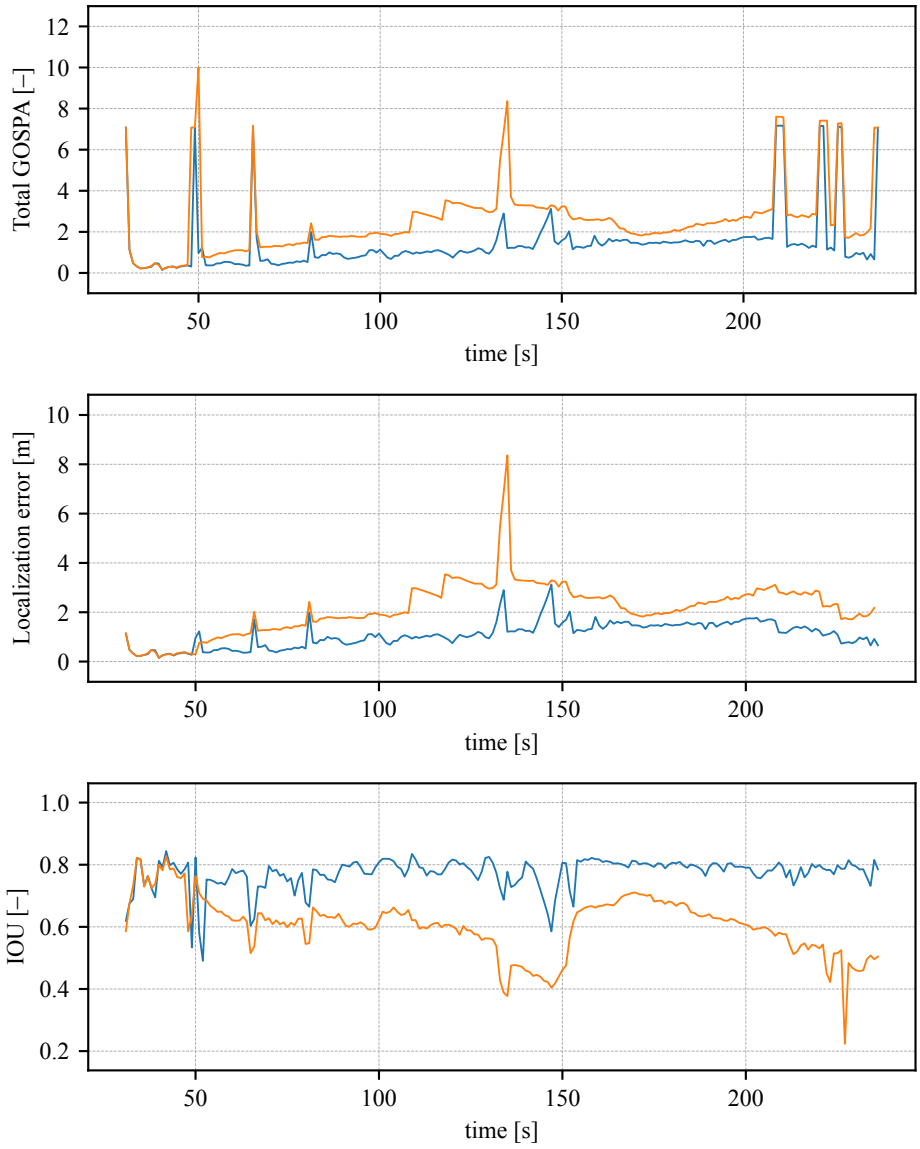

Figure 4.6: Figure with legend continues on next page.

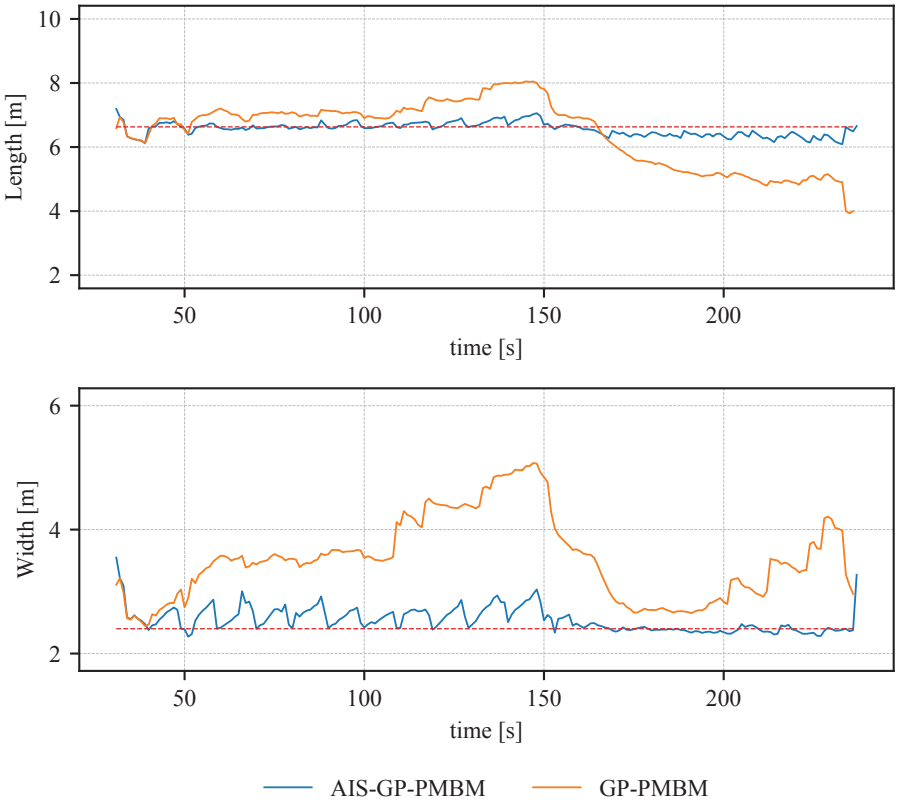

Figure 4.6: Comparison between the GP-PMBM and the AIS-GP-PMBM for the simulated data set. We use the GOSPA parameters $p = 2$, $c = 10$, and $\gamma = 5$. From top to bottom, the plots show the total GOSPA-T, the localization error, the IOU, the averaged estimated lengths, and the averaged estimated widths.

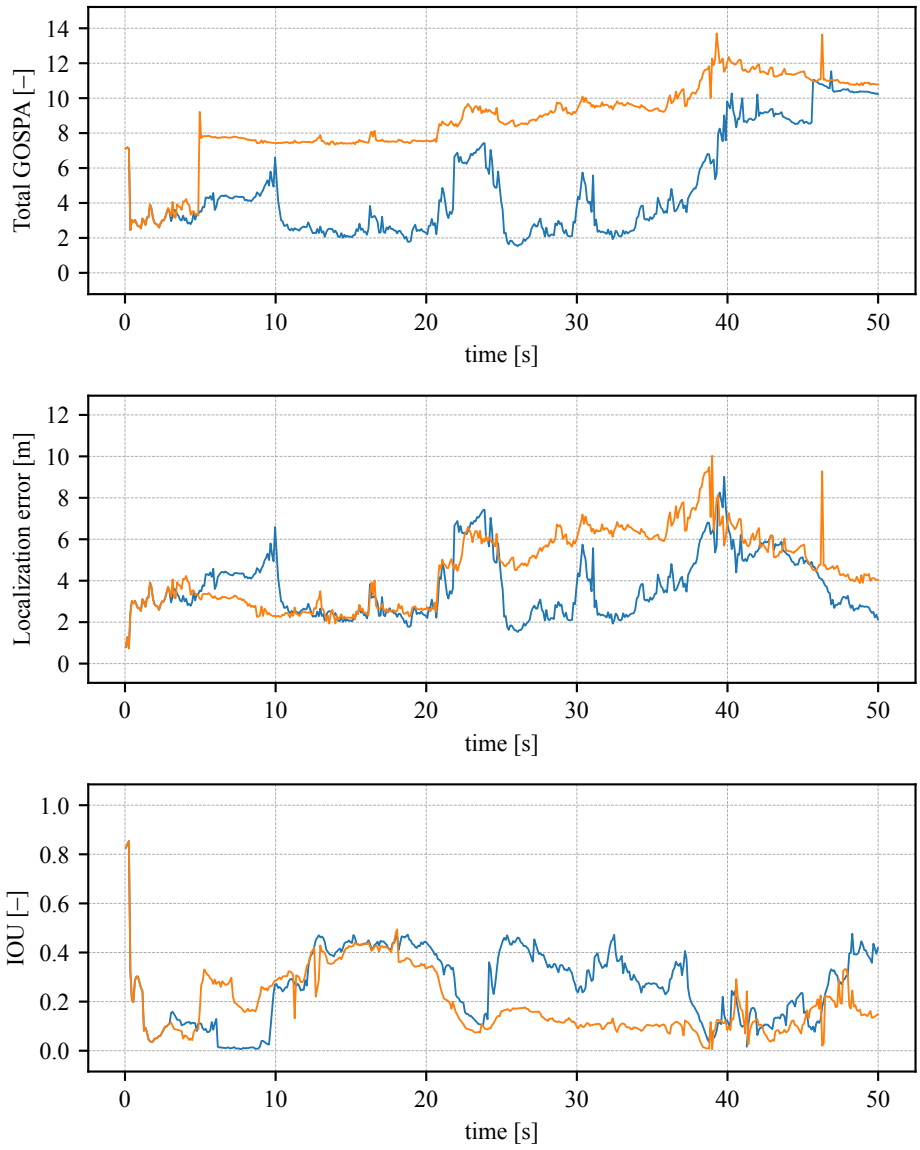

Figure 4.7: Figure with legend continues on next page.

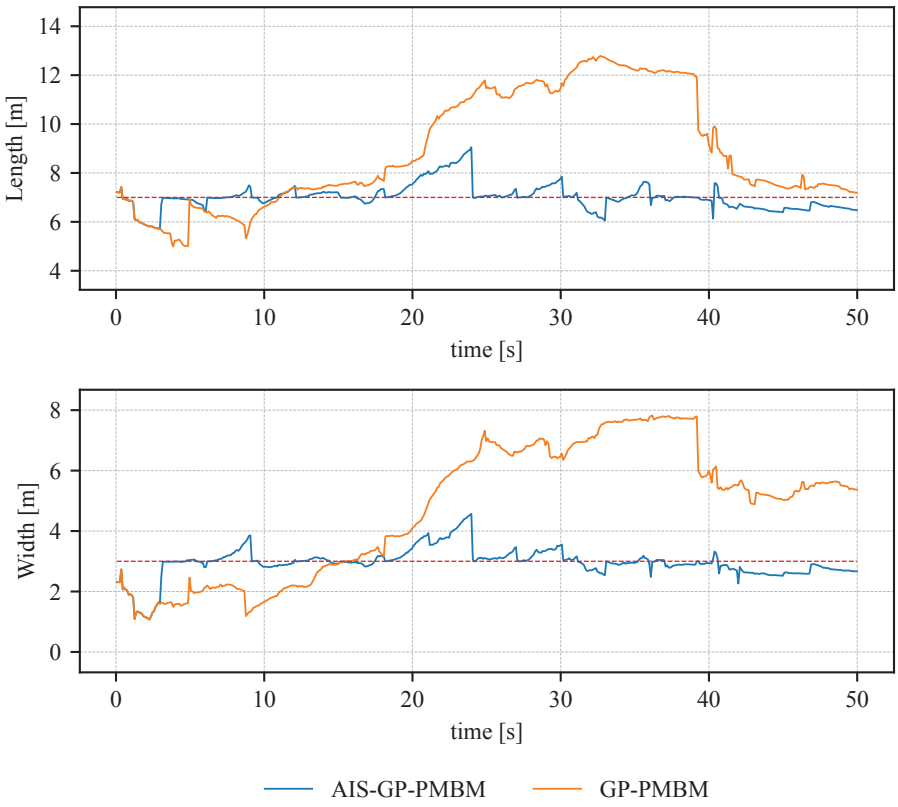

Figure 4.7: Comparison between the GP-PMBM and the AIS-GP-PMBM for the experimental data set. We use the GOSPA parameters $p = 2$, $c = 10$, and $\gamma = 5$. From top to bottom, the plots show the total GOSPA-T, the localization error, the IOU, the averaged estimated lengths, and the averaged estimated widths.

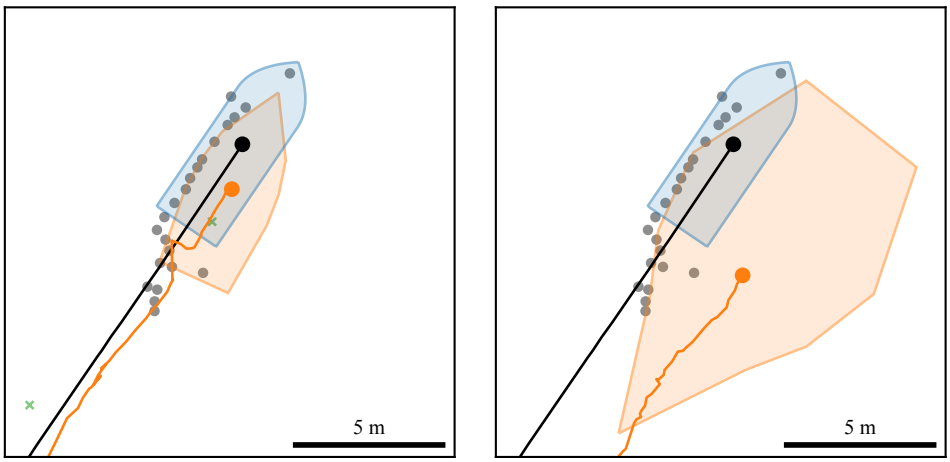

(a) (b)

Figure 4.8: A snapshot of the extent of one of the targets 30 seconds into the scenario, at time step 283. The output from the AIS-GP-PMBM is shown in (a), and that of the GP-PMBM in (b). The estimated extent is shown in orange, and the ground truth in blue. Additional elements are explained in Figure 4.4. When not using AIS information, the extent estimate is both blown up and flipped orientation-wise.

even though the heading is not directly updated with AIS information. Furthermore, false alarms are present for both trackers. For the pure LiDAR, tracker the false alarms are persistent throughout the whole scenario, whereas they arrive later in the scenario for the AIS-GP-PMBM. With a lot of clutter, and especially with wake clutter, this is difficult to mitigate without more advanced clutter models.

4.8 Chapter summary

We have presented a method for including target-provided information in PMBM trackers, with AIS messages as an example of such information. We do this for PMBM and its IMM-PMBM extension for point target tracking, and the GP-PMBM for extended objects. Through simulations and experiments we have shown that the AIS-PMBM and AIS-IMM-PMBM trackers perform better than their counterparts which do not use AIS measurements. Furthermore, we have compared the presented methods to other trackers which utilize AIS messages and shown that they perform well in comparison.

Because one of our use cases for the AIS-IMM-PMBM tracker is to provide target estimates for a collision avoidance system, we have also tested them in a closed-loop experiment. The results show that the AIS-IMM-PMBM tracker is able

to provide good estimates and enables the collision avoidance method to respond to the target movements. The experiments also show the feasibility of using PMBM in applications which demand real-time performance.

Furthermore, we show how use of target-provided information in extended object tracking can mitigate some of the persistent problems when using only exteroceptive sensors. We use AIS, together with LiDAR, and use information regarding the ship dimensions to help estimate the extent. By updating the extent with the ship dimensions, we avoid that the extent changes its size when the LiDAR measurement quality is poor. We show that this can improve performance when a lot of clutter measurements are present, and that it helps the tracker recover from occlusion effects.

There are some potential lines of future research. Firstly, AIS messages can be utilized more extensively to improve target estimates. For example, they contain information about course, heading, and antenna placement, which can be useful in several ways. Furthermore, better wake clutter modeling can solve some of the problems encountered in Section 4.7.2. Several works present more general clutter models [48, 133], and also target-dependent clutter models for wake clutter mitigation [19, 64]. Such models could be utilized to improve the performance of the AIS-GP-PMBM in the presence of wake clutter.

Acknowledgment

The collision avoidance experiment was conducted in co-operation with Maritime Robotics, who provided personnel and infrastructure.

Part II

Use of AIS in Maritime Autonomy Applications

Chapter 5

Autonomous Marine Collision Avoidance With Sensor Fusion of AIS and Radar

This chapter is based on the publication:

[70] A. G. Hem, E. F. Brekke, G. D. K. M. Kufoalor, and I. H. Kingman. Autonomous Marine Collision Avoidance With Sensor Fusion of AIS and Radar. Submitted to the 15th IFAC Conference on Control Applications in Marine Systems, Robotics and Vehicles (CAMS 2024)

Changes from the submitted publication involve shortening of the introduction section.

This paper is submitted for publication and is therefore not included.

Chapter 6

Validation of AIS Information With Exteroceptive Sensor Fusion in Autonomous Operations

This chapter is based on the publication:

[66] A. G. Hem and E. F. Brekke. Validation of AIS Information With Exteroceptive Sensor Fusion in Autonomous Operations. To appear in *IEEE Intelligent Transportation Systems Magazine*. doi: 10.1109/MITS.2024.3389869

Changes from the submitted publication involve shortening of the introduction section.

6.1 Introduction

To use wrong, or even deliberately falsified, information in an autonomous system is obviously problematic. It can cause degraded situational awareness, which could lead to wrong decisions and dangerous situations. Thus, it is important to be able to detect and handle errors in the incoming Automatic identification system (AIS) messages. The assumption that the vessel receiving the AIS messages is also equipped with additional sensors, provides us with the possibility of using a solution that utilizes exteroceptive sensor data to check the AIS information. The unreliable nature of AIS has led to the development of many methods for detecting and handling errors. A recent survey [144] considers different methods for detecting anomalies in AIS tracks. As we want to allow for the use of AIS messages in combination with exteroceptive sensors, the methods most relevant for our case are the ones that rely on such sensors. Bloisi et al. [13] considers fusing of camera data and AIS information for situations

where radar is unavailable. We, however, are concerned with situations where we can use radar. Katsilieris et al. [82] and d'Afflisio et al. [32], for example, detect false AIS messages using radar data. These methods use statistical hypothesis testing and are useable regardless of the availability of historical data. Their objective is to check if AIS messages provide incorrect positional information to detect spoofing. However, [82] and [32] only consider single targets and focus on classifying the spoofing scenario. This differs from our use case, where we want to validate correct data from multiple targets and discard wrong data without considering it further. We also need methods that consider errors in additional AIS message fields such as speed and course.

This leads to the contributions of this chapter. We present a methodology for validating AIS information for use together with exteroceptive sensors in maritime autonomy applications. The validation method relies on trackers which fuse AIS data and exteroceptive sensor data. We present methods for detecting errors in the transmitted position, speed, course, and rate of turn. This allows for use of AIS messages in target tracking, long-term prediction, and collision avoidance without risking consequences from erroneous messages. We use simulated data to determine the efficiency of the detector and demonstrate the method on a real-world dataset where false AIS data has been injected.

This chapter is structured as follows. Section 6.2 presents relevant details regarding AIS. That is, what parts of the available information we need to consider, target tracking methods that use AIS, and reliability concerns. Section 6.3 presents the statistical tests used to detect any errors in the relevant AIS message fields. Section 6.4 presents methods for validating AIS messages and outlines how these can be used with a target tracker. Section 6.5 presents the test setup and results, before Section 6.6 concludes.

6.2 Background

6.2.1 The Automatic Identification System (AIS)

In 2004, the International Maritime Organization (IMO) mandated the use of AIS to increase situational awareness for maritime vessels and improve their exchange of information. Vessels using AIS are categorized either as Class A or Class B vessels. Class A mainly applies to commercial ships, whereas most other vessels are class B.

There are many different AIS messages, but we are concerned with those providing dynamic information. These are the AIS messages with ID numbers 1, 2, or 3 for Class A vessels, and ID 18 for Class B vessels. The messages contain roughly the same information, except that the Class B dynamic messages omit the turn rate and are not required to transmit course information. Furthermore, in this

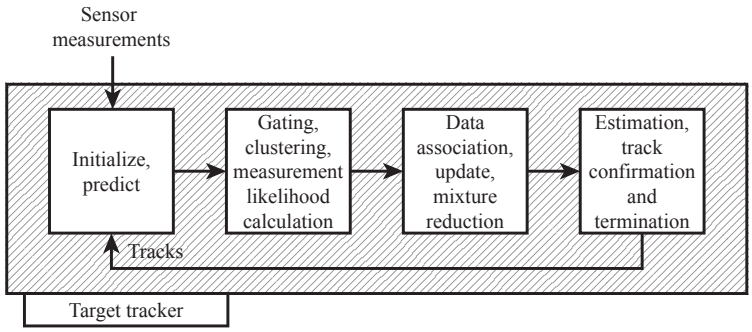

Figure 6.1: A general target tracker structure.

work, we are only interested in position, MMSI number, speed, course over ground, and rate of turn. A thorough technical description of the AIS messages, and the system as a whole, is found in [73].

6.2.2 A brief introduction to target tracking

To explain some of the concepts used in the following sections, we will briefly introduce a generic target tracker. Figure 6.1 shows a block diagram of a tracker structure.

We initialize new tracks on the incoming sensor measurements in the first block in Figure 6.1. The new tracks represent the hypothesis that the measurements come from previously undetected targets. Furthermore, the already existing tracks are predicted to the point in time when the measurements arrive. This prediction is usually performed with some filtering method, such as a Kalman filter or a particle filter, together with a kinematic model. Furthermore, a kinematic model is used to predict the state of the new tracks. Two widely used models are the constant velocity (CV) [8, p. 274] model, and the coordinated turn (CT) [8, p. 467] model.

The second block receives the new and predicted tracks from the first block. Here we perform gating, and then clustering, on the predicted tracks to reduce computational loads. Gating is a method where all measurements unlikely to be associated with a given track are ignored, by only considering measurements in the vicinity of the track. We also calculate the measurement likelihood for each track and measurement pair within the clusters.

Next, we perform data association by using the measurement likelihoods. The specific method depends on the tracker. If a probabilistic data association (PDA) filter [9] is used, each possible association between a track and a measurement is assigned a weight used when updating the track. The track-to-measurement weight decides how much each measurement contributes to the update when the

different hypotheses are merged into a single estimate. Other methods, such as multiple hypothesis tracking (MHT) [115], calculate weights to decide the most likely track-to-measurement associations.

Lastly, we estimate the target states. Here the most likely state of each track is calculated for use in other parts of the system. Furthermore, some methods estimate the existence probabilities of each track. The existence probability quantifies whether the potential targets the individual tracks follow are likely to exist or not. Termination or confirmation of tracks is done by, for example, evaluating the existence probabilities.

For the problem of multi-sensor target tracking, we can divide the solutions into track-to-measurement and track-to-track methods. The track-to-measurement methods fuse the data on a measurement level, whereas the track-to-track methods establish tracks on measurements from individual sensors before fusing the estimates.

Recently, several target tracking methods which incorporate AIS messages have been developed. Both [44] and Chapter 3 present trackers which fuse radar measurements and AIS messages on a measurement level. [44] uses loopy belief propagation with a particle filter, and Chapter 3 uses a joint integrated probabilistic data association (JIPDA) [105] tracker and interacting multiple models (IMM) [15] with an extended Kalman filter. IMM allows us to use more than one kinematic model to describe the movement of the targets. The tracker in Chapter 3 is the one we use to test the methods presented in this chapter.

6.2.3 Reliability of AIS messages

The AIS protocol detects errors from bit flips during transmission with cyclic redundancy checks and drops erroneous messages. Other errors, however, are hard to catch. Thus, the problem of potentially erroneous AIS messages has been the subject of significant research efforts. In [60], Harati-Mokhtari et al. investigate both static and dynamic information transmitted in AIS messages and conclude that the AIS messages, in general, can not be fully trusted. Ray et al. [114] present the risks associated with the AIS protocol and establish a methodology with which to develop error-mitigating techniques. In [75], Iphar et al. build upon the work in [114] and investigate the reliability issues in AIS. They perform a risk analysis and identify and classify several problems which can arise with falsified AIS data.

A more general analysis of cybersecurity challenges at sea is presented in [1]. Here, in addition to the potential for manipulated AIS messages, the problem of the global navigation satellite system (GNSS) being spoofed or jammed is pointed out. As the AIS system relies on GNSS for positioning, GNSS spoofing would also affect the transmitted AIS messages. These limitations have made the prospect of an improved AIS protocol tempting. The new VHF Data Exchange System (VDES) [71], supported by IMO, is thought to serve as a second-generation AIS.

It addresses many of the problems with the current AIS protocol but is not widely adopted yet.

6.3 Detection of AIS message errors

The fields in the AIS message have a limited range of correct values. Thus, it is possible to check if the received data lie outside the permitted intervals before performing more sophisticated tests. Table 6.1 summarizes the possible values for the relevant fields, and we discard any outside of these ranges as a preliminary step.

Table 6.1: Intervals for the AIS message fields providing useable data, from [73], and information regarding conversion from the bit values to usable values.

AIS message field	Interval	Conversion
MMSI number	See [74]	Becomes a nine-digit decimal number
Longitude	0000000_{16}-$66FF300_{16}$	An additional bit allows for using two's complement, with west being positive and east negative. Each step is 1/10000 minute.
Latitude	0000000_{16}-$337F980_{16}$	An additional bit allows for using two's complement, with north being positive and south negative. Each step is 1/10000 minute.
Course over ground	0 - 3599	Each step is 0.1 degrees
Speed over ground	0 - 1022	Each step is 0.1 knots
Rate of turn	0 - 256	The field value is calculated as $4.733\sqrt{\lvert\omega\rvert}$ for turn rate ω. A value between 0 and $+126$ or 0 and -126 indicates that the vessel is turning right or left, respectively. Similarly, a value of $+127$ or -127 indicates that the vessel is turning right or left at more than 5° per 30 seconds.

6.3.1 Error causes

We identify the undesirable events relevant to our use case from [75, Table 5] and summarize them in Table 6.2. These are the events that can feed the system incorrect information or the ones that remove part of the system input. Of these two options, the first is the most problematic. Whereas many errors can result in AIS messages

Table 6.2: The relevant events from [75, Table 5].

Problem	Consequence
Dynamic data from a target are false. AIS data reception is not reliable.	The system can be fed incorrect information, or messages can be dropped.
Dynamic data from a target are not available. AIS data reception is not working.	The system must operate using only exteroceptive sensors.

not being received, such as faulty antennas or bit flips leading to dropped messages, these situations are only a minor concern. If the tracker misses part of the AIS input, it can still rely on the exteroceptive sensors. However, if the tracker uses false AIS information in computations, this will likely degrade the estimates found using the exteroceptive sensors.

6.3.2 Error types

Before designing the error detection methods, we classify the types of errors we can expect to encounter. There are five fields in which errors can occur, and we assume that the speed and course is converted to a velocity vector before its use in a tracker. Because of this, we consider errors in velocity instead of course and speed separately. Furthermore, we also distinguish between what we denote as the primary and the additional information provided by the AIS messages. The primary information provided by AIS messages is the position and the MMSI number, whereas the additional information is the velocity and rate of turn.

The primary information is used for data association and is generally not easily discarded, as opposed to the additional information. Position errors in the AIS messages can lead to false tracks, and consecutive errors are likely to result in a false track. Figure 6.2b shows this, with a new track initialized on a sequence of false AIS messages. Furthermore, we assume that the underlying target tracker is designed to process the MMSI numbers probabilistically, so any errors in the MMSI information will be within the intended design of the tracker. This is the case for the trackers described in both [44] and Chapter 3.

Additional information is used mainly to improve the target estimates.

Errors in the course and speed can result in wrong velocity estimates, and errors in the turn rate can result in wrong angular velocity estimates. This can give inaccurate predictions, and in the worst case, track loss. Both sporadic and consecutive errors can result in these problems, with consecutive errors posing the

larger threat. Figure 6.2c shows an example of a scenario where consecutive errors in the course and speed result in track loss.

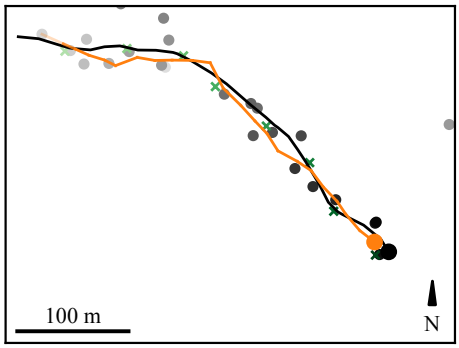

(a) A track which follows a target with correct AIS messages originating from an actual target.

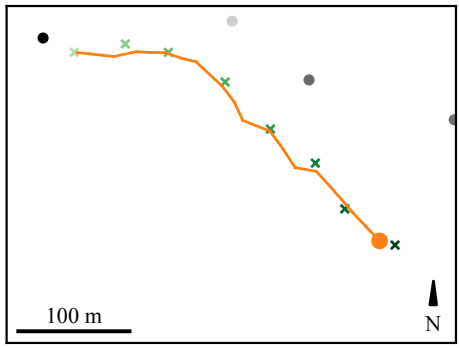

(b) A track established on AIS messages not corresponding to an actual target.

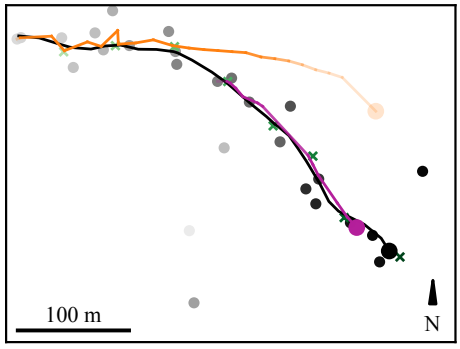

(c) A track being initialized on AIS messages reporting zero-valued speed, before losing the target and having to be re-established.

Figure 6.2: Tracking of a single target using the tracker from Chapter 3. The tracker input contains AIS messages (green crosses) and radar measurements (gray circles). The targets are black lines, with the dot signifying the end position. The tracks are colored lines ending in a dot, with their opacity corresponding to the estimated existence probability. As can be seen, erroneous AIS messages can make the tracker output false tracks or degrade the tracking performance.

6.3.3 Hypothesis test for validating velocity and rate of turn information

We use a hypothesis test to validate the velocity information in the AIS messages. The two hypotheses are as follows:

- H_0: The velocity measurement comes from the distribution $\mathcal{N}(\hat{\mathbf{x}}_v, \hat{\mathbf{P}}_v)$ where $\hat{\mathbf{x}}_v$ is the estimated velocity mean and $\hat{\mathbf{P}}_v$ is the estimated velocity covariance from the target submitting the AIS message.
- H_1: The velocity measurement does not come from the above distribution.

Furthermore, we use a test statistic describing the available data to evaluate the above hypotheses. We define it as

$$\zeta_v^{tj} = (\mathbf{z}_v - \hat{\mathbf{x}}_v)^\top \hat{\mathbf{P}}_v^{-1} (\mathbf{z}_v - \hat{\mathbf{x}}_v) \tag{6.1}$$

which is the squared Mahalanobis distance between the measurement $\mathbf{z}_v$ and the estimate $\hat{\mathbf{x}}_v$. ζ_v^{tj} is distributed according to a χ^2 distribution with the degrees of freedom equal to the velocity dimension. We do not necessarily know which target that has transmitted the AIS message. Thus, we must perform data association to evaluate the hypothesis. The data association can be done in several ways, for example, by using the measurement likelihoods as weights. We use the normalized product of the measurement likelihoods l^{tj} and the existence probabilities r^t as weights, which in JIPDA-style trackers are readily available. The weights become

$$w^{tj} = \frac{l^{tj} r^t}{\sum_t l^{tj} r^t}. \tag{6.2}$$

Using this, we find the final test statistic as a weighted average over the distances between the measurement and the tracks as

$$\Lambda^j = \sum_t w^{tj} \zeta_v^{tj} \tag{6.3}$$

When we have found the test statistic Λ^j, we can evaluate the hypothesis using the χ^2 distribution. Using a chosen significance level α, we can reject the hypothesis if

$$\Lambda^j > \chi^2_{\alpha,\nu} \tag{6.4}$$

where ν is the degrees of freedom for the χ^2-distribution.

To evaluate the turn rate information from the AIS messages, we formulate hypotheses similar to the ones used for the velocity information:

- H_0: The angular velocity measurement comes from the distribution $\mathcal{N}(\hat{x}_\omega, \hat{\sigma}^2_\omega)$ where $\hat{x}_\omega$ is the estimated angular velocity mean, and $\hat{\sigma}^2_\omega$ is the estimated angular velocity variance from the target submitting the AIS message.
- H_1: The angular velocity measurement does not come from the above distribution.

To perform the test, we need an estimate of the angular velocity of the targets from the tracker. We can obtain this with, for example, a CT model in an IMM framework or extended object tracking. The tracker we use, from Chapter 3, uses IMM with a CT model. As seen in Table 6.1, a target can either transmit a turn rate or only an indicator that the vessel is turning right or left at a rate larger than 5° per 30 seconds. For the case where the rate is transmitted with measurement z_ω, we calculate the test statistic as

$$\zeta_\omega^{tj} = \frac{(z_\omega - \hat{x}_\omega)^2}{\hat{\sigma}_\omega^2}. \tag{6.5}$$

In the second case, we calculate the test statistic by applying the cumulative distribution function of the standard normal distribution $\Phi(\cdot)$ as

$$\zeta_\omega^{tj} = \begin{cases} 1 - \Phi(Z) & \text{if } z_\omega = 5° \text{ per 30 seconds} \\ \Phi(Z) & \text{if } z_\omega = \text{-}5° \text{ per 30 seconds} \end{cases} \tag{6.6}$$

where

$$Z = \frac{z_\omega - \hat{x}_\omega}{\hat{\sigma}_\omega}. \tag{6.7}$$

For the case where the precise rate of turn is transmitted, the test statistic (6.3) is again χ^2-distributed with some significance level α and 1 degree of freedom. When only a turn indication is transmitted, we accept the hypothesis if $\Lambda^j \leq \alpha$.

Remark 6.1. If the exteroceptive measurements and the AIS messages are fused based on track-to-track fusion, we can, for example, update the AIS tracks with the AIS message first, before validating the information in the message with the radar track estimates. In that case, we replace the measurement likelihoods l^{tj} with the track-to-track association probabilities. ◁

6.3.4 Likelihood ratio test for detecting false AIS tracks

To detect wrong positional information, we formulate the following hypotheses:

- H_0: The position reported by the AIS message corresponds to the position of a target that is detectable by exteroceptive sensors. $P_D(\mathbf{x}_k^t|H_0) = P_D(\mathbf{x})$. Here, $P_D(\mathbf{x}_k^t)$ is the probability of detection of the target $\mathbf{x}_k^t$ at time step k.

- H_1: The position reported by the AIS message does not correspond to the position of a target that is detectable by exteroceptive sensors. Thus, any exteroceptive detections would have to be the result of clutter, and $P_D(\mathbf{x}_k^t|H_1) = \lambda(\mathbf{x}_k^t)$. Here, $\lambda(\mathbf{x}_k^t)$ is the probability of a false alarm where the target $\mathbf{x}_k^t$ is located.

The likelihood ratio test considers the likelihood of the data output from the tracker, conditioned on the hypotheses H_0 or H_1. First, we have that

$$\begin{aligned} l(\mathbf{x}_k^t|H_i) &= \sum_{a_k} l(\mathbf{x}_k^t, a_k|H_i) \\ &= \sum_{a_k} l(\mathbf{x}_k^t|a_k, H_i)p(a_k). \end{aligned} \tag{6.8}$$

Here, $p(a_k)$ is the probability of association hypothesis a_k, that is, the probability that some specific measurement j originates from the target that track t estimates. The likelihood $l(\mathbf{x}_k^t|a_k, H_i)$ is concerned with whether the data is reasonable with regards to the association between target and measurement in a_k and the hypothesis H_i. We formulate the likelihood for a single association as

$$l(\mathbf{x}_k^t|a_k, H_i) = \begin{cases} P_D(\mathbf{x}_k^t|H_i) & \text{if } a_k^t > 0 \\ (1 - P_D(\mathbf{x}_k^t|H_i)) & \text{if } a_k^t = 0 \end{cases} \tag{6.9}$$

With an application such as target tracking, where we receive more information over time, it is natural to use a sequential likelihood ratio test (SLRT) to test for false AIS messages. Wald first presented the SLRT in [135], which in our case is defined as

$$\mathcal{L}_k^t = \prod_{i=k_0}^{k} \frac{l(\mathbf{x}_i^t|H_0)}{l(\mathbf{x}_i^t|H_1)} \tag{6.10}$$

where k_0 is the time step when the track was initialized. By setting thresholds for when to decide on H_0 or H_1, the test ensures decision-making with a fixed maximum false alarm rate P_{FA} and a fixed minimum error detection probability P_{ED}. The thresholds decide how many samples are needed to make a decision, as the test only stops when either threshold is reached. From [135], the thresholds are

$$T_1 \leq \frac{P_{ED}}{P_{FA}}, \quad T_0 \geq \frac{1 - P_{FA}}{1 - P_{ED}}. \tag{6.11}$$

Remark 6.2. When using track-to-track fusion, we can replace the association probabilities in (6.8) with the association probabilities between the radar tracks and the AIS track to perform a similar test. ◁

6.4 Design of a safe target tracker using AIS

With the methods from Section 6.3, we design a target tracker that safely utilizes AIS information. First, we define some concepts to help us better describe the methods.

Definition 6.1. Quarantined AIS tracks are deemed to originate from false AIS messages. It will not be allowed to associate with exteroceptive sensor measurements and will be marked as quarantined until terminated. We do not immediately terminate quarantined tracks, so they can continue to associate with future false AIS messages.

Definition 6.2. Validated AIS tracks are tracks where the AIS messages are deemed not to be false. It will be allowed to associate with exteroceptive sensor measurements and marked as validated until terminated.

Definition 6.3. Preliminary AIS tracks are neither validated nor quarantined. These are tracks that have not provided enough information for the system to make a decision.

We divide the AIS validation framework into two parts: validation of individual AIS message fields and validation of AIS messages. A tracker structure that includes the two additional parts is shown in Figure 6.3.

6.4.1 Individual AIS message field validation

For the individual fields in the AIS message, we first check if they are within the allowed intervals in Table 6.1 before we use the tests in Section 6.3.3. The tests are performed on the velocity and turn rate, and they use the estimated velocities and angular velocities of the tracks associated with the incoming AIS message. New tracks initialized on AIS data should only use the positional data, as there is no way to validate the additional information without a previous estimate. Algorithm 5 shows how to perform the tests when an AIS message is received.

6.4.2 AIS track validation

For the validation and quarantining of AIS tracks we no longer only look at individual AIS messages, and use the SLRT described in Section 6.3.4. After some time, the AIS tracks are either validated or quarantined based on the result from the SLRT. Note that we also assume that the underlying target trackers use validation gating, which will remove any AIS messages with very large positional errors as a preliminary step. Thus, these need not be accounted for. Algorithm 6 shows the workflow for validating AIS tracks. We assume that classifying an AIS track as quarantined or validated also classifies all the AIS messages associated with the track.

Algorithm 5 Additional information validation

Require: Cluster with AIS measurement $\mathbf{z}^j$, measurement likelihoods l^{tj} calculated without additional information, existence probabilities r^t, threshold χ^2_α
  if information in $\mathbf{z}^j$ outside the intervals in Table 6.1 **then**
    discard additional information
  end if
  for track t in cluster **do**
    $w^{tj} \leftarrow r^t l^{tj}$ ▷ Weight given track t
    $\zeta^{tj} \leftarrow (\mathbf{z}^j - \hat{\mathbf{x}}^t)^\top (\mathbf{P}^t)^{-1} (\mathbf{z}^j - \hat{\mathbf{x}}^t)$ ▷ Test statistic
  end for
  $w^{tj} \leftarrow w^{tj} / \sum_t w^{tj}$ ▷ Normalize weights
  $\Lambda^j \leftarrow \sum_t w^{tj} \zeta^{tj}$ ▷ Combine test statistics
  if $\chi^2_\alpha \leq \Lambda^j$ **then**
    keep additional information
  else
    discard additional information
  end if

Algorithm 6 AIS track validation

Require: Marginal association probabilities p^{tj}, detection probability density $P_D(\mathbf{x})$, clutter density $\lambda(\mathbf{x})$, thresholds T_1 and T_0, previous SLRT $\mathcal{L}^t_{k-1}$
  for track t **do**
    $\mathcal{L} \leftarrow l(\mathbf{x}^t_i|H_0)/l(\mathbf{x}^t_i|H_1)$
    $\mathcal{L}^t_k \leftarrow \mathcal{L}^t_{k-1} \times \mathcal{L}$
    if $T_1 < \mathcal{L}^t_k$ **then**
      validate track t
    else if $T_0 \geq \mathcal{L}^t_k$ **then**
      quarantine track t
    end if
  end for

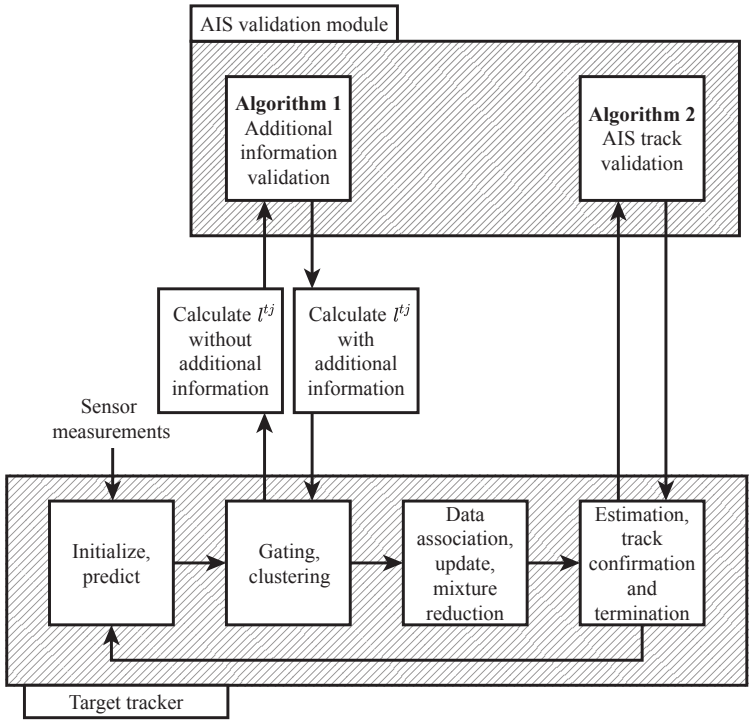

Figure 6.3: An outline of the program flow for validating and using AIS messages.

6.5 Results

6.5.1 Tracker implementation

The tracker used to provide the results is the multi-target tracker described in Chapter 3, of which an implementation is available at [69]. Furthermore, in addition to the AIS message validation framework, we expand the method to include rate of turn data in the update of the estimates. The update is done in one of two ways, depending on the type of turn rate information a target transmits. We use two CT models, with a positive and negative turn rate, respectively, to get the desired behavior. If the target only transmits the turn direction, we perform the update by giving the kinematic model s with the corresponding turning direction a high probability when calculating the measurement likelihood:

$$l(\mathbf{z}|\mathbf{x}, s) = f_{z_s}(z_s|s)l(\mathbf{z}|\mathbf{x}) \tag{6.12}$$

where

$$f_{z_s}(z_s|s) = \begin{cases} P_{\text{turn}} & \text{if } s = z_s \\ 1 - P_{\text{turn}} & \text{if } s \neq z_s \end{cases}. \tag{6.13}$$

Here, P_{turn} denotes the probability that the target turns in the direction corresponding to model s given a match with measurement z_s. P_{turn} serves the purpose of not letting a single kinematic model obtain all the probability mass. We also use (6.12) if the target transmits the turn rate, but then the estimated turn rate in the corresponding model is also updated.

6.5.2 Performance analysis on simulated data

We want to know how effectively the tracker detects AIS measurement errors. For this purpose, we look at the receiver operating characteristic (ROC) curves for the different error types: false tracks, sporadic and repeating zero- and random-valued velocity errors, and sporadic or repeating zero-valued turn rate errors.

The results are based on 150 simulated scenarios that include radar measurements and AIS messages. The scenarios consist of ten targets maneuvering within a surveillance area with a radius of 1000 meters. The reader is referred to Chapter 3 for a more detailed description of the modeling of targets and measurements. Four of the ten targets created for each scenario represent the trajectories of false AIS messages. We randomly selected the four false track targets from the set of ten and removed all radar measurements generated by them. We kept the AIS messages transmitted by the removed targets, and they served as the false AIS message input. For the sporadic error tests, each AIS message had a 5% chance of containing erroneous velocity information. Furthermore, for the repeating error tests, each AIS message had a 1% chance of being the start of a sequence of messages with erroneous velocity information. An AIS message in an erroneous message sequence had a 5% chance of being correct and ending the sequence.

The ROC curves in Figure 6.4 show the probability of error detection P_{ED} as a function of the probability of false alarm P_{FA}. We see in Figure 6.4a and Figure 6.4c that the tracker can detect the errors in the velocity and false tracks with high probability. However, the turn rate error detection, seen in Figure 6.4b, is not as effective. This reflects the general difficulty in estimating the turn rate using noisy measurements. We are nevertheless able to detect a meaningful percentage of the errors. Furthermore, we see in Figure 6.4c that the SLRT test for false tracks becomes more precise with more radar scans, but it still performs well when using only three samples.

We also investigate the average error between the undetected wrong velocity measurements and the true state. The results are shown in Figure 6.5. The error decreases sharply for low significance levels before leveling off. The curve indicates

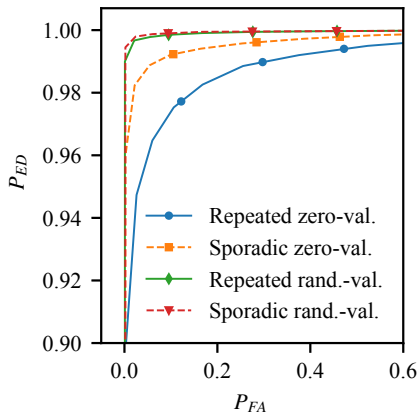

(a) Velocity error test success rates for the different error combinations.

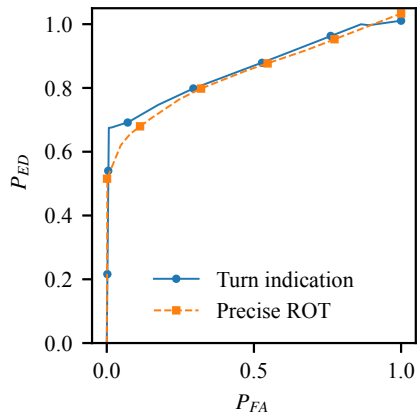

(b) The turn rate error test success rates for the two different measurement types.

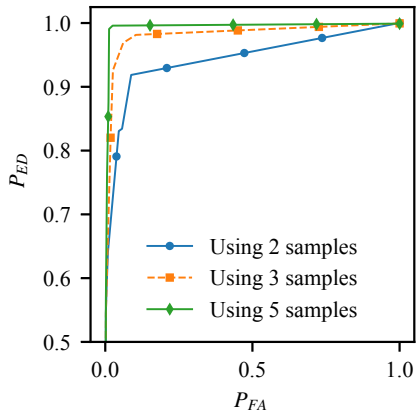

(c) The false track test success rates, with a fixed number of samples before a decision is made.

Figure 6.4: ROC curves for the velocity error detector (a), angular velocity error detector (b), and the false track SLRT detector (c).

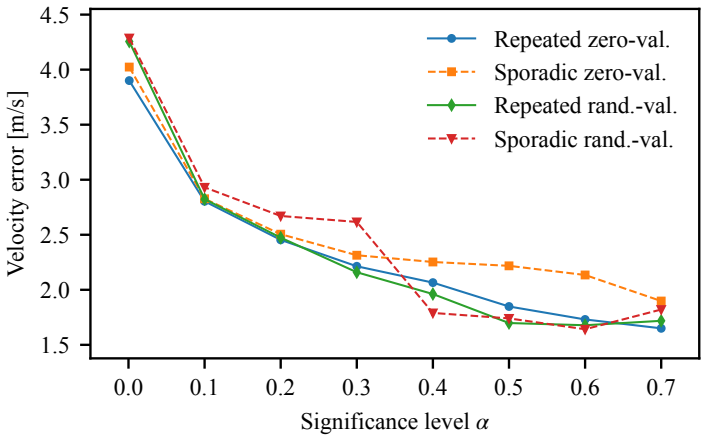

Figure 6.5: Average error between the velocity of an AIS message where the wrong velocity info is undetected, and the true velocity of the originating target.

that the rejection of erroneous measurements is limited by the uncertainty of the target estimate. However, significant errors are likely to be detected as these would differ enough even from an uncertain estimate.

6.5.3 Demonstration on experimental data

Figure 6.6 shows how the method performs when applied to an experimental data set with injected false AIS messages. The data were collected as part of the preparation for a demonstration of autonomous collision avoidance in the Trondheim Fjord and shows two targets crossing while the ship with the radar maneuvers to avoid collision. False AIS messages were added by copying the collected messages from the targets and offsetting these in space and time. Furthermore, the velocity reported by the AIS messages from the actual targets had a 5% chance of obtaining a random value within the permitted interval.

As shown, the false and erroneous AIS messages confuse the tracker when not accounted for, and if used as input to a collision avoidance system it would likely paralyze the decision-making process. Tracks are initialized on the false AIS messages, and the erroneous velocities also lead to false tracks. However, when utilizing the tracker structure in Figure 6.3, the tracker is able to only report estimates on the true targets and disregard the false AIS information.

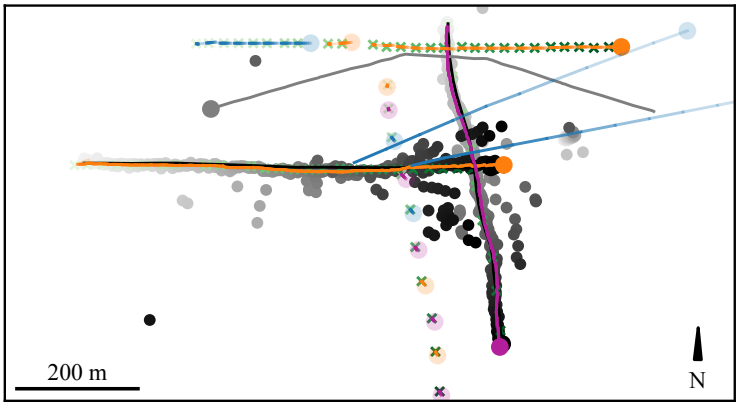

(a) Tracker output when not accounting for potentially wrong AIS messages.

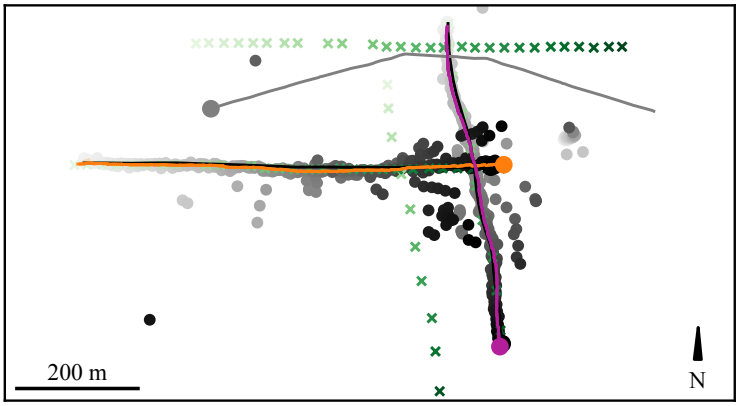

(b) Tracker output when accounting for potentially wrong AIS messages.

Figure 6.6: Tracking of multiple targets using the tracker from Chapter 3. Two targets are shown, in orange and purple, and the ship with the radar is shown in grey. The other graphical elements are explained in Figure 6.2.

6.6 Chapter summary

We have introduced techniques for safely utilizing AIS messages in target tracking while validating them for use in other system components. The validation allows the tracking component to supply reliable AIS data to other functions used in unmanned or autonomous maritime vessels, such as collision avoidance methods. We evaluate AIS messages within a statistical framework that uses additional exteroceptive sensors to provide trustworthy estimates. We have discussed the most relevant AIS fields for the tracking system and developed methods to validate them. Our results demonstrate the effectiveness of the hypothesis tests for detecting wrong velocity information and that the tracker is capable, but somewhat less effective, when it comes to detecting errors in the rate of turn. We also show that the tracker can validate or quarantine erroneous tracks with a high probability of detection and a low probability of false alarms based on only a few exteroceptive sensor scans. Lastly, we demonstrate the method on experimental data, where the tracker is able to handle the injected false AIS information.

Acknowledgement

The experimental data were collected in co-operation with Maritime Robotics, who provided personnel and infrastructure.

Part III

Selected Topics in Target Tracking With Exteroceptive Sensors

Chapter 7

Compensating Radar Rotation in Target Tracking

This chapter is based on the publication:

[67] A. G. Hem and E. F. Brekke. Compensating radar rotation in target tracking. In *Proc. Sensor Data Fusion: Trends, Solutions, Applications*, Bonn, Germany, Oct. 2022

Changes from the original publication involve shortening of the introduction.

7.1 Introduction

Most target tracking methods assume that all measurements from a single scan arrive simultaneously, for example, [108], and the one described in Chapter 2. However, such an assumption will only be a simplifying approximation when the measurements originate from a rotating radar. Rather than explicitly considering the timing of the individual measurements, additional induced measurement noise usually masks the timing inaccuracies of the models. Furthermore, the possible performance degradation of approximating the detections as simultaneous has not been investigated in the literature, which raises the question of how big of an impact the radar rotation has on the target estimates.

Some research has previously been conducted concerning radar rotation and handling measurements with individual time stamps. In [59], Habtemariam et al. presented a data association method for rotating radars to reduce tracking latency in systems with long time intervals between radar measurement updates. The model presented by Gaglione et al. in [44] processes AIS measurements using individual time stamps. The radar measurements, however, are handled assuming

simultaneous arrival. Several researchers have also done work on spinning radars for use in navigation and Simultaneous Localization And Mapping (SLAM) [129]. In [27], Burnett et al. consider the impact of Doppler effects and motion distortion in spinning radar navigation, showing that spinning radar navigation methods should consider motion distortion. A similar problem, that of using rolling shutter cameras in navigation, was investigated in [92]. For the more general problem of filtering of measurements with time delays, solutions include using a multi-rate Kalman filter [76], and extrapolation of the delayed measurements [89]. Albeit SLAM and radar navigation places different demands on the sensors in use, the previously mentioned research indicate that unaccounted radar rotation can degrade performance.

This work investigates the impact of the assumption of simultaneously arriving measurements when the measurements come from a spinning radar. The investigation is carried out within the framework of a joint integrated probabilistic data association (JIPDA) tracker [105].

The main contribution of this chapter lies in investigating the effects of radar rotation and finding ways of mitigating these effects. First, we identify problems caused by radar rotation. Then, we present a method where the model does not include a synchronization assumption before we present an implementation containing the mitigating solutions. Lastly, we evaluate the performance loss from assuming simultaneously arriving measurements and the effect of the proposed solutions.

7.2 Problem formulation

The problem consists of tracking an unknown number of targets using noisy measurements from a radar, with the presence of false alarms and missed detections. The radar is assumed to be situated on top of a maneuvering ship, rotating with a fixed rotation rate. The target births are modeled according to a Poisson Point Process (PPP) model with density $b(\mathbf{x})$. Furthermore, targets are assumed to survive from time t_k to time t_{k+1} with probability $P_S(\mathrm{t}_k, \mathrm{t}_{k+1})$. The movement of a target from time t_k to time t_{k+1} is assumed to follow a transition density $f_{\mathbf{x}}(\mathbf{x}_{k+1}|\mathbf{x}_k)$.

The probability of a target being detected is P_D, chosen as a constant for simplicity, and the measurement has likelihood function $f_{\mathbf{z}}(\mathbf{z}|\mathbf{x})$. Clutter measurements arrive according to a homogeneous stationary PPP with parameter λ. We make the standard assumptions of a target only being the origin of a single measurement and that a measurement can not originate from more than one target. The exact position and pose of the ownship, that is, the ship carrying the radar, is assumed to be known, and the movements of the targets and ownship are assumed to be constant between time steps. The above model choices are the same as in Chapter 2, where further details can be found.

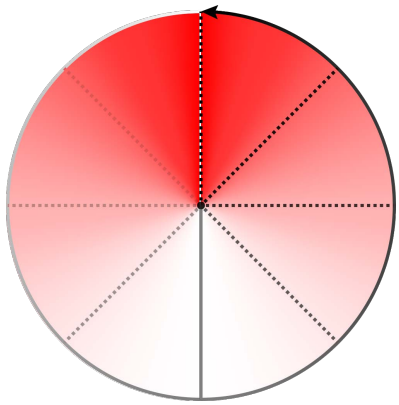

Figure 7.1: A radar surveillance area, with the radar rotating with a counter-clockwise motion. The measurements are time stamped at the time when the radar has finished half its rotation. The red color represents the bias of the targets, and a clearer red signifies a larger bias.

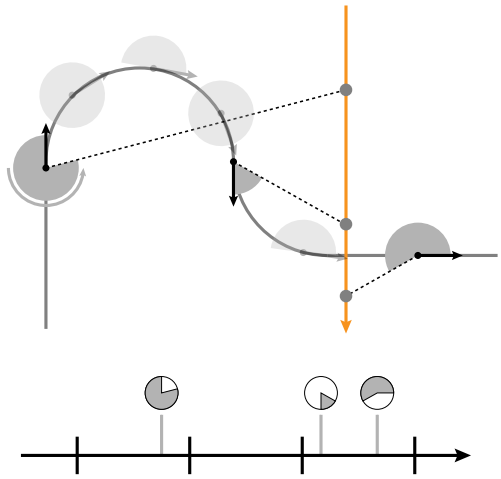

Figure 7.2: How a double detection can occur. The gray discs indicate the progress of the counter-clockwise radar rotation, and the radar is located on an ownship following the gray path. The target follows the orange path, and gray dots mark its detections. Detection times can be seen on the timeline.

The measurements are assumed to arrive at intermediate time steps k_i between time steps $k-1$ and k. This assumption differs from the one usually applied in the models, for example, [108], because the radar measurements are not assumed to originate from the same point in time, that is, simultaneously. Each bearing resolution cell provides an intermediate time step, indexed by i. The intermediate time step k_i belonging to a measurement is defined by the radar spoke in which the measurement lies. This formulation allows us to investigate the effects of radar rotation on the estimates.

7.3 Error sources from non-compensated radar rotation

We identified three effects caused by radar rotation:

- When making the simplifying assumption that all radar detections from a radar measurement batch are simultaneous, a bias is introduced in the system.
- A radar sweep will not necessarily correspond to a single, complete sweep of the surveillance area.
- The prediction intervals between time steps are not constant.

The consequences of these effects are, together with how they occur, explained here.

7.3.1 Bias

When not accounted for, radar rotation will result in a bias. This bias comes from the movement of targets between their detection and when the radar batch is time stamped. The targets at the start and end of the rotation have moved at most half a radar period dt relative to their detection. Figure 7.1 illustrates this.

Proposition 7.1. Assuming that the targets in the surveillance area are uniformly distributed and that the radar measurements are time stamped to the midpoint of the radar rotation, an approximation of the induced bias B per target over time is given as

$$B = \frac{\mathrm{dt}_R}{4}\bar{v} \tag{7.1}$$

where dt_R is the radar period, and $\bar{v}$ is chosen as the average speed of the targets. The choice of $\bar{v}$ could also be, for example, the speed of a specific target, or a perceived maximum target speed, depending on the application.

Proof. The duration a target can move as a function of where in the radar rotation it is detected can be written as $\mathrm{dt}_R\theta/(2\pi)$, where θ is the relative bearing. Furthermore, the probability of a target existing in a segment of length $d\theta$ is $d\theta/(2\pi)$ because the targets are uniformly distributed. Integrating, we get that the average time a target can move relative to its detection is

$$\int_{-\pi}^{\pi} \frac{\mathrm{dt}_R}{2\pi}\theta\frac{d\theta}{2\pi} = \left[\frac{\mathrm{dt}_R}{8\pi^2}\theta^2\right]_{-\pi}^{\pi} = \frac{\mathrm{dt}_R}{4}. \tag{7.2}$$

Knowing this, the bias B becomes the average velocity $\bar{v}$ times $\mathrm{dt}_R/4$. □

The targets will have almost correct timing relative to other targets in their vicinity, except for the angle where the radar finishes the collection of a batch of measurements. Around this angle, measurements collected almost a whole radar period apart from each other end up being processed as part of the same cluster.

7.3.2 Multiple detections of the same target

Another consequence of radar rotation is that two detections of the same target can occur in the same radar scan. The time difference between detections at the start and at the end of the radar rotation is a radar period dt_R. Thus, a target could move so that the radar detects it both at the start and end of its rotation. We denote this phenomenon as a double detection. As shown in Figure 7.2, the ownship movement exacerbates this problem due to the radar not covering the surveillance area exactly once. A target moves in a straight line while the ownship makes several turns. The target is detected each time the radar is pointed at it, but due to the ownship movement its detection is not present in each radar measurement batch. The timeline in Figure 7.2 shows the times the target is detected, and the fixed interval ticks indicate when the radar has completed a rotation, which also means that the radar outputs a radar measurement batch. The radar detects the target in the first batch, but in the second, it does not. In the next batch, however, two detections are present. This phenomenon breaches the assumption that each target can only be the origin of a single measurement and may cause undesired and unpredictable behavior.

Proposition 7.2. The number of double detections can be approximated as

$$N_{double} = \frac{N_m A_\phi}{2} + \mathrm{C} \tag{7.3}$$

where N_m is the number of target detections and C is a constant. A_ϕ is a fraction of the surveillance area decided by

$$A_\phi = \frac{2\pi + \Delta\phi}{2\pi}. \tag{7.4}$$

Here, $\Delta\phi$ is the total change in heading for the ownship during a radar period.

Proof. Under the assumption that the sum of the changes in heading over time will be approximately zero, the area that will be scanned twice by the radar will be $A_\phi/2$ times the surveillance area. Thus, the number of double detections becomes $N_{double} = N_m A_\phi/2$, plus a constant C which represents the double detections caused by target movement, rather than ownship movement. □

7.3.3 Variable prediction intervals

Prediction intervals are assumed to be constant, but this is an approximation when the target or the ownship is moving. Unless the target remains in the same bearing resolution cell, the time between detections will slightly differ from the radar period. This fact means fast-moving targets will cause more significant prediction interval errors, as will the ownship movement. When the ownship heading changes, so does the bearing cell the target detection occurs in. Thus, a large $\Delta\phi$ will give a large prediction timing error. If the prediction interval is inaccurate, this will result in less accurate predictions and, ultimately, less accurate estimates.

7.4 Removing the measurement synchronization assumption

Here, we present the changes required for a JIPDA to consider the exact time stamps. This will reduce the bias and provide precise prediction intervals. The changes consist of modifying the prediction, and the detection part of the update. That is, we use the same prior as in the regular JIPDA. For the update, the expressions for a new target and a missed detection remain unchanged. The notation and structure of the mathematical explanation is based upon that of Chapter 2, albeit somewhat simplified.

7.4.1 Prediction

When predicting the target estimates, often denoted as tracks, one usually assumes that the prediction is for an interval of constant or nearly constant length. We do not make the same assumption here. As a result, the varying prediction intervals influence the Markov chain modeled existence model and the kinematic model. The only impact on the kinematic modeling is through the implementation of the model, and we reserve the details for later. However, the prediction of the existence probability requires using continuous-time Markov chains.

Proposition 7.3. The prediction for a target estimate for a general time interval $\mathrm{t}_{k_i} - \mathrm{t}_{k-1} \triangleq \mathrm{dt}$, that is, between time steps $k-1$ and k_i, the existence probability for track t becomes

$$r^t_{k_i|k-1} = P_S^{\mathrm{dt}} r^t_{k-1} \tag{7.5}$$

and the kinematic state becomes

$$f^t_{k_i|k-1}(\mathbf{x}) = \int f^t_{k_i}(\mathbf{x}|\tilde{\mathbf{x}}) f^{\mathrm{dt}}_{k-1}(\tilde{\mathbf{x}}) \mathrm{d}\tilde{\mathbf{x}}. \tag{7.6}$$

Proof. We model the existence as a two-state Markov chain with a generator matrix

$$G = \begin{bmatrix} P_S^* - 1 & 1 - P_S^* \\ 0 & 0 \end{bmatrix}. \tag{7.7}$$

It is well-known [88, p. 198] that a generator matrix on this form results in the transition matrix

$$P(\mathrm{dt}) = \exp\left(\mathrm{dt} \begin{bmatrix} P_S^* - 1 & 1 - P_S^* \\ 1 & 1 \end{bmatrix} \right). \tag{7.8}$$

We are only concerned about the survival probability, that is, the probability of remaining in the first state. It becomes $p_{11}(\mathrm{dt}) = e^{(P_S^* - 1)\mathrm{dt}} \triangleq P_S^{\mathrm{dt}}$ where $P_S = e^{P_S^* - 1} < 1$ is a constant. Furthermore, (7.6) is the prediction equation from Chapter 2, with prediction interval up to an intermediate time step k_i.

□

7.4.2 Detection

The mathematical formulation for updating a track is elaborately explained in Chapter 2. The reader is referred there for information regarding initialization of new tracks and the updating of a track based on a missed detection. However, the change in how the measurements are modeled demands a change in how we conduct updates based on detections.

Proposition 7.4. Suppose we have a measurement $\mathbf{z}_j$, which arrives at the intermediate time step k_i. The detection of track dt through measurement j is weighted as

$$w^{tj} = P_D r^{tj}_{k_i|k-1} l^{tj}. \tag{7.9}$$

The weight w^{tj} is used to calculate the probability of the association hypotheses which include the association between track dt and measurement j. The association

hypotheses consist of feasible subsets of associations between tracks and measurements. The expressions for the existence probability $r_{k_i}^{tj}$, and the kinematic state $f_{k_i}^{tj}(\mathbf{x})$, conditional on the association between track dt and measurement j, becomes

$$r_{k_i}^{tj} = 1 \tag{7.10}$$

$$f_{k_i}^{tj}(\mathbf{x}) = f_{\mathbf{z}}(\mathbf{z}|\mathbf{x}) f_{k_i|k-1}^{t}(\mathbf{x})/l^{tj} \tag{7.11}$$

where

$$l^{tj} = \int f_{\mathbf{z}}(\mathbf{z}|\tilde{\mathbf{x}}) f_{k_i|k-1}^{t}(\tilde{\mathbf{x}}) \mathrm{d}\tilde{\mathbf{x}}. \tag{7.12}$$

Proof. The above equations are simply the update from Chapter 2 for time step k_i, where the predictions from $k-1$ to k_i are acquired by use of (7.5) and (7.6). The assumption that the targets are the origin of at most one measurement in the span of a radar revolution, and that a measurement comes from at most one target, enables us to use the same data association methods as in any tracker assuming synchronized measurements. □

Remark 7.1. The association hypotheses can be formulated as in Chapter 2. They will, however, consist of associations from different intermediate time steps. While this may seem a problem at first glance, we can use the weights the same way as if the measurements arrived simultaneously. The weights between measurement j and possible tracks are calculated based on the possible tracks' predictions up to the same intermediate time step, the one when measurement j arrives. Thus, comparisons between a single measurement and several tracks are made based on weights calculated from the same intermediate time step. Furthermore, any comparisons between a single track and several measurements are made on weights calculated based on the intermediate time steps of the measurements. In sum, this means that the time differences between measurements always will be accounted for when comparing weights so that we can use them the same way we would in a regular JIPDA. ◁

7.4.3 Mixture reduction

The existence probability and kinematic state are predicted from time step k_i to k to be in sync with states updated based on other association hypotheses. This prediction is performed using (7.5)-(7.6) with the values acquired from (7.10)-(7.11), giving $r_{k|k_i}^{tj}$ and $f_{k|k_i}^{tj}(\mathbf{x})$. Recognizing that these quantities serve the same purpose as if they had been the result of a measurement arriving at time t_k, the updated estimate after time step k can be calculated as one would do if assuming simultaneous measurements.

7.5 Implementation

The tracker used to create the results in this chapter is implemented similarly to the one described in Chapter 2. We perform mixture reduction as in Chapter 2, only without the multiple kinematic models. While the tracker from Chapter 2 implements a visibility state and multiple kinematic models, for simplicity, the tracker used here does not. We use the same method for initialization and termination of tracks, and Murty's method [104] is used together with the auction method [87] to make the data association problem computationally tractable. Furthermore, we make the same Gaussian-linear assumptions and use the same measurement model. Some necessary details, however, must be considered when using a model with time stamps for the individual measurements.

7.5.1 Kinematic model

A (nearly) Constant Velocity (CV) model [8, p. 269] is used to model the movement of the targets. Because of the varying prediction intervals, we use the discretized continuous formulation of the model. The discrete CV model formulation is only suitable for constant prediction intervals. For the individual dimensions, the transition matrix $\mathbf{F}$ and process noise covariance matrix $\mathbf{Q}$ are

$$\mathbf{F} = \begin{bmatrix} 1 & \mathrm{dt} \\ 0 & 1 \end{bmatrix} \text{ and } \mathbf{Q} = \begin{bmatrix} \mathrm{dt}^3/3 & \mathrm{dt}^2/2 \\ \mathrm{dt}^2/2 & \mathrm{dt} \end{bmatrix} q, \tag{7.13}$$

where q is the process noise intensity [8, p. 270].

7.5.2 Measurement gating

When gating the measurements that have arrived between time steps $k-1$ and k, it would be advantageous to do so without performing predictions up until the time of arrival for all measurements. For this purpose, a gating method can be used, which first uses the maximum realistic velocity of a target to find a subset of the measurements before predicting up to and gating only the measurement subset. The first part of the procedure is described in [141], while the second part amounts to evaluating

$$(\hat{\mathbf{z}}_{k_i} - \mathbf{z}_j)\mathbf{S}_{k_i}^{-1}(\hat{\mathbf{z}}_{k_i} - \mathbf{z}_j) < g^2 \tag{7.14}$$

using the output of a Kalman filter for each intermediate time step k_i.

7.5.3 Measurement batches

In Section 7.3, we showed that highly maneuvering ownships could result in increased double detections. However, given some preconditions, we can counter the effects shown in Figure 7.2. The idea is to define a radar measurement batch as the result of a complete rotation relative to the surveillance area rather than a rotation relative to the ownship. This approach assumes knowledge regarding the ownship heading. Denoting the rotation of the radar since the last intermediate time step as $\Delta\phi$, and the ownship's change in heading as $\Delta\psi$, the measurement batches are created according to Algorithm 7.

Algorithm 7 Measurement batch creation

Require: Radar measurement set Z_{k_i}, ownship heading change $\Delta\phi$, radar angle change $\Delta\psi$, surveillance area coverage θ

$Z_k \leftarrow Z_k \cup Z_{k_i}$

$\theta \leftarrow \theta + \Delta\phi + \Delta\psi$

if $\theta \geq 2\pi$ **then**

$k \leftarrow k + 1$

$Z_k \leftarrow \{\}$

$\theta \leftarrow \theta - 2\pi$

end if**return** $Z_{1:k}$, θ

7.6 Simulation setup and results

7.6.1 Simulation setup

We assume that measurements come from a rotating radar, with a bearing resolution of 2° and a range resolution of 2 meters. Each radar period contains $360^\circ \div 2^\circ = 180$ possible time stamps for radar detections. Unless otherwise stated, the radar periods are $T = 2$ seconds long. A radar measurement batch is either the result of one complete radar rotation or the output of Algorithm 7. The simulator also supplies simultaneously arriving measurements to act as a benchmark.

The simulated scenarios consist of seven targets and a single ownship, all maneuvering inside a circle with a 2000-meter radius for 100 seconds. The data is created according to the assumptions in Section 7.2. The ownship and the targets move according to the CV model from Section 7.5.1, with the process noise intensity q determining the magnitude of their acceleration changes. Furthermore,

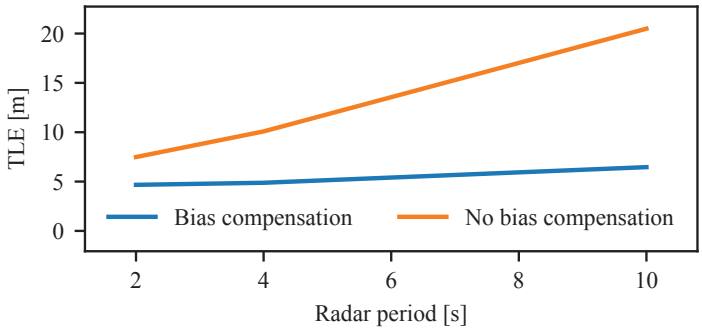

Figure 7.3: TLE values with increasing radar period, tracking a single target with constant maneuvering index [8, p. 287].

the simulator creates the measurements according to the measurement model from Chapter 2.

The scenarios are created to be somewhat challenging for the tracker, with parameters found in Table 7.1. Evaluation of the results is done by use of the Track localization error (TLE) [29] and the Optimal sub-pattern assignment (OSPA) metric [122]. 250 data sets were used for evaluation.

Quantity	Symbol unit	Value
Target process noise intensity q	$[\mathrm{m}^2/\mathrm{s}^3]$	2.0
Probability of detection P_D	[-]	0.8
Survival probability P_S	[-]	0.99
Cartesian measurement noise σ_{xy}	[m]	6.6
Range measurement noise σ_{r}	[m]	8.0
Bearing measurement noise σ_θ	[°]	1.0
Clutter intensity λ	$[1/\mathrm{m}^2]$	5×10^{-7}
Initial velocity std.	[m/s]	10

Table 7.1: Tracker parameters.

7.6.2 Effect of decreased radar frequency

As described in Section 7.3.1, the radar rotation will, when not accounted for, introduce a bias in the estimates. Simulations with increasing radar periods, shown in Figure 7.3, demonstrate that a longer radar period results in a more significant error. This result is expected, as the detected targets will be able to move further

away from where they were detected when the radar frequency is low. This linear increase is also in line with the approximation in (7.1).

7.6.3 Effect of increased ownship acceleration

We see from Table 7.2 that double detections happen more often with more ownship maneuvers. These results align with the explanation in Section 7.3.2. We can also see from Table 7.2 that the prediction interval timing errors increase similarly. As explained in Section 7.3.3, large changes in the ownship heading during a radar period will give large prediction interval timing errors.

Process noise intensity q [$\mathrm{m^2 s^{-3}}$]	0.5	1.0	1.5	2.0
Timing error [s]	0.155	0.256	0.303	0.340
Double detections [%]	1.83	3.06	4.44	4.84

Table 7.2: Timing errors and double detections for increasing ownship process noise intensity q.

7.6.4 Effect of compensating methods

We evaluated four slightly different JIPDA trackers: one with none of the changes implemented, one with only the changes from Section 7.4, one with only the measurement batch solution in Section 7.5, and lastly, one with both.

Figure 7.4 shows the performance of the different JIPDA combinations. We include the results for when the measurements arrive simultaneously as a benchmark for the best possible performance. We see from the OSPA values that the overall tracking performance decreases with increasing maneuvering. When only considering the TLE, however, we see that the positional error of the estimates is almost invariant to ownship movement. This indicates that the performance degradation seen with increasing ownship maneuvering is due to track jumps and problems with data association. Furthermore, Figure 7.4 shows that the best result is achieved when we both compensate for the bias and use the batch creation method.

While the methods give a performance increase individually or in combination, the performance is still clearly better when all measurements arrive simultaneously. This performance gap is in part because we evaluate the performance at the end of the radar period, so when the radar rotates we evaluate the performance based on tracks predicted from some time within the radar period.

Table 7.3 shows the average NEES [8, p. 168] values. It is clear that the bias from the radar rotation, when not accounted for, makes the tracker overconfident. In practice the overconfidence would be mitigated by increasing the measurement

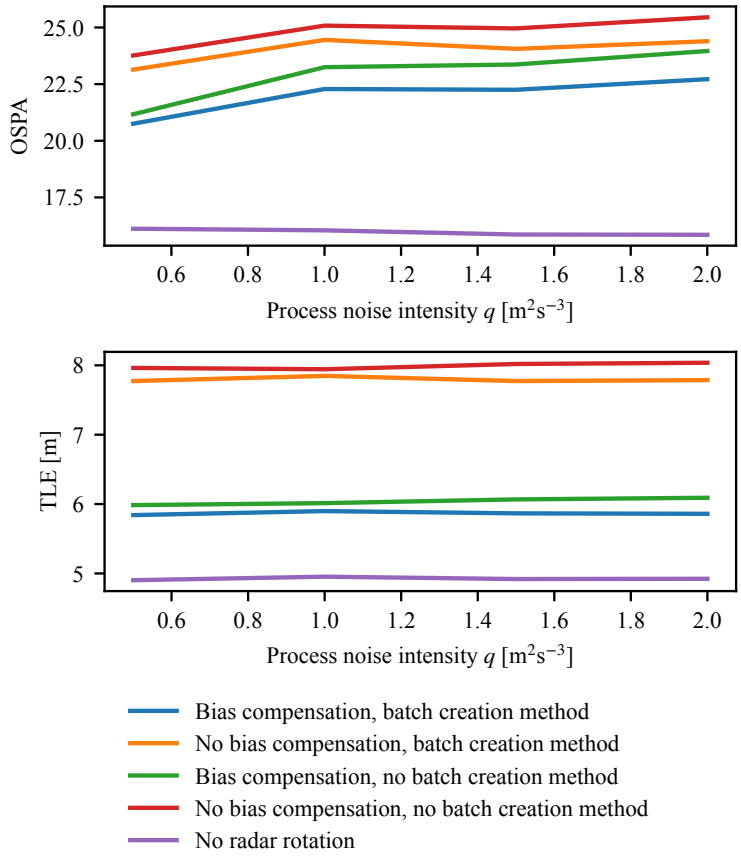

Figure 7.4: OSPA and TLE for the different methods when using a spinning radar and when the radar delivers all measurements simultaneously. The results show the methods with increasing ownship process noise intensity q.

	Average NEES
95% confidence interval	$[3.99, 4.01]$
No rotation	3.80
With bias compensation, no batch creation method	3.70
Without bias compensation, no batch creation method	12.30

Table 7.3: Average NEES values with ownship process noise intensity $q = 1.0$.

covariance. When the bias is compensated the average NEES values are similar to when the radar measurements arrive simultaneously.

7.7 Chapter summary

In this work, we have shown the assumption of measurements from a rotating radar arriving simultaneously to become dubious under certain conditions. We have presented approximate expressions for the magnitude of the introduced bias, and the number of double detections. Furthermore, methods for mitigating these problems have been presented together with implementation details. Results show both the impact of the radar rotation and the performance gain achieved through the proposed methods.

Future investigations should include how the effects from Section 7.3 impact the fusion of measurements from different sources. Furthermore, showing the effects in experimental data will further clarify their impact.

Chapter 8

Target Tracking With Existence Modeling in the Presence of Wakes

This chapter is based on the publication:

[64] A. G. Hem, H.-G. Alvheim, and E. F. Brekke. WakeIPDA: Target Tracking With Existence Modeling in the Presence of Wakes. In *Proc. IEEE 26th International Conference on Information Fusion*, Charleston, SC, USA, June 2023

Changes from the original publication involve shortening and editing of the introduction section.

8.1 Introduction

Wakes are likely to be present when tracking targets in the maritime domain. Intuitively, the wake clutter measurements usually occur in an area behind the target and depend on the target state. When not modeling the wake clutter, the tracker will often mistake it for target measurements, resulting in poor tracking performance [4]. In addition, the wake clutter measurement density is usually much higher than the regular clutter density, further complicating the tracking problem.

The tracking algorithm presented in this chapter builds upon the integrated probabilistic data association filter (IPDA) [106]. The IPDA is an extension of the probabilistic data association filter (PDA) [9], to which the IPDA adds an extra state modeling the existence probability of the target. The IPDA is a well-known and widely used algorithm and has been applied to the tracking of surface [137] and underwater [18] targets in the maritime domain.

Several methods for mitigating the effects of wake clutter already exist. [133] presents a general framework for handling measurements that do not adhere to

the usual assumptions made when developing target tracking methods. In [117], Rødningsby and Bar-Shalom present a PDA method designed to track a diver with bubbles in its wake. They expand the PDA to support multiple targets in [118], and finally, multiple sensors in [119]. Furthermore, [19] presents a multi-target tracking method that uses marginalization to mitigate the wake clutter effects. The use of marginalization distinguishes it from the methods presented in [117–119].

The methods mentioned in the previous paragraph are based on the PDA, or its multi-target equivalent, the joint PDA. They do not take advantage of the IPDA's ability to model the existence probability of the target. We present a novel single-target tracking method called WakeIPDA, which enables the IPDA to account for wake clutter. This extension helps mitigate the undesirable effects of wake clutter and takes advantage of the information inherent in wake clutter measurements through their implicit confirmation of target existence.

The chapter is structured as follows. In Section 8.2 we formulate our problem before we present the WakeIPDA algorithm in Section 8.3. Furthermore, Section 8.4 presents two ways of modeling the wake of a target. In Section 8.5 we present the simulation setup and results, and finally, in Section 8.6 we conclude the chapter.

8.2 Problem formulation

We want to track a target using noisy measurements that may include false alarms and missed detections. The targets are assumed easily separable, such that it is reasonable to use a single-target method. We model the birth of the target as a Poisson point process (PPP) with a density of $\lambda_b(\mathbf{x})$. See Chapter 2 for more details regarding target birth. The target is assumed to have a constant detection probability of P_D, while its movement from one time step to the next follows a transition density of $f_{\mathbf{x}}(\mathbf{x}|\cdot)$. We assume that the probability of a target surviving from one time step to the next is P_S.

The likelihood function for the target measurement is $f_{\mathbf{z}}(\mathbf{z}|\mathbf{x})$. Furthermore, we make the assumption that a single target can only be the origin of one target measurement and a single target measurement must come from only one target. Regular clutter measurements are generated from a homogeneous stationary PPP with constant intensity $\lambda_0(\mathbf{z}) = \lambda_0$. We also model wake clutter as a homogeneous stationary PPP, but with intensity λ_w within the wake area and intensity 0 outside it. The combined spatial clutter likelihood is thus $\lambda(\mathbf{z}|\mathbf{x}) = \lambda_0 + \lambda_w(\mathbf{z}|\mathbf{x})$. The clutter PPPs are the same as in [19], where more details are provided.

8.3 The WakeIPDA algorithm

We now present a mathematical description of the WakeIPDA algorithm, based on the above problem formulation. The input to the algorithm is a Bernoulli, serving as our prior, and a set of measurements Z. The prior is defined as

$$f_{k-1}(X_{k-1}) = \begin{cases} 1 - r_{k-1} & \text{if } X_{k-1} = \emptyset \\ r_{k-1} f_{k-1}(\mathbf{x}) & \text{if } X_{k-1} = \{\mathbf{x}\} \end{cases}. \tag{8.1}$$

Furthermore, the algorithm consists of two steps: a prediction step, and an update step. The prediction step is the same as in the IPDA [106].

Proposition 8.1. The prediction between time steps $k-1$ and k is given by

$$r_{k|k-1} = P_s r_{k-1} \tag{8.2}$$

for the existence probability, and

$$f_{k|k-1}(\mathbf{x}) = \int f_{\mathbf{x}}(\mathbf{x}|\tilde{\mathbf{x}}) f_{k-1}(\tilde{\mathbf{x}}) \mathrm{d}\tilde{\mathbf{x}}. \tag{8.3}$$

for the kinematic state.

We continue by presenting the update step, which is where the model changes have to be taken into account.

8.3.1 Including wake clutter in the measurement likelihood

The inclusion of target-dependent wake clutter means that the measurement likelihood needs to be reformulated. Rather than only being a function of the target measurement, as is usually the case, the measurement likelihood becomes a function of all the measurements.

Proposition 8.2. The measurement likelihood for a set Z_k of measurements at time k is given by

$$f_k(Z_k|X_k = \emptyset) = \Big(\prod_{\mathbf{z} \in Z_k} \lambda_0 \Big) \exp\Big(- \int_{\mathcal{A}} \lambda_0 \mathrm{d}\tilde{\mathbf{z}} \Big) \tag{8.4}$$

if a target is not present, and

$$\begin{aligned} f_k(Z_k|X_k = \{\mathbf{x}\}) = (1 - P_D) \exp\Big(- \int_{\mathcal{A}} \lambda(\tilde{\mathbf{z}}|\mathbf{x}) \mathrm{d}\tilde{\mathbf{z}} \Big) \prod_{\mathbf{z} \in Z_k} \lambda(\mathbf{z}|\mathbf{x}) + \\ \exp\Big(- \int_{\mathcal{A}} \lambda(\tilde{\mathbf{z}}|\mathbf{x}) \mathrm{d}\tilde{\mathbf{z}} \Big) \sum_{j=1}^{m_k} P_D f_{\mathbf{z}}(\mathbf{z}_k^j|\mathbf{x}) \prod_{\mathbf{z} \in Z_k \setminus \mathbf{z}_k^j} \lambda(\mathbf{z}|\mathbf{x}) \end{aligned} \tag{8.5}$$

if a target is present. Here, m_k is the number of measurements at time k, and $\mathbf{z}_k^j$ is the jth measurement in Z_k. The area $\mathcal{A}$ is some general area from which the measurements are obtained, usually corresponding to a validation gate [8, p. 263].

Proof. See Appendix I. □

8.3.2 Including wake clutter in the IPDA update

The reformulation of the measurement likelihood means that we need to reformulate the IPDA update as well.

Proposition 8.3. The update given a set Z_k of measurements at time k is given by

$$f_k(X_k|Z_{1:k}) = \begin{cases} 1 - r_k & \text{if } X_k = \emptyset \\ r_k \sum_{j=0}^{m_k} \beta_k^j f_k(\mathbf{x}|Z_{1:k}, a_k = j) & \text{if } X_k = \{\mathbf{x}\} \end{cases} \tag{8.6}$$

where a_k is the measurement assignment. The β-values are found as

$$\beta_k^j = \frac{p_k^j r_k^j}{r_k} \tag{8.7}$$

which are calculated using the association probabilities p_k^j, the association conditional existence probabilities r_k^j, and the existence probability r_k. The association probabilities describe the probability of the measurement assignment, and are given by

$$p_k^j = \begin{cases} r_{k|k-1} l_c^0 + (1 - r_{k|k-1}) \exp\Big(- \int_{\mathcal{A}} \lambda_0 \mathrm{d}\tilde{\mathbf{z}} \Big) \prod_{\mathbf{z} \in Z_k} \lambda_0 & \text{if } j = 0 \\ r_{k|k-1} l_c^j & \text{otherwise.} \end{cases} \tag{8.8}$$

For a missed detection have that

$$l_c^0 = \int f_{k|k-1}(\tilde{\mathbf{x}})(1 - P_D) \exp\Big(- \int_{\mathcal{A}} \lambda(\tilde{\mathbf{z}}|\tilde{\mathbf{x}}) \mathrm{d}\tilde{\mathbf{z}} \Big) \prod_{\mathbf{z} \in Z_k} \lambda(\mathbf{z}|\tilde{\mathbf{x}}) \mathrm{d}\tilde{\mathbf{x}}. \tag{8.9}$$

and for a detection we have that

$$l_c^j = l^j \int f_{k|k-1}(\tilde{\mathbf{x}}) P_D \exp\Big(- \int_{\mathcal{A}} \lambda(\tilde{\mathbf{z}}|\tilde{\mathbf{x}}) \mathrm{d}\tilde{\mathbf{z}} \Big) \prod_{\mathbf{z} \in Z_k \setminus \mathbf{z}_k^j} \lambda(\mathbf{z}|\tilde{\mathbf{x}}) \mathrm{d}\tilde{\mathbf{x}} \tag{8.10}$$

where

$$l^j = \int f_{\mathbf{z}}(\mathbf{z}_k^j|\tilde{\mathbf{x}}) f_{k|k-1}(\tilde{\mathbf{x}}) \mathrm{d}\tilde{\mathbf{x}}. \tag{8.11}$$

The association conditional existence probabilities are 1 for $j \geq 1$, that is, when conditioned on a detection, and

$$r_k^0 = \frac{r_{k|k-1} l_c^0}{p_k^0} \tag{8.12}$$

when conditioned on a missed detection. Finally, the existence probability is given by

$$r_k = p_k^0 r_k^0 + \sum_{j=1}^{m_k} p_k^j. \tag{8.13}$$

Proof. See Appendix II. □

8.4 Wake modeling

In [19, Section III], a method for mitigating the effect of wake clutter in a JPDA tracker by use of marginalization is presented. It is based on a non-parametric point of view, while we will use a parametric approach. The difference lies in how the values of λ_0 and λ_w are determined; through adaptive estimation (non-parametric) or tuning (parametric). The assumptions we make regarding the clutter distributions in Section 8.2 are nevertheless the same as in [19]. Furthermore, we present two different models for the clutter distribution of the measurements gated by the track. The first of these two models is also presented in [19] and is briefly summarized here. We denote the models as Model I and II, and when used with the WakeIPDA the resulting methods are denoted as WakeIPDA-1 and WakeIPDA-2.

Ultimately, what we want is to calculate (8.9) and (8.10). The first of these integrals is the clutter likelihood for missed detections, while the second is the clutter likelihood for detections. Because we also consider the existence probability, these expressions differ from the ones found in [19]. No analytical solutions exist for these integrals, but we can approximate them by use of, for example, importance sampling. We do, however, need expressions for $\lambda(\mathbf{z}|\mathbf{x})$ and $\exp(-\int_{\mathcal{A}} \lambda(\tilde{\mathbf{z}}|\mathbf{x}) \mathrm{d}\tilde{\mathbf{z}})$.

8.4.1 Model I

The first model, as it is defined in [19], assumes that all the edges of the wake area are outside the validation gate $\mathcal{G}$. Such a situation is shown in Figure 8.1. Based on this geometry, we need two things to calculate the clutter likelihoods. First, we need to know the area of the wake within the validation gate. Second, we need to

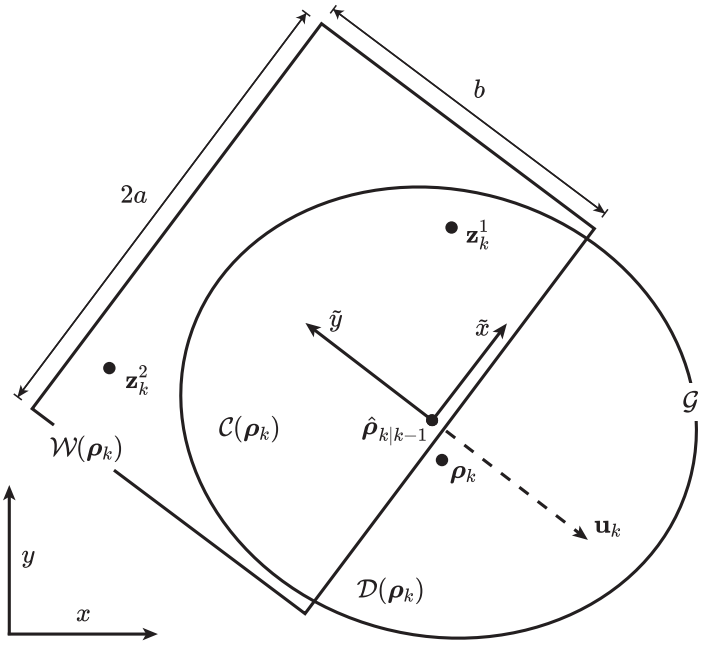

Figure 8.1: Illustration of the geometry of a validation gate and a wake area, with two measurements $\mathbf{z}_k^1$ and $\mathbf{z}_k^2$ in the wake area, of which one is inside and one outside the validation gate. The ellipsoid is the validation gate $\mathcal{G}$, and the rectangle is the wake area $\mathcal{W}(\boldsymbol{\rho}_k)$ of size $2a \times b$. The rectangle divides the validation gate into two areas of size $\mathcal{C}(\boldsymbol{\rho}_k)$ and $\mathcal{D}(\boldsymbol{\rho}_k)$, while $\boldsymbol{\rho}_k$ is the target position. Also shown is the wake-oriented coordinate system $(\tilde{x}, \tilde{y})$, the predicted target position $\hat{\boldsymbol{\rho}}_{k|k-1}$, and the target velocity vector $\mathbf{u}_k$.

know whether the measurements in the gate are inside or outside the wake area. By sampling from the prediction $f_{k|k-1}(\mathbf{x})$ we can get an estimate for the wake position $\boldsymbol{\rho}_k$, and by using a wake-oriented coordinate system $(\tilde{y}, \tilde{x})$ we can easily determine whether the measurements are inside the wake area or not. Details regarding the coordinate transformations can be found in [19]. We use the area of the wake within the validation gate to calculate the expected number of clutter measurements. The area is denoted as $\mathcal{C}(\boldsymbol{\rho}_k)$ in Figure 8.1, and is given in [19].

Using $\mathcal{C}(\boldsymbol{\rho}_k)$ we can calculate the expected number of clutter measurements as

$$\exp\Big(-\int_{\mathcal{G}} \lambda(\tilde{\mathbf{z}}|\mathbf{x})\mathrm{d}\tilde{\mathbf{z}}\Big) = \exp(-\lambda_w \mathcal{C}(\boldsymbol{\rho}_k)) \exp(-\lambda_0 V_{\mathcal{G}}) \tag{8.14}$$

where $V_{\mathcal{G}}$ denotes the area of the validation gate. Furthermore, we have that

$$\lambda(\mathbf{z}|\mathbf{x}) = \begin{cases} \lambda_w + \lambda_0 & \text{if } z_k^{\tilde{y}} \geq \rho_k^{\tilde{y}} \\ \lambda_0 & \text{otherwise} \end{cases}. \tag{8.15}$$

8.4.2 Model II

The second model assumes that all wake measurements originating from a target are gated by the track following said target. This can be enforced by gating an area around the predicted wake position in addition to the regular validation gate. In Figure 8.1 this would mean that the gated area is the union of the wake area $V_{\mathcal{W}}$ and the gate area $V_{\mathcal{G}}$, and that both measurements $\mathbf{z}_k^1$ and $\mathbf{z}_k^2$ would be gated. We assume that we know the length of the sides of the wake region $\mathcal{W}(\boldsymbol{\rho}_k)$, and the expected number of clutter measurements in the gate becomes

$$\exp\Big(- \int_{\mathcal{G}\cup\mathcal{W}} \lambda(\tilde{\mathbf{z}}|\mathbf{x}) \mathrm{d}\tilde{\mathbf{z}} \Big) \;=\; \exp(-\mathcal{D}(\boldsymbol{\rho}_k)\lambda_0) \exp(-V_{\mathcal{W}}(\lambda_w \,+\, \lambda_0)) \tag{8.16}$$

where we get the area $\mathcal{D}(\boldsymbol{\rho}_k)$ from [19]. Furthermore, we get the spatial likelihood as

$$\lambda(\mathbf{z}|\mathbf{x}) = \begin{cases} \lambda_w + \lambda_0 & \text{if } \rho_k^{\tilde{y}} + b \geq z_k^{\tilde{y}} \geq \rho_k^{\tilde{y}} \\ & \text{and } |z_k^{\tilde{x}}| \leq \rho_k^{\tilde{x}} + a \\ \lambda_0 & \text{otherwise} \end{cases} \tag{8.17}$$

where we again use the wake-oriented coordinate system to decide if a measurement is within the wake area.

Figure 8.2 shows the expected number of gated wake measurements compared to the actual number for the two models, when tested on simulated data. For WakeIPDA-1, the expected number is calculated by using $\exp(-\lambda_w \mathcal{C}(\boldsymbol{\rho}_k))$ from (8.14), while for WakeIPDA-2 it is calculated using $\exp(-\lambda_w V_{\mathcal{W}})$ from (8.16). For WakeIPDA-1, the expected and actual numbers are close for high detection probabilities, while the actual and expected number of wake measurements diverge as the detection probability decreases. The covariances of the estimates increase with decreasing detection probability, which results in increasing validation gate sizes and more expected wake measurements. For WakeIPDA-2, the expected number of wake measurements is independent of the validation gate size, due to the user defined wake area. Furthermore, Model II allows for the gating of all the wake measurements and thus the use of these for estimating target existence. When using Model I only a subset of the wake measurements are gated.

Both models require us to find some appropriate values for the parameters λ_w and λ_0. Deciding the regular clutter density parametrically by means of tuning is

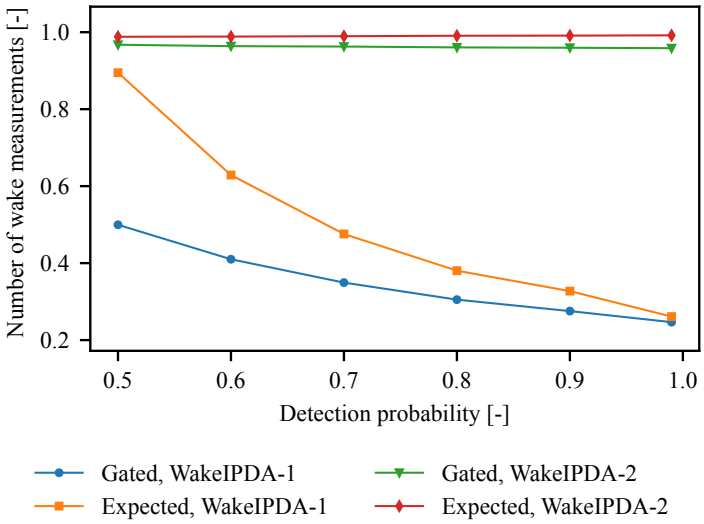

Figure 8.2: The expected number of gated wake measurements per time step, and the actual number of gated measurements, for the two models when used in combination with the method from Section 8.3. The values are averaged over all time steps, and plotted for a range of different detection probabilities. The actual number of wake measurements is, on average, one per time step.

shown to work well on experimental data in Chapter 2, and the possibility of using a parametric approach the wake density is discussed in [19]. In WakeIPDA-2 we also need to find values for the wake area, which can, for example, be based on average wake sizes [111].

8.5 Simulation setup and results

8.5.1 Simulation setup

We use simulated data sets to test the performance of the method. A discrete (nearly) constant velocity (CV) model [8, p. 274] with process noise $\sigma^2 = 0.35^2$ $\mathrm{m}^2/\mathrm{s}^4$ is used to model the movement of the targets. We model the measurements as having both polar and Cartesian noise. The polar noise represents the error in the range and bearing measurements, and the Cartesian noise represents other sources of error, such as clustering errors and receiver noise. We set the bearing std. dev. σ_θ to $1.0°$, the range std. dev. σ_r to 8.0 m, and the Cartesian std. dev. σ_xy to 6.6 m. We convert from polar to Cartesian coordinates using the method described in [95]. The clutter

densities are $\lambda_0 = 10^{-7}$ and $\lambda_w = 10^{-4}$, and we define a 100 by 100 meter large wake area. The gate size is set to $g = 4.0$.

We initialize new tracks on measurements that have not been gated by any track. Furthermore, we categorize tracks as either confirmed or preliminary. Tracks are confirmed when they have a probability of existence larger than a threshold $T_c = 0.99$ and are preliminary otherwise. Preliminary tracks are only allowed to associate with measurements not gated by any confirmed tracks. We terminate tracks as described in [143], with a termination threshold $T_t = 0.01$, and we set the survival probability P_S to 0.99. Lastly, we assume that the relevant densities are Gaussian, and we use a Kalman filter to estimate the target states.

The integrals in (8.9) and (8.10) are calculated using importance sampling with $N = 1000$ samples. This is done by sampling the distribution of the predicted state, and for each sample the values of (8.14) and (8.15), or (8.16) and (8.17) are calculated, depending on which model is used. The value of the integrals are then found by taking the mean of the values found by sampling. For a missed detection we have that

$$l_c^0 = (1 - P_D)\frac{1}{N}\sum_{i=1}^{N} \exp\Big(- \int_{\mathcal{A}} \lambda(\tilde{\mathbf{z}}|\mathbf{x}_i)\mathrm{d}\tilde{\mathbf{z}}\Big) \prod_{\mathbf{z}\in Z_k} \lambda(\mathbf{z}|\mathbf{x}_i) \tag{8.18}$$

and for a detection we have that

$$l_c^j = l^j P_D \frac{1}{N}\sum_{i=1}^{N} \exp\Big(- \int_{\mathcal{A}} \lambda(\tilde{\mathbf{z}}|\mathbf{x}_i)\mathrm{d}\tilde{\mathbf{z}}\Big) \prod_{\mathbf{z}\in Z_k \setminus \mathbf{z}_k^j} \lambda(\mathbf{z}|\mathbf{x}_i) \tag{8.19}$$

where $\mathbf{x}_i \sim f_{k|k-1}(\mathbf{x})$. When using Gaussian densities, l^j from (8.11) can be calculated analytacially.

8.5.2 Results

We obtained the results by running the algorithm on 2000 datasets, each containing a single moving target in a surveillance region with a radius of 1000 meters. The sensor, a radar, was placed in the center of the surveillance region. We assessed the performance according to the measures described in [29]. These are the track localization error (TLE), track probability of detection (TPD), track fragmentation rate (TFR), and track false alarm rate (TFAR).

In Figure 8.3 and Figure 8.4, we see the results for the WakeIPDA with the two wake models and for the regular IPDA. The results are shown for a range of detection probabilities. Note that the wake clutter is independent of the detection probability, and the expected number of wake measurements is constant across all scenarios. We see that both WakeIPDA variants perform better than the regular IPDA in most

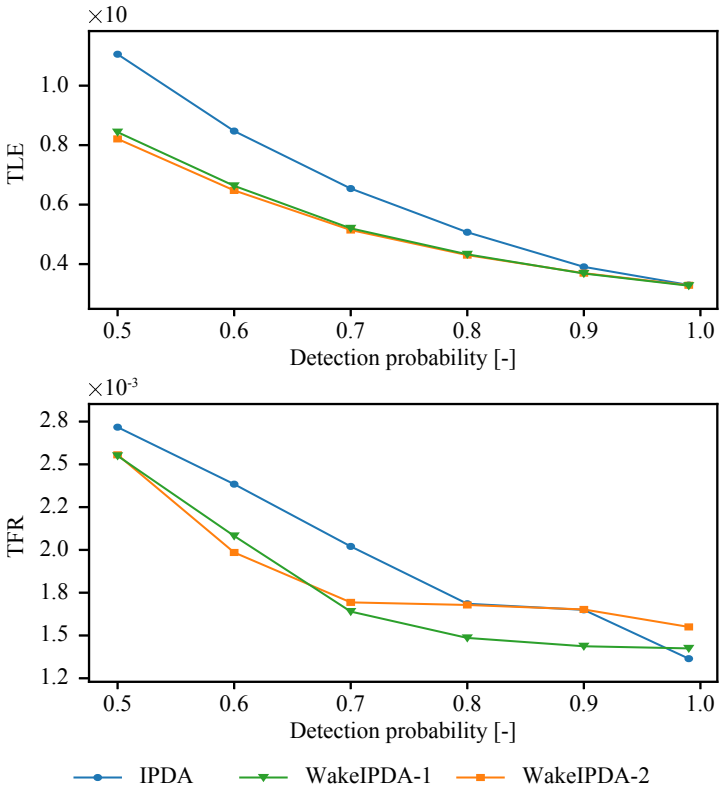

Figure 8.3: TLE and TFR values when tracking a single maneuvering target at varying probabilities of detection. The expected number of wake measurements is one per time step, irrespective of the detection probability. The results are averaged over 2000 datasets.

aspects and that the WakeIPDA-2 generally performs better than the WakeIPDA-1. The differences for low detection probabilities are the largest, likely due to how the IPDA can easily mistake the wake clutter measurements for detections when the target itself is not detected. Mistaking wake clutter for detections can degrade the estimate quality and may cause track loss. The WakeIPDA, on the other hand, is more robust to this error. However, we see that the WakeIPDA-1 has a higher TFAR than the other methods. We can find a cause for this in Figure 8.2; the large validation gates can lead to the tracker overestimating the existence probability of new tracks due to misclassification of regular clutter as wake clutter.

Figure 8.5 shows the target existence probability for the different methods. The left plot shows the existence probability averaged over the same datasets as

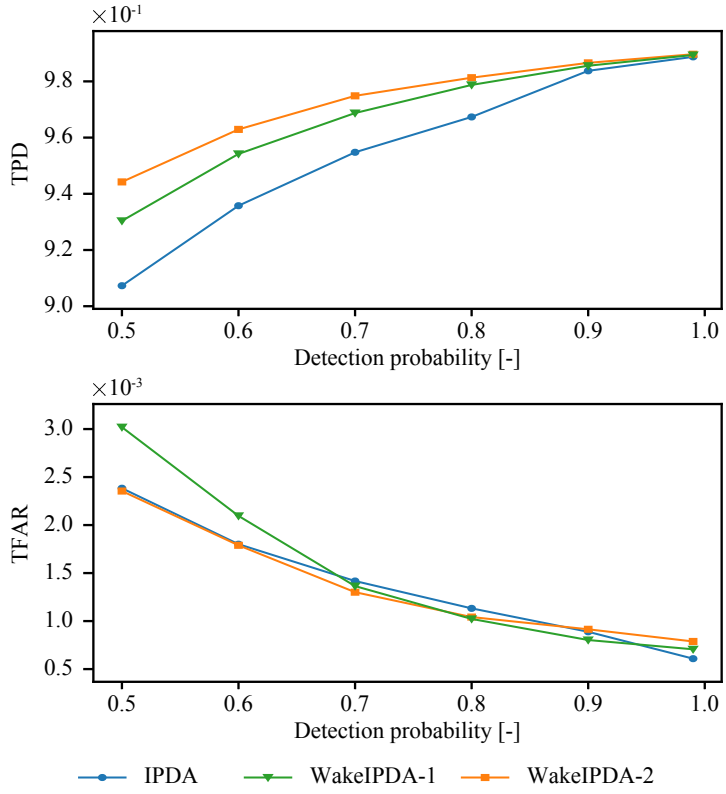

Figure 8.4: TPD and TFAR values when tracking a single maneuvering target at varying probabilities of detection. The expected number of wake measurements is one per time step, irrespective of the detection probability. The results are averaged over 2000 datasets.

above. The rightmost figure shows the existence probabilities for $P_D = 0.9$, at different times in the scenarios. The values are averaged over all datasets, and with a moving average taken over $N = 5$ time steps. We see that the WakeIPDA-2 has a markedly higher probability of existence than the other methods. Furthermore, the WakeIPDA-1 results in a slightly higher existence probability than the IPDA. The difference between the two methods is likely because the WakeIPDA-2 has a more accurate estimate for the number of wake measurements, resulting in a more precise estimate of the existence probability. Furthermore, both WakeIPDA versions result in higher existence probabilities in the early parts of the track duration. This effect is likely the cause of the TPD improvement of the WakeIPDA over the IPDA, see Figure 8.4, as the IPDA is more likely to fail to initialize a track.

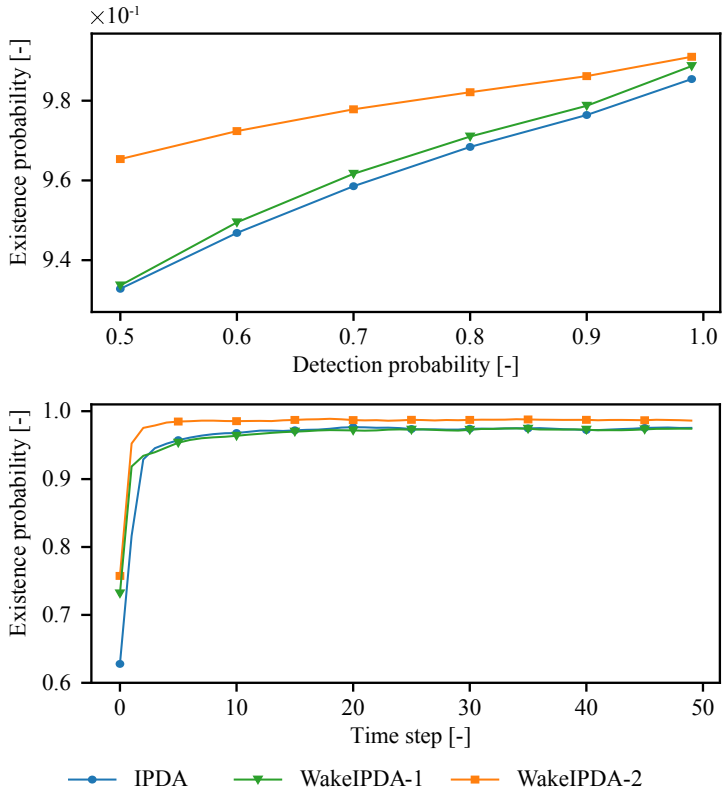

Figure 8.5: The probability of existence for the two WakeIPDA versions and the regular IPDA. The left plot shows the average probability of existence over 2000 datasets for a range of different detection probabilities. The right plot shows the moving average with $N = 5$ for $P_D = 0.9$, averaged over 2000 datasets.

8.6 Chapter summary

We have presented a new method, the WakeIPDA, which expands upon the single-target IPDA by modeling target-dependent wake clutter. We consider two different ways of modeling the wake clutter, and simulation results show that the WakeIPDA performs better than the regular IPDA when using either model. The models differ in that Model I rests on the assumption that all edges of the wake area are outside the validation gate of a track, while Model II uses the wake size directly. Model II performs best on the simulated data but is aided by specific knowledge regarding the wake measurements. On experimental data, the difference will likely be less clear-cut. Furthermore, by exploiting the presence of wake measurements, the WakeIPDA

applies a higher existence probability to the target than the IPDA, especially right after track initialization.

Further work may include extending the method to handle multiple targets and validating it on experimental data.

Appendix I

For the case where there is no target present, the measurement set is a Poisson RFS with intensity λ_0. The cardinality distribution is Poisson with mean $|Z| = \int_{\mathcal{A}} \lambda_0 \mathrm{d}\tilde{\mathbf{z}}$, and the individual measurements are independent and identically distributed (i.i.d.) according to $\lambda_0/|Z|$. Then, from [98, p. 366] we have

$$f(Z|X = \emptyset) = \Big(\prod_{\mathbf{z} \in Z_k} \lambda_0 \Big) \exp\Big(- \int_{\mathcal{A}} \lambda_0 \mathrm{d}\tilde{\mathbf{z}} \Big). \tag{8.20}$$

where we for ease of notation have omitted the time indices. When target dependent clutter may be present, we obtain a general expression for the measurement likelihood from [133, Proposition 1]. For a set Z of m measurements we have that

$$\begin{aligned} f(Z|X = \{\mathbf{x}\}) = (1 - P_D)\rho(|Z||\mathbf{x})|Z|! \prod_{\mathbf{z} \in Z} c(\mathbf{z}|\mathbf{x}) + \\ P_D \rho(|Z| - 1|\mathbf{x})(|Z| - 1)! \sum_{j=1}^{m} f_{\mathbf{z}}(\mathbf{z}^j|\mathbf{x}) \prod_{\mathbf{z} \in Z \setminus \mathbf{z}^j} c(\mathbf{z}|\mathbf{x}). \end{aligned} \tag{8.21}$$

The cardinality distibution $\rho(|Z||\mathbf{x})$ is Poisson

$$\rho(|Z||\mathbf{x}) = \frac{\Big(\int_{\mathcal{A}} \lambda(\tilde{\mathbf{z}}|\mathbf{x}) \mathrm{d}\tilde{\mathbf{z}} \Big)^{|Z|}}{|Z|!} \exp\Big(- \int_{\mathcal{A}} \lambda(\tilde{\mathbf{z}}|\mathbf{x}) \mathrm{d}\tilde{\mathbf{z}} \Big) \tag{8.22}$$

while the clutter measurements are i.i.d. as

$$c(\mathbf{z}|\mathbf{x}) = \frac{\lambda(\mathbf{z}|\mathbf{x})}{\int_{\mathcal{A}} \lambda(\tilde{\mathbf{z}}|\mathbf{x}) \mathrm{d}\tilde{\mathbf{z}}}. \tag{8.23}$$

By inserting (8.22) and (8.23) into the measurement likelihood we obtain

$$\begin{aligned} f(Z|X = \{\mathbf{x}\}) = \exp\Big(- \int_{\mathcal{A}} \lambda(\tilde{\mathbf{z}}|\mathbf{x}) \mathrm{d}\tilde{\mathbf{z}} \Big) \Big((1 - P_D) \\ \times \prod_{\mathbf{z} \in Z} \lambda(\mathbf{z}|\mathbf{x}) + P_D \sum_{j=1}^{m} f_{\mathbf{z}}(\mathbf{z}^j|\mathbf{x}) \prod_{\mathbf{z} \in Z \setminus \mathbf{z}^j} \lambda(\mathbf{z}|\mathbf{x}) \Big). \end{aligned} \tag{8.24}$$

This concludes our proof. □

Appendix II

The posterior we want to end up with is a Bernoulli mixture

$$f_k(X_k|Z_{1:k}) = \sum_{j=0}^{m_k} p_k^j f_k(X_k|Z_{1:k}, a_k = j). \tag{8.25}$$

where p_k^j are the weights of the mixture components, which sum to one. Bayes rule states that we can find the posterior target density according to

$$f_k(X_k|Z_{1:k}) = \frac{1}{N} f_k(Z_k|X_k) f_k(X_k|Z_{1:k-1}). \tag{8.26}$$

The Bernoulli RFS of a predicted track is defined as

$$f_k(X_k|Z_{1:k-1}) = \begin{cases} 1 - r_{k|k-1} & \text{if } X_k = \emptyset \\ r_{k|k-1} f_{k|k-1}(\mathbf{x}) & \text{if } X_k = \{\mathbf{x}_k\} \end{cases}. \tag{8.27}$$

By combining the above with the measurement likelihood from Appendix I we obtain

$$f_k(X_k|Z_{1:k}) = \frac{1}{N}(1 - r_{k|k-1}) \exp\Big(- \int_{\mathcal{A}} \lambda_0 \mathrm{d}\tilde{\mathbf{z}} \Big) \prod_{\mathbf{z} \in Z_k} \lambda_0(\mathbf{z}) \tag{8.28}$$

for the case when $X_k = \emptyset$. When $X_k = \{\mathbf{x}\}$ we have

$$\begin{aligned} f_k(X_k|Z_{1:k}) = \frac{1}{N} r_{k|k-1} f_{k|k-1}(\mathbf{x}) \exp\Big(- \int_{\mathcal{A}} \lambda(\tilde{\mathbf{z}}|\mathbf{x}) \mathrm{d}\tilde{\mathbf{z}} \Big) \times \\ \Big((1 - P_D) \prod_{\mathbf{z} \in Z_k} \lambda(\mathbf{z}|\mathbf{x}_k) + \sum_{j=1}^{m_k} P_D f_{\mathbf{z}}(\mathbf{z}_k^j|\mathbf{x}_k) \prod_{\mathbf{z} \in Z_k \setminus \mathbf{z}_k^j} \lambda(\mathbf{z}|\mathbf{x}_k) \Big). \end{aligned} \tag{8.29}$$

The measurement likelihood for an individual association hypothesis is given by

$$f_k(Z_k|\mathbf{x}_k, a_k = 0) = (1 - P_D) \exp\Big(- \int_{\mathcal{A}} \lambda(\tilde{\mathbf{z}}|\mathbf{x}) \mathrm{d}\tilde{\mathbf{z}} \Big) \prod_{\mathbf{z} \in Z} \lambda(\mathbf{z}|\mathbf{x}_k) \tag{8.30}$$

for a missed detection, and

$$f_k(Z_k|\mathbf{x}_k, a_k = j > 0) = P_D \exp\Big(- \int_{\mathcal{A}} \lambda(\tilde{\mathbf{z}}|\mathbf{x}) \mathrm{d}\tilde{\mathbf{z}} \Big) f_{\mathbf{z}}(\mathbf{z}_k^j|\mathbf{x}_k) \prod_{\mathbf{z} \in Z_k \setminus \mathbf{z}_k^j} \lambda(\mathbf{z}|\mathbf{x}_k) \tag{8.31}$$

for a detection. The posterior target density conditional on a single association hypothesis is

$$f_k(\mathbf{x}_k|Z_{1:k}, a_k = j) = \frac{f_k(Z_k|\mathbf{x}_{k|k-1}, a_k = j) f_{k|k-1}(\mathbf{x})}{l_c^j} \tag{8.32}$$

where

$$l_c^j = \int f_k(Z_k|\tilde{\mathbf{x}}, a_k = j) f_{k|k-1}(\tilde{\mathbf{x}}) \mathrm{d}\tilde{\mathbf{x}}. \tag{8.33}$$

Inserting this in Equation (8.29) we get

$$f_k(X_k|Z_{1:k}) = \frac{1}{N}\Big(r_{k|k-1} l_c^0 f_{k|k-1}(\mathbf{x}|Z_{1:k}, a_k = 0) + \sum_{j=1}^{m_k} r_{k|k-1} l_c^j f_k(\mathbf{x}|Z_{1:k}, a_k = i)\Big) \tag{8.34}$$

when $X_k = \{\mathbf{x}\}$. We now see that $p_k^j r_k^j \propto r_{k|k-1} l_c^j$. Furthermore, as $r_k^j = 1$ for $j > 0$, we have $p_k^j = r_{k|k-1} l_c^j / N$ for $j > 0$. Some more work is needed to find the existence probability r_k^0. We find it by dividing the unnormalized joint probability of a target existing and a missed detection occurring, by the probability of a missed detection. This gives

$$r_k^0 = \frac{r_{k|k-1} l_c^0}{r_{k|k-1} l_c^0 + (1 - r_{k|k-1}) \exp\Big(- \int_{\mathcal{A}} \lambda_0 \mathrm{d}\tilde{\mathbf{z}}\Big) \prod_{\mathbf{z} \in Z_k} \lambda_0(\mathbf{z})} \tag{8.35}$$

where the denominator is the sum of the probability of a missed detection occurring when a target exists, and a missed detection occurring when a target does not exist. Recognizing that the numerator in (8.35) is $p_k^0 r_k^0 \times N$, we find $p_k^0 \times N$ as the denominator. By remembering that $\sum_{j=0}^{m_k} p_k^j = 1$ we can find the normalization constant N.

Lastly, we want to have the final Bernoulli on the form

$$f_k(X_k|Z_{1:k}) = \begin{cases} 1 - r_k & \text{if } X_k = \emptyset \\ r_k \sum_{j=0}^{m_k} \beta_k^j \times f_k(\mathbf{x}_k|Z_{1:k}, a_k = j) & \text{if } X_k = \{\mathbf{x}_k\} \end{cases} \tag{8.36}$$

while we now have

$$f_k(X_k|Z_{1:k}) = \sum_{j=0}^{m_k} r_k^j p_k^j f_k(\mathbf{x}_k|Z_{1:k}, a_k = j) \text{ if } X_k = \{\mathbf{x}_k\}. \tag{8.37}$$

This is done by dividing the association probabilities and association-conditional existence with the posterior existence probability

$$\beta_k^j = \frac{p_k^j r_k^j}{r_k}. \tag{8.38}$$

Lastly, we find the posterior existence probability as

$$r_k = p_k^0 r_k^0 + \sum_{j=1}^{m_k} p_k^j \tag{8.39}$$

which leaves us with all the needed expressions. □

Chapter 9

Concluding Remarks

This thesis has presented ways of including target-provided information in target tracking, with a special focus on the application of AIS messages in maritime target tracking. It is also concerned with specific problems encountered when using exteroceptive sensors in the maritime domain. This chapter summarizes and discusses the contributions of the thesis and proposes possible future work.

9.1 Summary and discussion

Rather than summarizing the contributions of this thesis on a chapter-by-chapter basis, this section will instead connect the dots by considering concepts and problems encountered throughout.

Including target-provided information in target tracking

The single overarching theme of this thesis is the use of AIS messages in maritime target tracking, and this demands models and methods that can handle such information. A more general view of the problem is how to use target-provided information. The problem is not only relevant for maritime target tracking, but it is also applicable to, for example, air traffic control. In the context of measurement modeling, most messages a tracker can receive from surrounding targets share similar properties that distinguish them from exteroceptive sensor measurements: It is difficult to know when the messages will arrive, and the messages contain identity information, are not false alarms, and have position and velocity data gathered by the target. The fact that not all targets provide information to the tracker is also important. The measurement model should be able to model these properties rather than making similar assumptions as for exteroceptive measurements. Chapter 3 presents such a model, and a variant of the tracker from Chapter 2 is derived to account for the

target-provided information. The model upholds the mathematical properties needed for it to be applicable for use in the tracker from Chapter 2, and the results in both Chapter 3 and Chapter 5 show that the model works well in practice.

The trackers in Chapter 2, and by extension, the one in Chapter 3, are derived from the PMBM filter equations. Chapter 4 exploits this to create a PMBM tracker that can handle target-provided information. In addition to bringing the tracking performance for point targets close to what can be considered state-of-the-art, we also present a method for utilizing target-provided information in extended object tracking. When the identity of a target is known, its dimensions are often possible to obtain, either through the messages themselves or through a database. The results show that the method improves the extent estimate significantly by using this information. The results in Chapter 4 show that the PMBM variant improves upon the tracker from Chapter 3, especially in difficult situations.

In Chapter 5, the point target PMBM with multiple models and support for AIS messages is used in a closed-loop collision avoidance experiment. The collision avoidance method predicts the positions of the targets to a possible future point of contact, which makes the course and speed estimates especially important. A key motivation behind developing the trackers in Chapters 2–4 was experiences with poor course and speed estimates from simpler tracking methods. The uplifting results suggest that the PMBM variant is well-suited for such applications. Furthermore, the same field test applies the JIPDA variant from Chapter 3, and the results confirm that it also works well for its intended usage in autonomous marine collision avoidance.

The hybrid state formulation

The hybrid state, as formulated in [8, p. 411], is used extensively in the first part of this thesis. A hybrid state is in itself a simple concept; it is a state that includes both continuous and discrete elements. Using Chapter 2 as a reference point, the continuous part of the hybrid state contains the position and velocity of the target, whereas the discrete parts inform what kinematic model the target is following and whether or not it is visible. This way of looking at the target state allows both a structured way of developing the target tracking methods and simplifies changes to the target state. As shown in Section 3.4 and throughout Chapter 4, when thinking about the target state in such a way, the extraction of the individual elements of the target state becomes somewhat detached from the tracking method itself. That is, under the assumption that the hybrid state upholds any restrictions placed on the target state by the tracking method, the tracker need not be concerned with what specific elements hide in the hybrid state.

When Chapter 3 introduces the ID as a part of the hybrid state, and in the two trackers in Chapter 4 that use different state spaces, the hybrid state formulation both simplifies and enables. If one, for example, would want to include the color of a

boat as part of the state, this could be done without changing much more than the implementation details. Even though this way of thinking about the target state is more of an enabler than a contribution in itself, it highlights how the hybrid state formulation can be used to simplify the development of target tracking methods for specific applications in a more plug-and-play manner.

Why we should distrust AIS messages but still use them

For any sensor, it is important to understand its limitations. The problems encountered when using exteroceptive sensors are well known, with their false alarms, missed detections, and measurement noise. The limitations take on a different form for target-provided information in general and AIS messages specifically. First and foremost, the system that utilizes the information is no longer in control of the information gathering. This contrasts the exteroceptive sensors, where, for example, the radar is mounted on the ship that uses its measurements. In the case of AIS messages, the surrounding ships collect information to populate the different fields in the messages with no way for the receiver to know how this information was obtained or if the transmitter even tries to provide correct data. Furthermore, the information is transmitted over a wireless channel and is vulnerable to noise and interference.

This thesis addresses the problem of potentially incorrect AIS information in several ways. The IDs received from surrounding targets are not always assumed to be correct. Instead, the received IDs are modeled with a slight possibility of being incorrect, as is done in Chapter 3 and Chapter 4, and thus the tracker does not always see a new target in each new ID. However, the trackers themselves do not consider incorrect information in other message fields, except for inaccuracies in the transmitted position and velocity. These inaccuracies are modeled similarly to how the models assume some error in the exteroceptive measurements. Under the assumption that all the information the trackers want to use can be wrong, Chapter 6 provides a method to validate the information in the AIS messages. By use of exteroceptive sensors, simple hypothesis tests can detect errors. The validation helps avoid the initialization of tracks from AIS messages where no target is present and removes incorrect velocity data that otherwise would be used to update the target estimates. The information given by the AIS messages has to be confirmed before it is allowed to affect the tracker estimates.

Nevertheless, data from surrounding targets can be an invaluable source of information. For AIS, the messages contain the identity of the transmitting ship, which can be hard to obtain by other means. The identities can help the tracker with the data association problem and find additional uses in other parts of the larger system. Furthermore, the received measurements of position, course, and speed are usually more precise than what exteroceptive sensors can provide. Thus, the

target estimates can be improved by the AIS messages, in some cases significantly. The AIS messages also report turn rate information that allows the tracker to detect a maneuver quicker than otherwise possible. Lastly, Chapter 4 shows how the AIS messages can improve target extent estimates by using the dimensions of the transmitting ship.

How to improve tracking performance?

The methods presented in Chapters 2–4 try to improve tracking performance by adding to existing trackers. Chapter 2 shows how modeling the visibility of targets, in combination with more elaborate modeling of target kinematics, helps when encountering difficult situations in the maritime domain. The same tracker is extended to include target-provided information, which can help when measurements from exteroceptive sensors are too poor to provide precise estimates regardless of the tracker. The tracker in Chapter 4 expands the PMBM to include target-provided information, and it is shown how this can help when the tracker is faced with difficult situations or poor measurements.

However, what should not be overlooked is how the trackers rely on assumptions that will always be an approximation of the real world. The best possible tracker will perform poorly when based on incorrect assumptions. The methods also rely on further approximations when derived from the models, and how the compounding approximations impact the tracking performance is not easy to assess. Limited knowledge of the environment and the targets that inhabit it argues for a conservative approach where the models are kept simple. Chapter 7 dissects one of the common assumptions made when modeling radar measurements, in that all detections in a single scan are from the same point in time. With any rotating sensor that outputs its measurements after each full rotation, this is an approximation. The results show that the approximation is not necessarily a problem but can be in some situations. The chapter illustrates how questioning assumptions and approximations can help improve performance, whereas the use of a visibility state, multiple kinematic models, and target-provided information shows how adding functionality can help. However, the distinction is not necessarily clear-cut, as the goal of adding to an existing model is often to come closer to the real world. Allowing a target to become temporarily invisible to a sensor is closer to reality than assuming it is always visible.

The reliance on models means that real-world data is paramount. Such experimental data can reveal situations not accounted for in the models and inspire new developments. Many of the research problems encountered in this thesis were discovered when analyzing experimental data in the Autosea project [26], where the trackers often struggled to handle occlusions, maneuvering targets, and wakes. Testing the trackers on experimental data also provides a way to assess their usefulness in practical applications, where assumptions such as uniformly distributed clutter or

constant probability of detection can be unreasonable. The focus on real-world data is brought further in Chapter 5, which uses the trackers from Chapter 3 and Chapter 4 in their intended application: autonomous marine collision avoidance. Their success in this application strongly indicates their usefulness in the maritime domain, which is difficult to assess with simulated data adhering to their model assumptions.

The persistent problems of wake clutter and occlusions

Wake clutter showed up in most of the chapters in this thesis, either by explicit mention or through the problems they cause. In Chapter 2, it is encountered in the Joyride data set, where track losses occur when not using the visibility model. In Chapter 3, it is the reason for poor radar measurement quality when a large target ship, Gunnerus, makes a sharp turn. Here, the tracker is able to provide precise estimates only when aided by AIS messages. In Chapter 4, the wake of a target that is tracked with a LiDAR confuses the extended object tracker. When the wake is assumed to be a part of the target, its extent and position estimates become poor. Again, the AIS messages, through the target dimensions available through them, help the tracker avoid a blown-up extent estimate. When the trackers were tested in a closed-loop collision avoidance experiment in Chapter 5, the wake of the ownship caused problems, requiring us to discard all measurements close to the ship. Lastly, Chapter 8 presents a method for tracking targets in the presence of wakes. Here, rather than relying on general performance-enhancing measures to avoid problems caused by wakes, we instead went to the root of the problem and modeled the wakes themselves.

By modeling the wake clutter, the tracker can follow a target without being impeded by it. Furthermore, wake clutter measurement indicates that a target is the cause of the wake and can be used to estimate the target's existence probability. The results indicate that the information inherent in the wakes can help with target initialization and the continued estimation of its existence probability. However, the tracker is a single-target tracker and will have trouble with multiple closely spaced targets. Furthermore, a lack of available experimental data sets makes it difficult to evaluate the wake-compensating methods in a maritime environment. Nevertheless, awareness of the problem enables us to design experiments to provide data, and it can also help identify useful data sets in other contexts.

The problem of occlusions is also present in most of the chapters in this thesis. Occlusions occur when a target is hidden in some way that makes the exteroceptive sensor unable to detect it. The visibility state from Chapter 2 and target-provided measurements can mitigate the problem. It helps the target estimates behave reasonably when the tracker does not receive any detections for a prolonged time, and the target-provided measurements provide information about the target regardless of the occlusion. Even though the absence of measurements will deteriorate

any estimate somewhat, results throughout the thesis indicate that the remedies successfully contain the issues caused by occlusions.

9.2 Future work

Possible directions for future work have been hinted at in individual chapters and are summarized here. The availability of robust ways associating between exteroceptive sensors and target-provided messages provides the opportunity to use such messages in a wide range of applications. Chapter 4 introduces the use of AIS in extended object tracking, but more elaborate ways of exploiting the information likely exist. The AIS messages contain more information than what this thesis has considered, and the use of this information should be explored further. In other subject areas, such as coastal surveillance, air traffic control, and similar problems where the surveillance area is large, trackers that combine the target-provided information with exteroceptive sensor data can be beneficial. Furthermore, topics such as long-term prediction and collision avoidance can benefit from robust data association between AIS and exteroceptive measurements and the possibility of validating the AIS message information.

Regarding the trackers themselves, there exist ways of taking their performance even further. The trajectory PMBM filter [57] has shown itself to be a powerful method, and identity-carrying messages should intuitively work well in such a tracker. Furthermore, generalization of the wake modeling in Chapter 8 to multi-target trackers such as JIPDA and PMBM would increase its utility in practical applications. Such an extension is not trivial, however, because the sampling needed to compute the likelihoods may be prohibitive when increasing the number of association hypotheses. Furthermore, the number of associations themself will increase when we no longer assume that each target only creates at most one detection. Nevertheless, Chapter 8 can serve as a starting point for such an extension through its basis in [133]. Collecting data to evaluate the trackers in a wider range of situations is also important, and can help identify problems not yet considered.

Lastly, it can be viable to investigate the use of target-provided information to mitigate some of the more practical problems often encountered when using target tracking methods. Tuning of tracking methods can often be difficult and situation-dependent. Here, the target-provided information can provide a way to evaluate the accuracy of the exteroceptive sensors automatically. Knowledge of target extent through AIS messages can also help with clustering when the targets in an area are of varying sizes. AIS can also provide a minimum number of targets in an area, which can help avoid clustering detections from several targets together.

References

[1] F. Akpan, G. Bendiab, S. Shiaeles, S. Karamperidis, and M. Michaloliakos. Cybersecurity challenges in the maritime sector. *Network*, 2(1):123–138, Mar. 2022.

[2] T. Ardeshiri, K. Granström, E. Özkan, and U. Orguner. Greedy reduction algorithms for mixtures of exponential family. *IEEE Signal Processing Letters*, 22(6):676–680, Nov. 2015.

[3] A. Asada, F. Maeda, K. Kuramoto, Y. Kawashima, M. Nanri, and K. Hantani. Advanced Surveillance Technology for Underwater Security Sonar Systems. In *Proc. IEEE OCEANS*, Aberdeen, Scotland, June 2007.

[4] M. Athans, R. Whiting, and M. Gruber. A Suboptimal Estimation Algorithm with Probabilistic Editing for False Measurements with Applications to Target Tracking with Wake Phenomena. *IEEE Transactions on Automatic Control*, 22(3):372–384, June 1977.

[5] M. Baerveldt, A. G. Hem, and E. F. Brekke. Comparing Multiple Extended Object Tracking with Point Based Multi Object Tracking for LiDAR in a Maritime Context. *Journal of Physics: Conference Series*, 2618(1):012011, Oct. 2023.

[6] M. Baerveldt, M. E. López, and E. F. Brekke. Extended target PMBM tracker with a Gaussian Process target model on LiDAR data. In *Proc. IEEE 26th International Conference on Information Fusion*, Charleston, SC, USA, June 2023.

[7] Y. Bar-Shalom and X. R. Li. *Multitarget-Multisensor Tracking: Principles and Techniques*. YBS Publishing, Storrs, CT, USA, 1995.

[8] Y. Bar-Shalom, X. R. Li, and T. Kirubarajan. *Estimation with Application to Tracking and Navigation*. John Wiley & Sons, Inc., Hoboken, NJ, USA, 2001.

[9] Y. Bar-Shalom and E. Tse. Tracking in a Cluttered Environment with Probabilistic Data Association. *Automatica*, 11:451–460, Sept. 1975.

[10] Y. Bar-Shalom, P. K. Willett, and X. Tian. *Tracking and Data Fusion: A Handbook of Algorithms*. YBS Publishing, Storrs, CT, USA, 2011.

[11] M. Baum and U. D. Hanebeck. Shape tracking of extended objects and group targets with star-convex RHMs. In *Proc. IEEE 14th International Conference on Information Fusion*, Chicago, IL, USA, July 2011.

[12] E. P. Blasch, R. Niu, and S. O'Rourke. Target tracking analysis for stone soup. In *Proc. IEEE 23rd International Conference on Information Fusion*, Virtual conference, July 2020.

[13] D. D. Bloisi, F. Previtali, A. Pennisi, D. Nardi, and M. Fiorini. Enhancing Automatic Maritime Surveillance Systems With Visual Information. *IEEE Transactions on Intelligent Transportation Systems*, 18(4):824–833, Apr. 2017.

[14] H. Blom and E. Bloem. Combining interacting multiple model and joint probabilistic data association for tracking multiple maneuvering targets in clutter. Technical report, National Aerospace Laboratory NLR, Aug. 2002.

[15] H. A. P. Blom and Y. Bar-Shalom. The interacting multiple model algorithm for systems with markovian switching coefficients. *IEEE Transactions on Automatic Control*, 33(8):780–783, Aug. 1988.

[16] A. Bole, A. Wall, and A. Norris. Automatic Identification System (AIS). In *Radar and ARPA Manual*, pages 255–275. Elsevier, Amsterdam, Netherlands, 2014.

[17] A. Bole, A. Wall, and A. Norris. Basic Radar Principles. In *Radar and ARPA Manual*, pages 1–28. Elsevier, Amsterdam, Netherlands, 2014.

[18] E. Brekke, O. Hallingstad, and J. Glattetre. Target Tracking in State Dependent Wake Clutter. In *Proc. IEEE OCEANS*, Sydney, NSW, Australia, May 2010.

[19] E. Brekke, O. Hallingstad, and J. Glattetre. Improved target tracking in the presence of wakes. *IEEE Transactions on Aerospace and Electronic Systems*, 48(2):1005–1017, Apr. 2012.

[20] E. F. Brekke and M. A. Chitre. Relationship between finite set statistics and the multiple hypothesis tracker. *IEEE Transactions on Aerospace and Electronic Systems*, 54(4):1902–1917, Aug. 2018.

[21] E. F. Brekke, E. Eide, B.-O. H. Eriksen, E. F. Wilthil, M. Breivik, E. Skjellaug, Ø. K. Helgesen, A. M. Lekkas, A. B. Martinsen, E. H. Thyri, T. Torben, E. Veitch, O. A. Alsos, and T. A. Johansen. milliAmpere: An Autonomous Ferry Prototype. *Journal of Physics: Conference Series*, 2311(1):012029, July 2022.

[22] E. F. Brekke and A. G. Hem. A long simulation scenario for evaluation of multi-target tracking methods. In *Proc. 3rd International Conference on Electrical, Computer, Communications and Mechatronics Engineering (ICECCME)*, Tenerife, Canary Islands, Spain, July 2023.

[23] E. F. Brekke, A. G. Hem, and L.-C. N. Tokle. The VIMMJIPDA: Hybrid state formulation and verification on maritime radar benchmark data. In *Proc. IEEE OCEANS*, Virtual conference, Oct. 2020.

[24] E. F. Brekke, A. G. Hem, and L.-C. N. Tokle. Multitarget Tracking With Multiple Models and Visibility: Derivation and Verification on Maritime Radar Data. *IEEE Journal of Oceanic Engineering*, 46(4):1272–1287, July 2021.

[25] E. F. Brekke, B. Kalyan, and M. A. Chitre. A novel formulation of the Bayes recursion for single-cluster filtering. In *Proc. IEEE Aerospace Conference*, Big Sky, MT, USA, Mar. 2014.

[26] E. F. Brekke, E. F. Wilthil, B.-O. H. Eriksen, D. K. M. Kufoalor, Ø. K. Helgesen, I. B. Hagen, M. Breivik, and T. A. Johansen. The Autosea project: Developing closed-loop target tracking and collision avoidance systems. *Journal of Physics: Conference Series*, 1357(1):012020, Oct. 2019.

[27] K. Burnett, A. P. Schoellig, and T. D. Barfoot. Do we need to compensate for motion distortion and doppler effects in spinning radar navigation? *IEEE Robotics and Automation Letters*, 6(2):771–778, Jan. 2021.

[28] B. Chen and J. K. Tugnait. Tracking of multiple maneuvering targets in clutter using IMM/JPDA filtering and fixed-lag smoothing. *Automatica*, 37:239–249, Feb. 2001.

[29] S. Coraluppi, D. Grimmett, and P. D. Theije. Benchmark Evaluation of Multistatic Trackers. In *Proc. IEEE 9th International Conference on Information Fusion*, Florence, Italy, July 2006.

[30] D. F. Crouse, P. Willett, K. Pattipati, and L. Svensson. A look at Gaussian mixture reduction algorithms. In *Proc. IEEE 14th International Conference on Information Fusion*, Chicago, IL, USA, July 2011.

[31] K. Cutlip. Spoofing: One Identity Shared by Multiple Vessels. https://globalfishingwatch.org/data/spoofing-one-identity-shared-by-multiple-vessels/, July 2016. Accessed on: Jan. 7, 2024.

[32] E. d'Afflisio, P. Braca, and P. Willett. Malicious AIS spoofing and abnormal stealth deviations: A comprehensive statistical framework for maritime anomaly detection. *IEEE Transactions on Aerospace and Electronic Systems*, 57(4):2093–2108, Aug. 2021.

[33] J. Dezert and Y. Bar-Shalom. Joint probabilistic data association for autonomous navigation. *IEEE Transactions on Aerospace and Electronic Systems*, 29(4):1275–1286, Oct. 1993.

[34] A. Dunaeva, T. Vaule, and R. Olsen. High-resolution Survey on Versatile USV with Norbit Winghead. https://www.hydro-international.com/case-study/high-resolution-survey-on-versatile-usv, Nov. 2020. Accessed on: Dec. 27, 2023.

[35] B. H. Eriksen, E. F. Wilthil, A. L. Flåten, E. F. Brekke, and M. Breivik. Radar-based maritime collision avoidance using dynamic window. In *Proc. IEEE Aerospace Conference*, Big Sky, MT, USA, Mar. 2018.

[36] B.-O. H. Eriksen, E. F. Wilthil, A. L. Flaten, E. F. Brekke, and M. Breivik. Radar-based maritime collision avoidance using dynamic window. In *Proc. IEEE Aerospace Conference*, Big Sky, MT, Mar. 2018. IEEE.

[37] M. Estensen. Mariner USV, Aker BioMarine's new fishing drone will optimize fishing in Antarctica. https://www.akerbiomarine.com/news/mariner-usv-aker-biomarines-new-fishing-drone-will-optimize-fishing-in-antarctica, July 2023. Accessed on: Dec. 27, 2023.

[38] Euronews. Self-driving electric ferry takes to the seas around Stockholm. https://www.euronews.com/next/2023/06/12/self-driving-electric-commuter-ferry-takes-to-the-water-on-its-maiden-voyage-in-stockholm, June 2023. Accessed on: Dec. 27, 2023.

[39] Federal Aviation Administration. Automatic Dependent Surveillance— Broadcast (ADS–B) Out Performance Requirements To Support Air Traffic Control (ATC) Service, June 2010.

[40] G. B. Folland. *Real Analysis: Modern Techniques and Their Applications*. John Wiley & Sons, Inc., Hoboken, NJ, USA, 1999.

[41] T. E. Fortmann, Y. Bar-Shalom, and M. Scheffe. Sonar Tracking of Multiple Targets Using Joint Probabilistic Data Association. *IEEE Journal of Ocean Engineering*, 3:173–184, June 1983.

[42] D. Gaglione, P. Braca, and G. Soldi. Belief Propagation Based AIS/Radar Data Fusion for Multi - Target Tracking. In *Proc. IEEE 21st International Conference on Information Fusion*, Cambridge, UK, July 2018.

[43] D. Gaglione, P. Braca, G. Soldi, F. Meyer, A. G. Hem, E. F. Brekke, and F. Hlawatsch. Comments on "Variations of Joint Integrated Data Association with Radar and Target-Provided Measurements". *Journal of Advances in Information Fusion*, 18(2):93–101, Dec. 2023.

[44] D. Gaglione, P. Braca, G. Soldi, F. Meyer, F. Hlawatsch, and M. Win. Fusion of Sensor Measurements and Target-Provided Information in Multitarget Tracking. *IEEE Transactions on Signal Processing*, 70(1):322–336, Dec. 2021.

[45] Á. García-Fernández. MTT. Accessed on: Jan. 7, 2024.

[46] A. F. Garcia-Fernandez, A. S. Rahmathullah, and L. Svensson. A Metric on the Space of Finite Sets of Trajectories for Evaluation of Multi-Target Tracking Algorithms. *IEEE Transactions on Signal Processing*, 68:3917–3928, June 2020.

[47] A. F. Garcia-Fernandez, J. L. Williams, K. Granström, and L. Svensson. Poisson multi-Bernoulli mixture filter: Direct derivation and implementation. *IEEE Transactions on Aerospace and Electronic Systems*, 54(4):1883–1901, Aug. 2018.

[48] Á. F. García-Fernández, Y. Xia, and L. Svensson. Poisson Multi-Bernoulli Mixture Filter With General Target-Generated Measurements and Arbitrary Clutter. *IEEE Transactions on Signal Processing*, 71:1895–1906, May 2023.

[49] K. Gilholm, S. Godsill, S. Maskell, and D. Salmond. Poisson models for extended target and group tracking. In O. E. Drummond, editor, *Proc. Signal and Data Processing of Small Targets*, volume 5913, San Diego, CA, USA, Aug. 2005.

[50] I. R. Goodman, H. T. Nguyen, and R. Mahler. *Mathematics of Data Fusion*. Springer, Dordrecht, Netherlands, 1997.

[51] K. Granström, M. Baum, and S. Reuter. Extended Object Tracking: Introduction, Overview and Applications. *Journal of Advances in Information Fusion*, 12(2):139–174, Dec. 2017.

[52] K. Granström, M. Fatemi, and L. Svensson. Poisson multi-bernoulli mixture conjugate prior for multiple extended target filtering. *IEEE Transactions on Aerospace and Electronic Systems*, 56(1):208–225, Feb. 2020.

[53] K. Granström and U. Orguner. Estimation and maintenance of measurement rates for multiple extended target tracking. In *Proc. IEEE 15th International Conference on Information Fusion*, Singapore, Singapore, July 2012.

[54] K. Granstrom and U. Orguner. A phd Filter for Tracking Multiple Extended Targets Using Random Matrices. *IEEE Transactions on Signal Processing*, 60(11):5657–5671, Nov. 2012.

[55] K. Granström, S. Reuter, M. Fatemi, and L. Svensson. Pedestrian tracking using Velodyne data - Stochastic optimization for extended object tracking. In *Proc. IEEE Intelligent Vehicles Symposium*, Los Angeles, CA, USA, June 2017.

[56] K. Granström, S. Reuter, D. Meissner, and A. Scheel. A multiple model PHD approach to tracking of cars under an assumed rectangular shape. In *Proc. IEEE 17th International Conference on Information Fusion*, Salamanca, Spain, July 2014.

[57] K. Granström, L. Svensson, Y. Xia, J. Williams, and A. F. Garcia-Fernandez. Poisson multi-Bernoulli mixture trackers: Continuity through random finite sets of trajectories. In *Proc. IEEE 21st International Conference on Information Fusion*, Cambridge, UK, July 2018.

[58] G. Grimmett and D. Stirzaker. *Probability and Random Processes*. Oxford University Press, Oxford, UK, 2001.

[59] B. Habtemariam, R. Tharmarasa, M. McDonald, and T. Kirubarajan. Measurement level AIS/radar fusion. *Signal Processing*, 106(1):348–357, Jan. 2015.

[60] A. Harati-Mokhtari, A. Wall, P. Brooks, and J. Wang. Automatic identification system (AIS): Data reliability and human error implications. *The Journal of Navigation*, 60(3):373–389, Sept. 2007.

[61] Ø. K. Helgesen, E. F. Brekke, H. H. Helgesen, and Ø. Engelhardtsen. Sensor combinations in heterogeneous multi-sensor fusion for maritime target tracking. In *Proc. IEEE 22nd International Conference on Information Fusion*, Ottawa, Canada, July 2019.

[62] Ø. K. Helgesen, K. Vasstein, E. F. Brekke, and A. Stahl. Heterogeneous multi-sensor tracking for an autonomous surface vehicle in a littoral environment. *Ocean Engineering*, 252:111168, May 2022.

[63] A. G. Hem. Maritime multi-target tracking with radar and asynchronous transponder measurements. Master's thesis, NTNU, Jan. 2021.

[64] A. G. Hem, H.-G. Alvheim, and E. F. Brekke. WakeIPDA: Target Tracking With Existence Modeling in the Presence of Wakes. In *Proc. IEEE 26th International Conference on Information Fusion*, Charleston, SC, USA, June 2023.

[65] 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. To appear in *IEEE Access*. doi: 10.1109/ACCESS.2024.3389824.

[66] A. G. Hem and E. F. Brekke. Validation of AIS Information With Exteroceptive Sensor Fusion in Autonomous Operations. To appear in *IEEE Intelligent Transportation Systems Magazine*. doi: 10.1109/MITS.2024.3389869.

[67] A. G. Hem and E. F. Brekke. Compensating radar rotation in target tracking. In *Proc. Sensor Data Fusion: Trends, Solutions, Applications*, Bonn, Germany, Oct. 2022.

[68] A. G. Hem and E. F. Brekke. Variations of joint integrated data association with radar and target-provided measurements. *Journal of Advances in Information Fusion*, 17(2):97–115, Dec. 2022.

[69] A. G. Hem and E. F. Brekke. Variations of joint integrated data association with radar and target-provided measurements [source code], 2023. Accessed: Jan. 1, 2023.

[70] A. G. Hem, E. F. Brekke, G. D. K. M. Kufoalor, and I. H. Kingman. Autonomous Marine Collision Avoidance With Sensor Fusion of AIS and Radar. Submitted to the 15th IFAC Conference on Control Applications in Marine Systems, Robotics and Vehicles (CAMS 2024).

[71] International Association of Marine Aids to Navigation and Lighthouse Authorities (IALA). IALA guideline G1117: VHF data exchange system (VDES) overview, Dec. 2022. Accessed: Feb. 3, 2023.

[72] International Maritime Organization (IMO). Resolution A.1106(29), Revised guidelines for the onboard operational use of shipborne Automatic identification systems (AIS), Feb. 2015.

[73] International Telecommunication Union Radiocommunication Sector (ITU-R). Recommendation ITU-R M. 1371-5: Technical characteristics for an automatic identification system using time division multiple access in the VHF maritime mobile frequency band, Feb. 2014. Accessed: Feb. 12, 2023.

[74] International Telecommunication Union Radiocommunication Sector (ITU-R). Recommendation ITU-R M.585-9 - assignment and use of identities in the maritime mobile service, May 2022. Accessed: Feb. 12, 2023.

[75] C. Iphar, A. Napoli, C. Ray, E. Alincourt, and D. Brosset. Risk analysis of falsified automatic identification system for the improvement of maritime traffic safety. In *Proc. 26th European Safety and Reliability Conference*, Glasgow, UK, Sept. 2016.

[76] K. Ito, B. M. Nguyen, Y. Wang, M. Odai, H. Ogawa, E. Takano, T. Inoue, M. Koyama, H. Fujimoto, and Y. Hori. Fast and accurate vision-based positioning control employing multi-rate Kalman filter. In *Proc 39th Annual Conference of the IEEE Industrial Electronics Society*, Vienna, Austria, Nov. 2013.

[77] T. A. Johansen, T. Perez, and A. Cristofaro. Ship Collision Avoidance and COLREGS Compliance Using Simulation-Based Control Behavior Selection With Predictive Hazard Assessment. *IEEE Transactions on Intelligent Transportation Systems*, 17(12):3407–3422, Dec. 2016.

[78] S. Julier and J. K. Uhlmann. General Decentralized Data Fusion With Covariance Intersection (CI). In D. Hall and J. Llinas, editors, *Handbook of Data Fusion*, pages 319–344. CRC Press, Boca Raton, FL, USA, 2001.

[79] M. Kaess, H. Johannsson, R. Roberts, V. Ila, J. J. Leonard, and F. Dellaert. ISAM2: Incremental smoothing and mapping using the Bayes tree. *International Journal of Robotics Research*, 31(2):216–235, Feb. 2012.

[80] R. E. Kalman. A New Approach to Linear Filtering and Prediction Problems. *Journal of Basic Engineering*, 82(1):35–45, Mar. 1960.

[81] M. Kastek, R. Dulski, M. Zyczkowski, M. Szustakowski, P. Trzaskawka, W. Ciurapinski, G. Grelowska, I. Gloza, S. Milewski, and K. Listewnik. Multisensor system for the protection of critical infrastructure of a seaport. In *Proc. Unattended Ground, Sea, and Air Sensor Technologies and Applications XIV*, volume 8388, Rome, Italy, May 2012. SPIE.

[82] F. Katsilieris, P. Braca, and S. Coraluppi. Detection of malicious AIS position spoofing by exploiting radar information. In *Proc. IEEE 16th International Conference on Information Fusion*, Istanbul, Turkey, July 2013.

[83] J. W. Koch. Bayesian approach to extended object and cluster tracking using random matrices. *IEEE Transactions on Aerospace and Electronic Systems*, 44(3):1042–1059, July 2008.

[84] L. Kretschmann, H.-C. Burmeister, and C. Jahn. Analyzing the economic benefit of unmanned autonomous ships: An exploratory cost-comparison between an autonomous and a conventional bulk carrier. *Research in Transportation Business & Management*, 25:76–86, Dec. 2017.

[85] D. K. M. Kufoalor, T. A. Johansen, E. F. Brekke, A. Hepsø, and K. Trnka. Autonomous maritime collision avoidance: Field verification of autonomous surface vehicle behavior in challenging scenarios. *Journal of Field Robotics*, 37(3):387–403, Nov. 2019.

[86] D. K. M. Kufoalor, E. Wilthil, I. B. Hagen, E. F. Brekke, and T. A. Johansen. Autonomous COLREGs-Compliant decision making using maritime radar tracking and model predictive control. In *Proc. 18th European Control Conference*, Naples, Italy, June 2019.

[87] H. W. Kuhn. The Hungarian method for the assignment problem. *Naval research logistics quarterly*, 2(1-2):83–97, Mar. 1955.

[88] K. Lange. *Applied Probability*. Springer, New York, NY, USA, 2010.

[89] T. Larsen, N. Andersen, O. Ravn, and N. Poulsen. Incorporation of time delayed measurements in a discrete-time Kalman filter. In *Proc. IEEE Conference on Decision and Control*, Tampa, USA, Dec. 1998.

[90] D. Last, P. Thomas, S. Hiscocks, J. Barr, D. Kirkland, M. Rashid, S. B. Li, and L. Vladimirov. Stone Soup: Announcement of beta release of an open-source framework for tracking and state estimation. In *Proc. SPIE*, volume 11018, Baltimore, ML, USA, May 2019.

[91] G. Li, L. Kong, W. Yi, and X. Li. Multiple Model Poisson Multi-Bernoulli Mixture Filter for Maneuvering Targets. *IEEE Sensors Journal*, 21(3):3143–3154, Jan. 2021.

[92] M. Li, B. H. Kim, and A. I. Mourikis. Real-time motion tracking on a cellphone using inertial sensing and a rolling-shutter camera. In *Proc. IEEE International Conference on Robotics and Automation*, Karlsruhe, Germany, May 2013. IEEE.

[93] X. Li and Y. Bar-Shalom. Design of an interacting multiple model algorithm for air traffic control tracking. *IEEE Transactions on Control Systems Technology*, 1(3):186–194, Sept. 1993.

[94] E. Liland. AIS Aided Multi Hypothesis Tracker - Multi-Frame Multi-Target Tracking Using Radar and the Automatic Identification System. Master's thesis, NTNU, 2018.

[95] M. Longbin, S. Xiaoquan, Z. Yiyu, S. Z. Kang, and Y. Bar-Shalom. Unbiased converted measurements for tracking. *IEEE Transactions on Aerospace and Electronic Systems*, 34(3):1023–1027, July 1998.

[96] C. Lundquist, K. Granström, and U. Orguner. An Extended Target CPHD Filter and a Gamma Gaussian Inverse Wishart Implementation. *IEEE Journal of Selected Topics in Signal Processing*, 7(3):472–483, June 2013.

[97] R. Mahler. Multitarget Bayes filtering via first-order multitarget moments. *IEEE Transactions on Aerospace and Electronic Systems*, 39(4):1152–1178, May 2003.

[98] R. Mahler. *Statistical Multisource-Multitarget Information Fusion*. Artech House, Norwood, MA, USA, 2007.

[99] R. Mahler. PHD filters for nonstandard targets, I: Extended targets. In *Proc. IEEE 12th International Conference on Information Fusion*, Heidelberg, Germany, July 2009.

[100] R. Mahler. *Advances in Statistical Multisource-Multitarget Information Fusion*. Artech House, Norwood, MA, USA, 2014.

[101] Maritime Robotics. World's first uncrewed freight route at sea in the Trondheimsfjord. https://www.maritimerobotics.com/articles/uncrewed-freight-route-trondheimsfjord, Feb. 2023. Accessed on: Dec. 27, 2023.

[102] T. Miao, E. E. Amam, P. Slaets, and D. Pissoort. Multi-target tracking and detection, fusing RADAR and AIS signals using poisson multi-bernoulli mixture tracking, in support of autonomous sailing. In *Proc. International Naval Engineering Conference & Exhibition*, Delft, Netherlands, Oct. 2020.

[103] K. P. Murphy. *Machine Learning - a Probabilistic Perspective*. MIT Press, Cambridge, MA, USA, 2012.

[104] K. G. Murty. An algorithm for ranking all the assignments in order of increasing cost. *Operations research*, 16(3):682–687, May 1968.

[105] D. Musicki and R. Evans. Joint integrated probabilistic data association: JIPDA. *IEEE Transactions on Aerospace and Electronic Systems*, 40(3):1093–1099, July 2004.

[106] D. Musicki, R. Evans, and S. Stankovic. Integrated probabilistic data association (IPDA). In *Proc. IEEE 31st Conference on Decision and Control*, Tucson, AZ, USA, Dec. 1992.

[107] D. Musicki, R. Evans, and S. Stankovic. Integrated Probabilistic Data Association. *IEEE Transactions on Automatic Control*, 39(6):1237–1241, July 1994.

[108] D. Musicki and S. Suvorova. Tracking in clutter using IMM-IPDA-based algorithms. *IEEE Transactions on Aerospace and Electronic Systems*, 44(1):111–126, Jan. 2008.

[109] P. C. Niedfeldt, K. Ingersoll, and R. W. Beard. Comparison and Analysis of Recursive-RANSAC for Multiple Target Tracking. *IEEE Transactions on Aerospace and Electronic Systems*, 53(1):461–476, Feb. 2017.

[110] Norwegian University of Science and Technology. Autosit - NTNU. https://www.ntnu.edu/autosit. Accessed on: Jan. 7, 2024.

[111] A. T. Nylund, L. Arneborg, A. Tengberg, U. Mallast, and I.-M. Hassellöv. In situ observations of turbulent ship wakes and their spatiotemporal extent. *Ocean Science*, 17(5):1285–1302, Sept. 2021.

[112] R. S. of ITU. Technical characteristics for an automatic identification system using time division multiple access in the VHF maritime mobile frequency band. Technical Report ITU-R M.1371-5, International Telecommunication Union, Feb. 2014.

[113] A. S. Rahmathullah, A. F. Garcia-Fernandez, and L. Svensson. Generalized optimal sub-pattern assignment metric. In *Proc. IEEE 20th International Conference on Information Fusion*, Xian, China, July 2017.

[114] C. Ray, R. Gallen, C. Iphar, A. Napoli, and A. Bouju. DeAIS project: Detection of AIS spoofing and resulting risks. In *Proc. IEEE OCEANS*, Genova, Italy, May 2015. IEEE.

[115] D. Reid. An algorithm for tracking multiple targets. *IEEE Transactions on Automatic Control*, 24(6):843–854, Dec. 1979.

[116] M. A. Richards. *Fundamentals of Radar Signal Processing*. McGraw Hill Professional, New York, NY, USA, 2005.

[117] A. Rødningsby and Y. Bar-Shalom. Tracking of Divers in a Noisy Background Using a Bubble Model. In *Proc. SPIE*, volume 6699, San Diego, CA, USA, Aug. 2007.

[118] A. Rødningsby, Y. Bar-Shalom, O. Hallingstad, and J. Glattetre. Multitarget Tracking in the Presence of Wakes. In *Proc. IEEE 11th International Conference on Information Fusion*, Cologne, Germany, June 2008.

[119] A. Rødningsby, Y. Bar-Shalom, O. Hallingstad, and J. Glattetre. Multitarget Multisensor Tracking in the Presence of Wakes. *Journal of Advances in Information Fusion*, 4(2):117–145, June 2009.

[120] A. R. Runnalls. Kullback-Leibler approach to Gaussian mixture reduction. *IEEE Transactions on Aerospace and Electronic Systems*, 43(3):989–999, July 2007.

[121] J. Å. Sagild, A. G. Hem, and E. F. Brekke. Counting technique versus single-time test for track-to-track association. In *Proc. IEEE 24th International Conference on Information Fusion*, Sun City, South Africa, Nov. 2021.

[122] D. Schuhmacher, B. T. Vo, and B.-N. Vo. A consistent metric for performance evaluation of multi-object filters. *IEEE Transactions on Signal Processing*, 56:3447–3457, Aug. 2008.

[123] M. Schuster, M. Blaich, and J. Reuter. Collision Avoidance for Vessels using a Low-Cost Radar Sensor. *IFAC Proceedings Volumes*, 47(3):9673–9678, Apr. 2014. 19th IFAC World Congress.

[124] G. Soldi, D. Gaglione, S. Raponi, N. Forti, E. d'Afflisio, P. Kowalski, L. M. Millefiori, D. Zissis, P. Braca, P. Willett, A. Maguer, S. Carniel, G. Sembenini, and C. Warner. Monitoring of Critical Undersea Infrastructures: The Nord Stream and Other Recent Case Studies. *IEEE Aerospace and Electronic Systems Magazine*, 38(10):4–24, Oct. 2023.

[125] R. Streit. *Poisson Point Processes: Imaging, Tracking, and Sensing*. Springer, New York, NY, USA, 2010.

[126] R. Streit. How I learned to stop worrying about a thousand and one filters and love analytic combinatorics. In *Proc. IEEE Aerospace Conference*, Big Sky, MT, USA, Mar. 2017.

[127] H. A. Tran, T. A. Johansen, and R. R. Negenborn. Collision avoidance of autonomous ships in inland waterways – A survey and open research problems. *Journal of Physics: Conference Series*, 2618(1):012004, Oct. 2023.

[128] Trondheim Havn. Tavern III. https://commons.wikimedia.org/wiki/File:Tavern_-III_%286242560007%29.jpg. Accessed on: Oct. 27, 2023.

[129] D. Vivet, P. Checchin, and R. Chapuis. Radar-only localization and mapping for ground vehicle at high speed and for riverside boat. In *Proc. IEEE International Conference on Robotics and Automation*, St Paul, MN, USA, May 2012.

[130] B.-n. Vo, M. Mallick, Y. Bar-shalom, S. Coraluppi, R. Osborne III, R. Mahler, and B.-t. Vo. Multitarget Tracking. In *Wiley Encyclopedia of Electrical and Electronics Engineering*, pages 1–15. John Wiley & Sons, Ltd, Hoboken, NJ, USA, 2015.

[131] B.-N. Vo, S. Singh, and A. Doucet. Sequential Monte Carlo methods for multitarget filtering with random finite sets. *IEEE Transactions on Aerospace and Electronic Systems*, 41(4):1224–1245, Oct. 2005.

[132] B.-T. Vo and B.-N. Vo. Labeled Random Finite Sets and Multi-Object Conjugate Priors. *IEEE Transactions on Signal Processing*, 61(13):3460–3475, July 2013.

[133] B.-T. Vo, B.-N. Vo, and A. Cantoni. Bayesian Filtering With Random Finite Set Observations. *IEEE Transactions on Signal Processing*, 56(4):1313–1326, Apr. 2008.

[134] N. Wahlstrom and E. Ozkan. Extended Target Tracking Using Gaussian Processes. *IEEE Transactions on Signal Processing*, 63(16):4165–4178, Aug. 2015.

[135] A. Wald. Sequential tests of statistical hypotheses. *The Annals of Mathematical Statistics*, 16(2):117–186, June 1945.

[136] X. Wang, S. Challa, and R. Evans. Gating techniques for maneuvering target tracking in clutter. *IEEE Transactions on Aerospace and Electronic Systems*, 38(3):1087–97, July 2002.

[137] X. Wang and D. Musicki. Low elevation sea-surface target tracking using IPDA type filters. In *Proc. International Radar Conference*, Adelaide, SA, Australia, Sept. 2003.

[138] J. Williams. Marginal multi-Bernoulli filters: RFS derivation of MHT, JIPDA, and association-based MeMBer. *IEEE Transactions on Aerospace and Electronic Systems*, 51(3):1664–1687, July 2015.

[139] J. L. Williams. Hybrid Poisson and multi-Bernoulli filters. In *Proc. IEEE 15th International Conference on Information Fusion*, Singapore, Singapore, July 2012.

[140] J. L. Williams. An efficient, variational approximation of the best fitting multi-Bernoulli filter. *IEEE Transactions on Signal Processing*, 63(1):258–273, Jan. 2015.

[141] E. F. Wilthil, Y. Bar-Shalom, P. Willett, and E. F. Brekke. Estimation of Target Detectability for Maritime Target Tracking in the PDA Framework. In *Proc. IEEE 22nd International Conference on Information Fusion*, Ottawa, Canada, July 2019.

[142] E. F. Wilthil, E. F. Brekke, and O. B. Asplin. Track Initiation for Maritime Radar Tracking with and without Prior Information. In *Proc. IEEE 21st International Conference on Information Fusion*, Cambridge, UK, July 2018.

[143] E. F. Wilthil, A. L. Flåten, and E. F. Brekke. A target tracking system for ASV collision avoidance based on the PDAF. In T. I. Fossen, K. Y. Pettersen, and H. Nijmeijer, editors, *Sensing and Control for Autonomous Vehicles*, volume 474, pages 269–288. Springer, Ålesund, Norway, 2017.

[144] K. Wolsing, L. Roepert, J. Bauer, and K. Wehrle. Anomaly detection in maritime AIS tracks: A review of recent approaches. *Journal of Marine Science and Engineering*, 10(1):1–19, Jan. 2022.

[145] Y. Xia. *Poisson Multi-Bernoulli Mixtures for Multiple Object Tracking*. PhD thesis, Chalmers Tekniska Hogskola (Sweden), 2022.

[146] Y. Xia, K. Granström, L. Svensson, A. F. Garcia-Fernandez, and J. L. Williams. Multiscan implementation of the trajectory Poisson multi-Bernoulli mixture filter. *Journal of Advances in Information Fusion*, 14(2):213–235, Nov. 2019.

ISBN 978-82-326-8028-3 (printed ver.)
ISBN 978-82-326-8027-6 (electronic ver.)
ISSN 1503-8181 (printed ver.)
ISSN 2703-8084 (online ver.)

NTNU

Norwegian University of
Science and Technology