NTNU

Norwegian University of
Science and Technology

Near-shore target tracking with clutter maps

Oskar Buset Asplin

Master of Science in Cybernetics and Robotics
Submission date: June 2018
Supervisor: Edmund Førland Brekke, ITK
Co-supervisor: Erik Wilthil, ITK

Norwegian University of Science and Technology
Department of Engineering Cybernetics

Problem description

Autonomous ships and autonomous surface vehicles (ASVs) in general need an automatic collision avoidance (COLAV) system for safe navigation. A key component of the COLAV system is a multi-target tracking method which estimates positions and velocities of other moving objects, referred to as targets. A key challenge is track initialization: The tracking method must make decisions regarding whether or not a sequence of radar measurements come from a target, or whether they are to be discarded as clutter. In the real world, clutter is highly non-stationary, and it tends to be particularly bothersome near land.

This project will build on a previous 5th year specialization project where the conventional M-out-of-N method was compared with the Integrated Probabilistic Data Association (IPDA) for track initialization. In this Master's thesis the goal will be to further adapt the IPDA for optimal performance when zones with high clutter densities are present. The project involves the following tasks:

1. Test a plain vanilla IPDA on real radar data recorded in the Autosea project.
2. Implement a clutter map and discuss strategies for determining/estimating the clutter density.
3. Use the clutter map to provide the clutter density used in the IPDA.
4. Analyze and discuss choices for confirmation and termination thresholds in the IPDA with clutter map.
5. Propose different performance measures for track initiation and discuss their suitability, with particular focus on their usefulness in a COLAV system.

Preface

This thesis was written as the final submission of the master's degree at the Department of Engineering Cybernetics at NTNU. I would like to thank my two supervisors, Edmund Brekke and Erik Wilthil, for their great guidance and support in the work of this thesis. Thanks to their excellent counselling I have been introduced to the field of tracking and assisted in all its terminology and methodology. A special thanks to Erik Wilthil for including and integrating me in the Autosea project and for introducing me to their current system framework.

The Autosea project provided me with a working tracking system of an implemented PDA and M/N as only track initiation method, written in Python. The IPDA integration was done in the specialization project in 2017. I was also provided with a commenced draft of a clutter map module. Parts of the work in the thesis were also reported in the conference paper in Appendix B. In this paper, I was responsible for the implementation and testing of the IPDA.

The academic cooperation and social atmosphere in room G234 was highly rewarding and invigorating during the work of this thesis.

Abstract

Target tracking is of vital importance for autonomous vehicles moving in areas with unpredictable traffic. The radar is an important sensor in target tracking for autonomous vessels, though, it suffers from occasional false alarms, which may be particularly troublesome near land. Track initiation methods have different approaches to filtering the clutter and detecting true targets among the measurements. Integrated Probabilistic Data Association (IPDA) is a method suitable for both data association and track initiation, though its assumption of uniform clutter is not fully adaptable to many real world situations. Near-shore environments tend to have an uneven clutter distribution at the same time as they are the subject to targets leaving shore, making efficient track initiation a difficult task. Clutter maps attempt to estimate clutter density to assist in the process of selecting or weighting measurements associated with a target.

Three clutter estimators were thoroughly tested and compared, which given some initialization time, improved the efficiency and accuracy of the IPDA in simulations. The Spatial estimator shows the best results for a shorter initialization length of 20 scans, while the Classic and Temporal estimators are able to achieve higher true detection rates when given longer initialization time. The Spatial estimator performs the best in simulated uniform clutter which correlates with its theoretical bias and quick convergence. The tests from real data sets recorded in Trondheimsfjorden suggest that the Classic estimator is most suitable for tracking scenarios where the radar is on board the ownship. For a stationary radar setup, which is more comparable to the simulations in this thesis, the Temporal estimator is a logical choice due to no theoretical bias when given sufficient estimation time. Excluding estimated target measurements from clutter estimation shows a positive effect on track initiation time, and should be tested and explored in further testing of clutter maps.

Sammendrag

Målfølging er kritisk for autonome kjøretøyer som beveger seg i områder med uforutsigbar trafikk. Radaren er en viktig sensor i målfølging for autonome fartøy, men den lider av sporadiske falske alarmer, noe som kan være spesielt problematisk nærme land. Målinitieringsmetoder har ulike tilnærminger til å filtrere clutteret og detektere sanne mål blant målingene. Integrated Probabilistic Data Association (IPDA) er en metode som passer for både dataassosiasjon og målinitiering, men antagelsen om uniformt clutter stemmer ikke helt overens med mange virkelige situasjoner. Områder nærme land har en tendens til å ha en ujevn clutterfordeling samtidig som det kan dukke opp mål som forlater kysten, noe som gjør effektiv målinitiering til en vanskelig oppgave. Clutterkart forsøker å estimere cluttertetthet for å bistå i prosessen med å plukke ut eller vekte målinger assosiert med et mål.

Tre clutterestimatorer ble grundig testet og sammenlignet, og gitt noe tid til initialisering, forbedret de effektiviteten og nøyaktigheten til en IPDA i simuleringer. Den romlige estimatoren viser de beste resultatene for en kortere initialiseringslengde på 20 skanninger, mens den klassiske og den temporale estimatoren kan oppnå høyere sanne gjenkjenningsrater når de får lengre initialiseringstid. Den romlige estimatoren gir de beste resultatene i simulert enhetlig clutter som korrelerer med dens teoretiske bias og raske konvergens. Tester fra ekte datasett registrert i Trondheimsfjorden antyder at den klassiske estimatoren er best egnet for målfølgings-scenarier hvor radaren er ombord på det autonome fartøyet. For et stasjonært radaroppsett, som er mer sammenlignbart med simuleringene i denne oppgaven, er den temporale estimatoren et logisk valg på grunn av null teoretisk bias når den blir gitt tilstrekkelig estimeringstid. Å ekskuldere estimerte målmålinger fra clutterestimering viser en positiv effekt på målinitieringstid, og bør testes og utforskes ved ytterligere testing av clutterkart.

Table of Contents

Acronyms

AIS	Automatic Identification System
ASV	Autonomous Surface Vehicle
CDF	Cumulative Distribution Function
COLAV	Collision Avoidance
GPS	Global Positioning System
IMM-PDA	Interacting Multiple Model Probabilistic Data Association
IPDA	Integrated Probabilistic Data Association
MHT	Multiple Hypothesis Tracking
NCV	Nearly Constant Velocity
NED	North East Down
NIS	Normalized Innovation Squared
PDA	Probabilistic Data Association
ROC	Receiver Operating Characteristic
SOC	System Operating Characteristic
SOI	Spread of Innovations
SPRT	Sequential Probability Ratio Test
TTP	True Track Probability

Glossary

Autosea	A collaborative research and development project between NTNU, DNV GL, Kongsberg and Maritime Robotics with aim to improve design and verification of control systems for ASVs
Clustering	Creating a shared measurement for several close distance measurements
Clutter	Noise in the form of false measurements
Gate	An area of the scan which is assumed to be affiliated with a certain track
Innovation	The distance from a measurement to the predicted measurement
Measurement	A location in the measurement space where something is detected
Ownship	Ownship is the vehicle or platform which the sensor is mounted on
Python	An object-oriented programming language
Scan	A limited surveillance area surrounding the radar position in which it receives detections
Target	An object with unknown system state which we want to track
Track	A sequence of measurements which is possibly from the same target
Tracking	The process of continuously estimating the system state of a target

Nomenclature

$(\hat{\cdot})_{k\|k}$	Estimated variable
$(\hat{\cdot})_{k\|k-1}$	Predicted variable
β	Association weights
c	Clutter map cell
χ	Course angle
$\boldsymbol{e}$	Estimation error
$\boldsymbol{F}$	State-transition model
$f(\boldsymbol{z})$	Probability density of the predicted measurement position
γ_G	Gate threshold
$\boldsymbol{H}$	Observation model
i	Measurement index
$\boldsymbol{K}$	Kalman gain
k	Scan index
L	Window length for clutter map estimators
λ	Clutter density
μ_k	Inverse clutter density measurement at scan k
N	Window size (averaging length) of a clutter map estimator
N_k	The averaged value of inverse clutter density
$\boldsymbol{\nu}$	Innovation
$\boldsymbol{P}$	Error covariance matrix

P_D	Detection probability
P_{DT}	True track detection probability
P_{FT}	False track probability
P_G	Gate probability
$\tilde{\boldsymbol{P}}_k$	Spread of innovations at scan k
$P\{\boldsymbol{x}_k\}$	Probability of existence at scan k
$\boldsymbol{Q}$	Covariance of the process noise
$\boldsymbol{R}$	Covariance of the observation noise
$\boldsymbol{S}$	Covariance of the innovation
σ_q	Process noise strength
σ_r	Measurement noise strength
T	Time period between scans
τ_k	Inter-arrival time at index k of measurement arrival
$\overline{\tau}_k$	The averaged value of inter-arrival times at index k
$V(c)$	Volume of clutter map cell c
V_k	Area of the validation region at scan k
$\boldsymbol{x}$	True system state
$\boldsymbol{z}_k$	Measurement at scan k
$\hat{\boldsymbol{z}}_k$	Predicted measurement at scan k
$\boldsymbol{Z}^k$	Set of measurements falling in the validation region up to and including scan k

Chapter 1

Introduction

1.1 Motivation and background

Target tracking is of key importance for autonomous surface vehicles (ASVs) moving in environments with traffic. Autonomous vehicles need to keep track of other moving objects in their vicinity to avoid any collisions. We refer to the objects we wish to keep track of as targets. Target tracking entails processing data from one or several exteroceptive sensors, such as radar, lidar or camera, to monitor the surroundings. The radar is especially suitable for object detection at sea, however it suffers from noise due to external conditions such as weather and land which can lead to false measurements. A significant amount of noise may be especially apparent in near-shore environments [20]. Returns from land stand for the majority of these measurements. Most of the returns from land are removed, but due to map inaccuracies, some measurements may be apparent from land.

An autonomous vessel needs to keep track of its surroundings in order to safely navigate through an area without any collisions or dangerous situations. Other moving objects like ships or small boats are especially important to keep track of as they may show unpredictable behaviour in terms of velocity and direction. Desired information about other ships can either be communicated between ships, or it has to be estimated through sensors aboard the ownship. One widespread way of communicating ship information with other ships and coastal authorities is called Automatic Identification System (AIS) [8]. All ships of a certain size or type are required to use AIS which communicates information about the ship such as position, course and speed. However, AIS is not required for all sea vessels and may also not be completely reliable in all situations. For this reason, having a way of tracking the surrounding ships and other objects using exteroceptive sensors is essential for a collision avoidance (COLAV) system at sea. Such an estimation process of other targets system state is referred to as target tracking.

A sequence of measurements associated with a possible target is called a track,

and a track can be either true or false depending on if there is a real target at that location. Returns from land or other measurements considered as noise can cause a target tracking system to initiate false tracks, which can cause problems for an autonomous vehicle navigating through these areas. Unwanted measurements in a radar scan are referred to as clutter. Measurements from real targets moving close to shore may be associated with measurements from land or other types of clutter, resulting in the real targets not being discovered. Clutter should therefore in some way be handled in target tracking to achieve a controlled tracking situation.

A tracking system needs a way of determining initial states for potential targets in the measurement area. A proper initiation method must be able to filter false measurements from measurements originating from true targets in order to initiate correct tracks. Initiation of a track based on false measurements may lead to unwanted navigation and behaviour from an autonomous vehicle. On the other hand, not detecting a true target may lead to a collision with the target.

The process of filtering out false measurements from a set of real world measurements is a difficult task. Some methods and algorithms attempts to do so, while others take all measurements into account and weight them differently according to the likelihood that they originate from a target. In both cases, knowing the spatial distribution and frequency of false measurements may be beneficial. Some methods rely on an assumption that the false measurements are uniformly spread, while others, such as clutter maps [17], attempt to map noise densities in several regions of the map.

A map of noise estimates across the measurement area may not be necessary far out in the middle of the ocean, but in areas close to shore there might exist considerable discontinuities in the amount of noise returns. Consequently, combining a clutter map with a tracking system can improve the performance of the system while operating in near-shore environments. Areas close to shore are also more likely to spawn new targets as ships or small boats leave dock. This is a complicating factor considering that turning up the estimated clutter density near land to filter more noise may also lead to undetected new targets from such areas. If measurements from a real target are classified as clutter we might get dangerous situations for autonomous ships relying on proper detection techniques of other surrounding vessels. Accurate and efficient track initiation is thus of vital importance in such areas to maintain a secure COLAV system.

1.2 Previous work

There exists a number of tracking methods and techniques which can be used in maritime COLAV systems. They differ in scope and area of specialization, as they attempt to solve different challenges of autonomous vehicle movement in maritime environments. Tracking systems are generally divided into two catogeries; single-target tracking, and multi-target tracking. Single-target tracking [2, 15] attempts to track a single target at a time and associate measurements to that target, while

multi-target tracking [14, 18, 9] considers all potential targets at a given time simultaneously. Associating measurements to a target (data association) is relevant to both single-target tracking and multi-target tracking. Some methods like nearest-neighbor [6] and Multiple Hypothesis Tracking (MHT) [18] single out one measurement for the target as a basis for estimating the target state, while others like Probabilistic Data Association (PDA) [10] use several nearby measurements and merge their contribution to a single estimate. PDA is a data association filter based on a Kalman filter [4] approach for estimating the state of a target. PDA has been a proclaimed and much used method since its introduction, and has since been extended to other more complex methods such as Integrated PDA (IPDA) [15], Joint PDA (JPDA) [2] and Interacting Multiple Model PDA (IMM-PDA) [2].

Initiating tracks precisely and efficiently is also an important part of a COLAV system, with several different approaches and existing methods. Popular track initiation methods such as M/N [2], IPDA [15], IMM-PDA [2], SPRT [9] and Perceivability [12] have been used together with the PDA to make a complete tracking system. Of these methods, all except M/N tries in some way to estimate detection probability or existence probability for a possible track to determine whether to initiate a track or not. M/N and IPDA were tested and compared in the specialization project [1] where it was demonstrated that IPDA achieves significantly better results in high clutter areas.

Some track initiation methods, such as the IPDA, have been further improved by combining them with noise estimation methods such as clutter maps [17] or semi-parametric clutter estimation [7]. A clutter map estimates unique clutter values for regions of the measurement area based on measurement data accumulated over some preceding duration of the tracking scenario. Estimating the clutter densities can assist track initiation by relying more on measurements from low clutter areas than from high clutter areas.

1.3 Objective and scope

In this thesis we will examine how track initiation can be improved in near-shore and other non stationary environments by the use of clutter maps. A regular IPDA will be used for track initiation combined with estimated noise (clutter) density values from the clutter map. Three different clutter map estimators [17] are implemented and tested for a set of scenarios and various parameter tuning. We want to determine strengths and weaknesses of the three clutter map estimators, and compare them to a non parametric IPDA. It is also desirable to identify types of scenarios in which a clutter map is most beneficial.

1.3.1 Contributions

The contributions of this thesis are:

- A review of single target tracking, with extra focus the IPDA method and clutter maps

- A study of how the theoretical biases of different clutter map estimators relate to specific parameter setups and clutter distributions
- A proposed technique for excluding target measurements from clutter estimation
- A comparison of a non parametric IPDA and a parametric IPDA with the three featured clutter map estimators in [17] in the form of:
  - A simulation study of a scenario where two targets are moving into, out of and alongside the transition between high and low clutter areas
  - Performance results on two real data sets recorded in Trondheimsfjorden. The first is a relatively simple COLAV experiment with three vessels present. The second is a longer scenario where the ownship follows another target to shore with several other targets are present.

1.4 Outline

The remainder of the thesis is organized as follows: The second chapter presents target tracking theory. This entails exploring methods of noise filtering, measurement association, track initiation and clutter maps. Chapter three explains how a tracking simulator was created and tuned in Python, and how a test scenario was designed to analyze theory from Chapter two. In Chapter four the biases of different clutter map techniques are evaluated and results from the simulations are presented to compare performance of an IPDA with clutter maps and a non parametric IPDA. Chapter five presents results from two scenarios of real data recorded in Trondheimsfjorden. The results from both the simulations and the real data sets are discussed in Chapter six, and an overall conclusion of the thesis is written in Chapter seven.

This thesis is a continuation of the specialization project [1] written during the fall of 2017.

Parts of this thesis are included in a conference article [19] accepted in FUSION 2018 (21st International Conference on Information Fusion) held at the University of Cambridge. The contribution to the article by the author of this thesis was the implementation and testing of an IPDA. The article is added in Appendix B.

Chapter 2

Tracking theory

Target tracking is the process in which we want to estimate the system states of any potential targets in an area of interest. The estimation is based on measurements from the surveillance area. A target can be a ship, a plane, or any physical object within the search area. In this report we will define the area of interest as a limited search area at sea and the targets will be objects detected at sea level. Hence, these objects will in most cases be ships or small boats. The states we want to estimate are position and velocity for each target.

A tracking system needs a set of measurements from the surveillance area. With high quality measurements, the tracking system will be able to make better estimates of target states. The process of establishing a set of measurements can be divided into several steps. First, we need a way to register objects within the area by receiving reflected energy using appropriate sensors. Then raw data from the sensors needs to be processed to make valid measurements which can be sent as input to a tracking system. Suggested steps of this process using a radar are shown in Figure 2.1 and are inspired by [20].

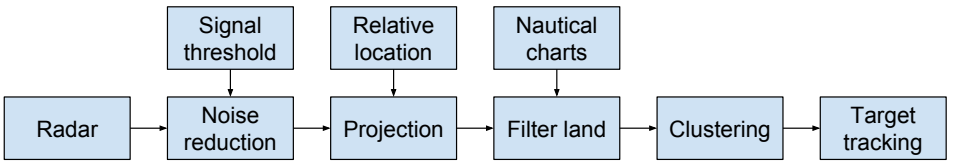

Figure 2.1: Block diagram of proposed processing of radar signals.

The steps shown in Figure 2.1 are explained in Section 2.1.

2.1 Object detection and management

One or several sensors are required to locate targets in the area. Radar, lidar and camera are commonly used sensors for object detection. In this report we will examine target tracking using radar scans. Radar has the advantage that it has a long range and it is not dependent on good lighting conditions like a camera. A radar scan also contains less information so it is easier to interpret. Less information containing only positional data of obstacles within a certain range may be more relevant to a COLAV system than high resolution pictures of the surveillance area. Still, a combination of several sensors could supplement each other.

2.1.1 Radar scanning

A circular radar scan contains information about objects in an area surrounding the transmitter [11]. A transmitter emits radio waves in predetermined directions, whilst a receiver records reflected waves from solid objects hit by the transmitted waves. The receiver is usually in the same position as the transmitter. The angle of the reflected signal and the time difference from a wave was transmitted until a reflected wave returns, is used to estimate the position of the object. A detected object on a scan is called a measurement. An illustration of radar scanning is shown in Figure 2.2 where the blue circles are transmitted waves, and black dots are physical objects in the surrounding area.

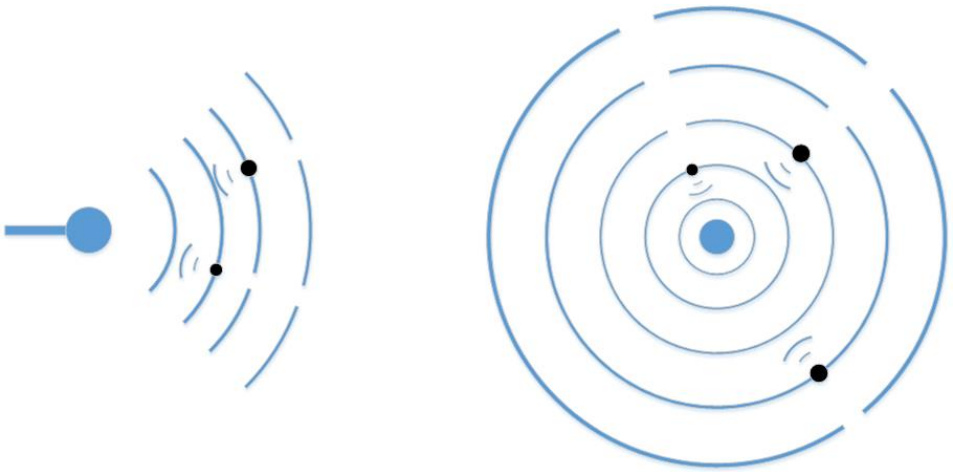

Figure 2.2: Illustration of radar transmitter and a circular scan.

2.1.2 Noise reduction

In a radar scan there is often noise which is referred to as clutter. The clutter derives from unwanted reflected signals in the scanning area. Clutter can be caused by waves, land, rain, snow, animals, insects or other disturbances. Much clutter can be filtered by having a threshold for the received signals, so that only signals with a certain strength results in valid measurements. However, setting a threshold

too high may lead to loss of wanted information. In a real scenario some amount of clutter is to be expected among the validated measurements. Clutter is therefore an important factor to take into consideration in a tracking system.

2.1.3 Projection

The valid measurements has a position relative to the radar location, but should be projected to a world fixed frame. Given the GPS position and attitude of the radar, the measurements can be projected to GPS coordinates. For an autonomous vehicle the radar system may be on board the vehicle, which makes it important to have an accurate navigation system for the vehicle in order to make proper estimates of other target positions.

2.1.4 Filter land

When there is land within the scanned area, the radar will receive a large amount of reflected signals from it. Since most land is already mapped on nautical charts, they should not be tracked by the tracking system. Returns from land should therefore be excluded from the measurements sent to the tracking system. By combining the radar search with a nautical map, we can filter these returns. If a measurement is on a location where it is land on the chart, it is removed from the set of valid measurements sent to the tracking system.

2.1.5 Clustering

Most target tracking methods are based on the assumption that at most one measurement per scan originates from the target. In reality, radar often detects more than one measurement for a single target due to high resolution compared to the target size. Hence, the measurements should be further processed to output only one measurement per target. Measurements within a close range of each other are therefore clustered together to produce one shared measurement. The resulting measurement is located at the centroid of the shared measurements as shown in Figure 2.3.

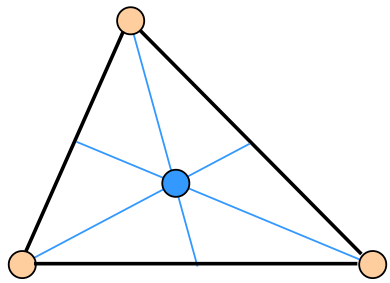

Figure 2.3: Clustering of three measurements.

2.1.6 Target management

Once we have a set of valid measurements in the search area, we need to evaluate which measurements originate from targets and filter the clutter. Doing analysis from a single scan can be ambiguous, so multiple scans are needed to properly analyze the situation. Thus, the measurements are evaluated over several consecutive scans to determine initial states for potential targets. We call this the initiation process. Once one or several targets are initiated, they can be tracked by single-target tracking or multi-target tracking methods. By using single-target tracking we look at each target individually to estimate its state over time. Multi-target tracking considers all measurements and targets simultaneously to establish all target estimates. Single-target tracking will be the focus in this report.

2.2 Single-target tracking

There are several challenges to consider when it comes to single target tracking. First one needs to initiate an estimate of the location and system dynamics of the target. Different approaches exist for finding an initial estimate, however they are often complicated if there is much clutter in the search input data. Once an initial estimate is given, the target has to be continuously tracked over time to keep up with the current state of the system. When only one valid measurement for the target is given as input, the Kalman filter can be used to filter measurement noise and estimate the current system state. However, in a real tracking scenario some amount of clutter is to be expected, and we need some method to associate correct data to the track.

Finding a subset of measurements that most likely originate from a specific target is often refereed to as an association problem. There exists several approaches to this association problem, including PDA [10], nearest neighbour [6] and track split. Nearest neighbour is the simplest of the three, as it takes the closest measurement to the current target estimate and uses it to calculate the next state. Track split is basically a single-target edition of MHT [18], which takes each target estimate and creates a tree of new possible estimates based on nearby measurements in each step. It is more common to use MHT for multi-target tracking, than to use it for single-target tracking in the form of track split. JPDA [2] is a popular method for data association in multi-target tracking which also handles situations where nearby targets share a subset of potential measurements. Since the PDA was introduced it has been one of the most popular algorithms for single-target tracking, and is thus the method we will use in this thesis.

A Probabilistic Data Association (PDA) filter attempts to find the most likely position of the target, by looking at measurements from a certain region of the search input. It expands the Kalman filter method to possibly include more than one measurement. Thus, it can can maintain single-target tracking in a scenario with several valid measurements and measurement noise.

In this section we will discuss Kalman filter and PDA filter, under the assump-

tion that a single track already has been initialized.

2.2.1 Kalman filter

The Kalman filter is used for estimating the state of a given system (target) based on incoming measurements. The filter is used to remove noise and estimate a system state based on a sequence of single measurements. It is an algorithm which first predicts the next state of the system based on previous estimation and measurements, and then updates the estimate when the next measurement is given. The theory of a discrete Kalman filter is explained in further detail in [4], but a summary of the basics of the algorithm are presented in this section.

It is assumed that the system state $\boldsymbol{x}_k$ and measurements $\boldsymbol{z}_k$ can be modeled as follows:

$$\boldsymbol{x}_{k+1} = \boldsymbol{F}_k \boldsymbol{x}_k + \boldsymbol{w}_k \tag{2.1}$$

$$\boldsymbol{z}_k = \boldsymbol{H}_k \boldsymbol{x}_k + \boldsymbol{v}_k \tag{2.2}$$

$\boldsymbol{w}_k$ and $\boldsymbol{v}_k$ are sequences of white noise with known covariances $\boldsymbol{Q}_k$ and $\boldsymbol{R}_k$. An initial estimate of the system is acquired, as the filter bases its predictions for the next state on the current state. The estimated state is denoted $\hat{\boldsymbol{x}}$. The estimation error is

$$\boldsymbol{e}_{k|k-1} = \boldsymbol{x}_k - \hat{\boldsymbol{x}}_{k|k-1} \tag{2.3}$$

and the associated error covariance is

$$\boldsymbol{P}_{k|k-1} = E[\boldsymbol{e}_{k|k-1} \boldsymbol{e}_{k|k-1}^\top] \tag{2.4}$$

Next, the measurement $\boldsymbol{z}_k$ is used to improve the estimate of the current state. The updated estimate is found using a linear blending of the noisy measurement and the prior estimate.

$$\hat{\boldsymbol{x}}_{k|k} = \hat{\boldsymbol{x}}_{k|k-1} + \boldsymbol{K}_k \big(\boldsymbol{z}_k - \boldsymbol{H}_k \hat{\boldsymbol{x}}_{k|k-1}\big) \tag{2.5}$$

The current error covariance is then updated

$$\boldsymbol{P}_{k|k} = \big(\boldsymbol{I} - \boldsymbol{K}_k \boldsymbol{H}_k) \boldsymbol{P}_{k|k-1} (\boldsymbol{I} - \boldsymbol{K}_k \boldsymbol{H}_k\big)^\top + \boldsymbol{K}_k \boldsymbol{R}_k \boldsymbol{K}_k^\top \tag{2.6}$$

To minimize the estimation error, we use a Kalman gain which is given as follows

$$\boldsymbol{K}_k = \boldsymbol{P}_{k|k-1} \boldsymbol{H}_k^\top \big(\boldsymbol{H}_k \boldsymbol{P}_{k|k-1} \boldsymbol{H}_k^\top + \boldsymbol{R}_k\big)^{-1} \tag{2.7}$$

Using (2.7) as $\boldsymbol{K}_k$ we can simplify (2.6) to

$$\boldsymbol{P}_{k|k} = \big(\boldsymbol{I} - \boldsymbol{K}_k \boldsymbol{H}_k\big) \boldsymbol{P}_{k|k-1} \tag{2.8}$$

Based on (2.1) we can set up a prior estimate of the next state. $\boldsymbol{w}_k$ has a zero mean and is not correlated with any of the previous $\boldsymbol{w}$'s. Therefore, the prior estimate of the next state is

$$\hat{\boldsymbol{x}}_{k+1|k} = \boldsymbol{F}_k \hat{\boldsymbol{x}}_{k|k} \tag{2.9}$$

The error for the prior estimate of the next step is

$$\boldsymbol{e}_{k+1|k} = \boldsymbol{x}_{k+1} - \hat{\boldsymbol{x}}_{k+1|k} \quad = \quad \boldsymbol{F}_k \boldsymbol{e}_k + \boldsymbol{w}_k \tag{2.10}$$

The expression for the error covariance of the next step is

$$\boldsymbol{P}_{k+1|k} = \boldsymbol{F}_k \boldsymbol{P}_k \boldsymbol{F}_k^\top + \boldsymbol{Q}_k \tag{2.11}$$

Now we have all the equations we need for setting up the Kalman filter. The filter is a loop consisting of estimating the next step, and then updating according to system measurements. An illustration of the loop is shown in Figure 2.4 inspired by [4].

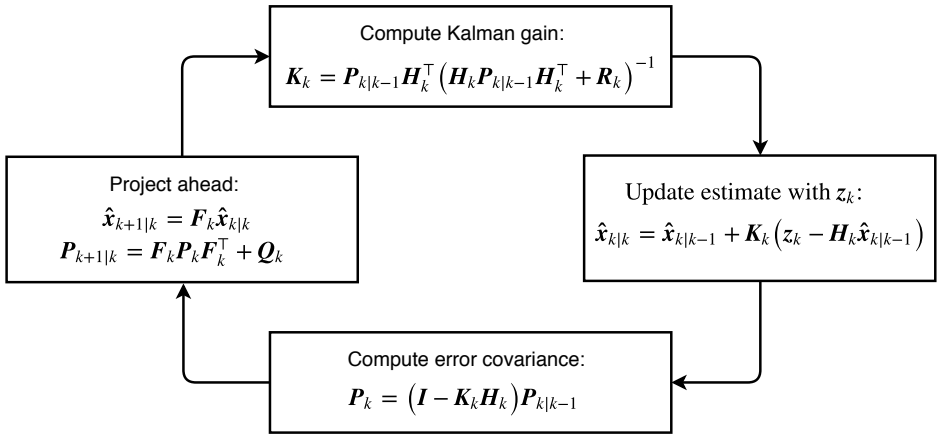

Figure 2.4: Kalman filter loop.

2.2.2 The Probabilistic Data Association filter (PDA)

A PDA filter is used when there is uncertainty in the measurement origin. This uncertainty is often the case for radar tracking of targets, where several measurements are close to the expected target position even though there is only one target in the area. At each time step, a validation region is set up around the prior estimate of the track location, and all measurements within this region are considered in the estimation of the posterior target state. The PDA calculates the association probabilities for each measurement in the validation region at the current time. The measurement-conditional estimates are then used to give a single posterior estimate of the target state. Next, the filter uses a normal Kalman filter approach for giving a prior estimate of the track in the next time step.

Assumptions for PDA [2]:

- There is only one target of interest in the search area.
- The track has been initialized.

- The past information is summarized approximately by

$$p[\boldsymbol{x}_k|\boldsymbol{Z}^{k-1}] = \mathcal{N}[\boldsymbol{x}_k; \hat{\boldsymbol{x}}_{k|k-1}, \boldsymbol{P}_{k|k-1}] \tag{2.12}$$

- At each time a validation region is set up around the prior estimate of the track.
- At most one of the validated measurements can be target-originated.
- There is a fixed probability of clutter in the gate.
- False measurements are independently and uniformly spatially distributed and independent across time.
- The target detections occur independently over time with known probability P_D.

where $\mathcal{N}$ denotes the normal distribution.

Validation gate

The validation gate of the PDA is determined based on the predicted covariance. Higher covariance means a bigger validation region. Only the measurements within the validation region are considered to possibly originate from the target in order to reduce needed calculation. The error covariance is calculated using (2.11) from the Kalman filter approach. Using the error covariance we find the covariance of the innovation corresponding to the true measurement.

$$\boldsymbol{S}_k = \boldsymbol{H}\boldsymbol{P}_{k|k-1}\boldsymbol{H}^\top + \boldsymbol{R} \tag{2.13}$$

This is used to find Normalized Innovation Squared (NIS). The innovation $\boldsymbol{\nu}_k$ is the distance from a measurement to the prior estimate of the track position $\hat{\boldsymbol{z}}_k$. Given the innovation covariance $\boldsymbol{S}_k$ at the current step, we find NIS as

$$NIS = (\boldsymbol{z}_k^i - \hat{\boldsymbol{z}}_k)^\top \boldsymbol{S}_k^{-1}(\boldsymbol{z}_k^i - \hat{\boldsymbol{z}}_k) \tag{2.14}$$

$$= \boldsymbol{\nu}_k^{i\top} \boldsymbol{S}_k^{-1} \boldsymbol{\nu}_k^i \leq \gamma_G \tag{2.15}$$

If NIS is beneath a given threshold γ_G, the measurement is within the gate. γ_G is calculated from a given gate probability P_G. The gate probability P_G is the probability that the target measurement will fall within the gate. The gating threshold γ_G is found from the inverse cumulative distribution function (CDF) of the χ^2-distribution with degrees of freedom corresponding to the dimension of the measurement [20]. An illustration of the validation gate is shown in Figure 2.5. The number of measurements within the gate is denoted m_k.

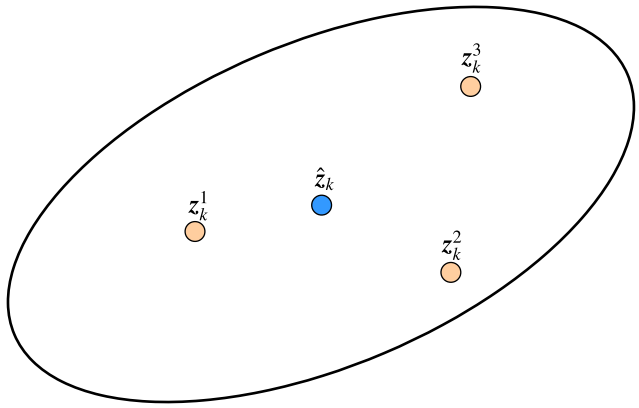

Figure 2.5: Illustration of the validation gate.

Hedging of measurements

When a (possibly empty) set of measurements is assigned to each target, an approximation of the posterior state of the target is made by hedging on the measurements. Each measurement is assigned a value according to the likelihood that it originated from the target. The probability that none of the measurements are valid is also calculated. The probabilities are calculated by the given formula from [2]

$$\beta_k^i = \begin{cases} \frac{1}{c}\exp(-\frac{1}{2}\boldsymbol{\nu}_k^{i\top}\boldsymbol{S}_k^{-1}\boldsymbol{\nu}_k^i), & i = 1 \ ... \ m_k \\ \frac{1}{c}\frac{2(1-P_D P_G)}{\gamma_G P_D} m_k, & i = 0 \end{cases} \tag{2.16}$$

where c is a normalization constant.

The combined innovation is

$$\boldsymbol{\nu}_k = \sum_{i=1}^{m_k} \beta_k^i \boldsymbol{\nu}_k^i \tag{2.17}$$

Once the probabilities of each validated measurement is calculated, the posterior estimates for predicted state and error covariance are calculated

$$\boldsymbol{K}_k = \boldsymbol{P}_{k|k-1}\boldsymbol{H}^\top \boldsymbol{S}_k^{-1} \tag{2.18}$$

$$\hat{\boldsymbol{x}}_{k|k} = \hat{\boldsymbol{x}}_{k|k-1} + \boldsymbol{K}_k\boldsymbol{\nu}_k \tag{2.19}$$

$$\boldsymbol{P}_{k|k} = \boldsymbol{P}_{k|k-1} - (1-\beta_k^0)\boldsymbol{K}_k\boldsymbol{H}\boldsymbol{P}_{k|k-1} + \tilde{\boldsymbol{P}}_k \tag{2.20}$$

where

$$\tilde{\boldsymbol{P}}_k = \boldsymbol{K}_k\left(\sum_{i=1}^{m_k}\beta_k^i\boldsymbol{\nu}_k^i\boldsymbol{\nu}_k^{i\top} - \boldsymbol{\nu}_k\boldsymbol{\nu}_k^\top\right)\boldsymbol{K}_k^\top \tag{2.21}$$

$\tilde{\boldsymbol{P}}_k$ is also called the Spread of Innovations (SOI).

Now we have all we need to track a target through the same recursive principles as with the Kalman filter.

2.2.3 Parametric PDA

Given an estimation of clutter density in specific areas of the map, we can adjust the PDA algorithm to include the additional information. When the tracking method assumes knowledge of the prior clutter density, we call it a parametric method. This terminology is in agreement with Musicki's definition of parametric methods [17], though there exists other definitions of the term as well.

When using PDA combined with a clutter map or another method to more accurately estimate clutter density, we say we are using a parametric PDA [16]. In the case where we have non-uniform clutter estimation in the map, for example by using a clutter map, the calculation of (2.16) in the PDA algorithm needs to be modified. This is done to account for specific clutter density linked to each measurement. We can calculate the probabilities β_k^i as in [16]

$$\beta_k^i = \begin{cases} \frac{1}{c} P_D p(\boldsymbol{z}_k^i | \theta_k^i)/\lambda_i, & i = 1 \ ... \ m_k \\ \frac{1}{c}(1 - P_D), & i = 0 \end{cases} \tag{2.22}$$

A parametric method should in theory perform better than a non parametric, but is naturally also harder to implement due to more added complexity. The extra complexity and adjustments needed can take time to implement and may slow down runtime, but should lead to more accurate results. A non parametric PDA will use all measurements within the gate at scan k under the assumption that the number of clutter measurements inside the gate is uniformly distributed. For each measurement, only the distance from the predicted position and covariance of the estimate $\hat{\boldsymbol{x}}_{k|k-1}$ will contribute to the new estimated position. On the other hand, a parametric PDA with a clutter map is able to distribute unique clutter densities to each measurement dependent on their geographical position. This will support the tracking system in being more critical to measurements from high clutter areas. An illustration of how a clutter map may provide more accurate results is shown in Figure 2.6

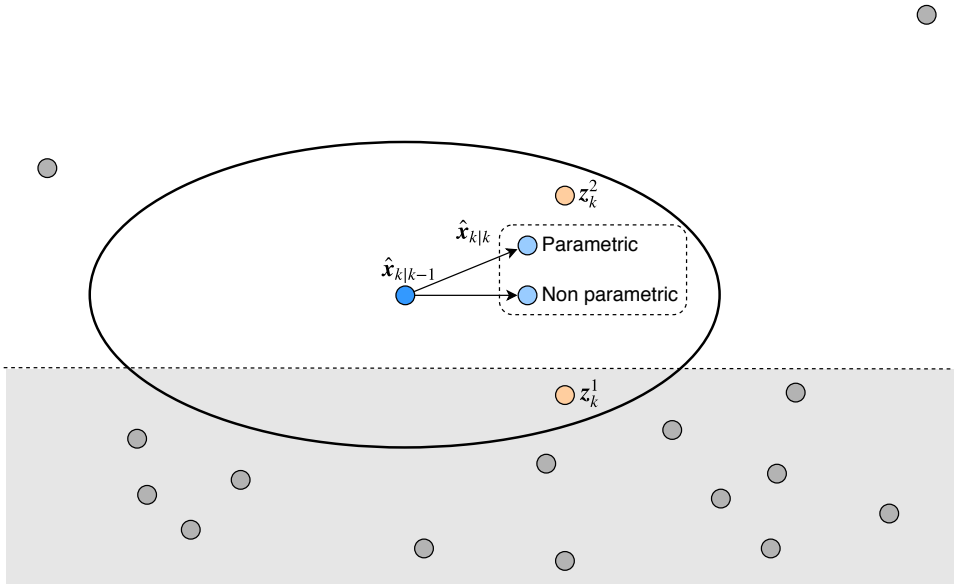

Figure 2.6: PDA estimated position near high clutter area. A parametric method will reduce the association probability of z_k^1, since it originates from an area with higher clutter density.

2.3 Track initiation

In most real world scenarios, the number of targets as well as their initial positions in the scanning area are not given. Hence, an initiation method is required to locate potential tracks. Once a set of potential tracks has been established, each target can be tracked using a single- or multi-target tracking method as discussed in Section 2.1.6. Depending on the chosen initiation method and the associated parameters as well as the tracking scenario, some amount of false tracks may be initiated. These false tracks need a way to be terminated. A proper termination method can continuously determine the validity of ongoing tracks. Suitable methods for track initiation and termination are vital to maintain a viable tracking process over time. Several methods attempt to initiate and terminate tracks through single-target-tracking or multi-target tracking, such as M/N [2], IPDA [15], IMM-PDA [2], SPRT [9] and Perceivability [12]. Before we elaborate how the different methods work, a normal way of categorizing tracks during their initiation phase is presented.

2.3.1 Pre-track categorization

All measurements of a radar scan can potentially come from valid targets in the searching area. Data from one scan must in some way be kept and compared with several succeeding scans in order to obtain actual targets and filter the clutter. Measurements within a small range of each other in successive scans increases the possibility that a target exists in that area. As in [2], measurements in each scan

can be linked with one of three pre-track categories; tentative tracks, preliminary tracks and confirmed tracks.

Tentative tracks are all measurements that are not associated with any existing preliminary or confirmed tracks. They are candidates for new tracks. In the succeeding scan, a validation gate is set up around each tentative track. If a measurement falls within the validation gate of a tentative track, a preliminary track is formed. Otherwise, the tentative track is dropped. The radius of the validation gate is given by the maximum velocity of a potential target and the measurement noise statistics.

Since a preliminary track has two measurements it can be tracked using for instance a PDA. Further conditions or logic can now be examined in order to confirm or terminate the preliminary track. This is the critical part of the track initiation problem. If the confirmation conditions are met, the track is confirmed and thus assumed to be a real target. Still there may be reasons for terminating some of the confirmed tracks in the future, so a termination routine should be operating concurrently with the tracking.

M/N-logic

The M/N approach [2] is a simple and intuitive way of initiating and/or terminating tracks. It checks for a number of M associated measurements in a total of N scans to evaluate initiation or termination. A preliminary track is tracked using a filter with a validation region in each step, like the PDA, for a maximum of N scans. If a measurement falls within the validation region of the preliminary track at least M times over the N scans, the track is confirmed. If not, the track is terminated. The same procedure can be used for terminating a confirmed track, and the chosen parameters for M and N need not be the same as for initiation. M/N is often used in tracking systems because of its easy implementation.

A comparison of M/N and IPDA was done in the specialization project [1]. The results from the report demonstrated that IPDA achieves significantly better results in high clutter areas, and the difference is noticeable for clutter densities higher than $\lambda = 1 \times 10^{-5} m^{-2}$. This corresponds to a rough estimate of 0.04 gated clutter measurements in each scan for a confirmed track.

IMM-PDA

An Interacting Multiple Model PDA (IMM-PDA) filter [2] can be used for initiation and termination of tracks, as well as track maintenance. It continuously estimates a True Track Probability (TTP) for each track, which is used for track initiation and track termination given predetermined thresholds. The estimation of TTP for track initiation is done by following a model for a detectable target and a model for an undetectable target. The TTP is updated in each step according to a Markov chain transition matrix. The initiation process is quite similar to the IPDA, but they differ in techniques of track maintenance and track termination. An IMM-PDA may have additional models after track initiation where other models of target

movement and detectability are modeled.

SPRT

Sequential Probability Ratio Test (SPRT) [9] is a method which calculates the ratio between the probabilities of a track existing in the scan, and no tracks existing in the scan. It is also referred to as sequential track extraction. The initiation method distinguishes two hypotheses from all measurements in a scan;

- All measurements in a scan consists of detections from one target and false detections
- All measurements in a scan are false detections

Under the assumption that there exists at most one target, and no special cases of suddenly spawned targets or a suddenly disappearing target, one of the two hypotheses has to be true. Both of these probabilities are calculated in each step, and the ratio between them is used to decide which hypothesis is accepted.

Perceivability

Perceivability [12] attempts to more accurately estimate the number of false measurements m_k within a track gate, and thus also the clutter density within the gate. It makes no attempt to map locations of clutter densities in the map, only the clutter density connected to a track and its gate for each scan.

There are three approaches to estimating m_k;

- Conditional Mean Estimation
- Maximum Likelihood Estimation
- Method of Moments

The Conditional Mean Estimation models three ways of estimating m_k based on the current scan and previous scans. Maximum Likelihood Estimation assumes the clutter density in each scan to be unknown, but nonrandom. Method of Moments is a heuristic approach to estimating the number of false measurements m_k.

2.3.2 Integrated Probabilistic Data Association (IPDA)

The IPDA includes calculation of probability of track existence in a PDA controller in order to initiate or terminate tracks. The model is described in depth in [15]. The track existence is modeled as a Markov process with two possible cases. The first, which is called Markov chain one, models two possibilities;

- The track exists
- The track does not exist

The second case, which is called Markov chain two models three possibilities;

- The track exists and is observable
- The track exists and is not observable
- The track does not exist

Markov chain one is usually used in track initiation, while both can be used for the termination process of confirmed tracks. Markov chain two can be used for track maintenance when the target trajectory is partly obscured or the target is otherwise not observable for parts of the tracking [15]. It has also proven better results for track maintenance when tested on real data in [3]. Since this report is focused on track initiation, Markov chain one is used in the tracking simulator.

Algorithm for IPDA using Markov chain one

Let $\boldsymbol{Z}^k$ be the set of measurements falling in the track window up to and including scan k. The a priori probability that the track exists at scan k is denoted by $P\{\boldsymbol{x}_k|\boldsymbol{Z}^{k-1}\}$ and the probability that the track does not exist is denoted by $1 - P\{\boldsymbol{x}_k|\boldsymbol{Z}^{k-1}\}$

$$P\{\boldsymbol{x}_k|\boldsymbol{Z}^{k-1}\} = p_{11}P\{\boldsymbol{x}_{k-1}|\boldsymbol{Z}^{k-1}\} + p_{21}\big(1 - P\{\boldsymbol{x}_{k-1}|\boldsymbol{Z}^{k-1}\}\big) \tag{2.23}$$

$$1 - P\{\boldsymbol{x}_k|\boldsymbol{Z}^{k-1}\} = p_{12}P\{\boldsymbol{x}_{k-1}|\boldsymbol{Z}^{k-1}\} + p_{22}\big(1 - P\{\boldsymbol{x}_{k-1}|\boldsymbol{Z}^{k-1}\}\big) \tag{2.24}$$

p_{11} is often called "survival probability" as it denotes the chance that a track at scan $k-1$ is still present at scan k. p_{21} denotes the chance that a track is present at scan k which was not present at scan $k-1$, and is often called "probability of birth". The Markov chain coefficients must satisfy

$$p_{11} + p_{12} = p_{21} + p_{22} \tag{2.25}$$

The standard procedures from PDA are used to find predicted state $\hat{\boldsymbol{x}}_{k|k-1}$, predicted covariance $\hat{\boldsymbol{P}}_{k|k-1}$ and validation gate. Using (2.3.2-1) in [2] we can calculate a priori probability density of the predicted measurement position $f(\boldsymbol{z}|\boldsymbol{Z}^{k-1})$

$$f(\boldsymbol{z}|\boldsymbol{Z}^{k-1}) = \mathcal{N}[\boldsymbol{z}_k; \hat{\boldsymbol{z}}_{k|k-1}, \boldsymbol{S}_k] \tag{2.26}$$

Let $\hat{m}_k$ denote the number of expected false tracks in the set of measurements at scan k. If the number of false measurements in each scan are modeled by a Poisson distribution over the validation region V_k, then

$$\hat{m}_k = \lambda V_k \tag{2.27}$$

where λ is a known parameter. For a non parametric IPDA where λ is unknown, $\hat{m}_k$ can be estimated by (2.28) below.

$$\hat{m}_k = \begin{cases} 0, & m_k = 0 \\ m_k - P_D P_G P\{\boldsymbol{x}_k|\boldsymbol{Z}^{k-1}\}, & m_k > 0 \end{cases} \tag{2.28}$$

V_k can be found by (2.3.2-4) in [2].

The posterior track existence using the following equation.

$$P\{\boldsymbol{x}_k|\boldsymbol{Z}^k\} = \frac{1-\delta_k}{1-\delta_k P\{\boldsymbol{x}_k|\boldsymbol{Z}^{k-1}\}} P\{\boldsymbol{x}_k|\boldsymbol{Z}^{k-1}\} \tag{2.29}$$

This value $P\{\boldsymbol{x}_k|\boldsymbol{Z}^k\}$ may be used to determine whether to initiate or terminate the track. We use δ_k to scale the a priori existence value according to data from the newest scan k. Combining (2.9) and (2.13) in [15] we evaluate the quantity δ_k. With a non parametric IPDA, δ_k is calculated as follows

$$\delta_k = \begin{cases} P_D P_G, & m_k = 0 \\ P_D P_G - P_D \frac{V_k}{\hat{m}_k} \sum_{i=1}^{m_k} f(\boldsymbol{z}_k^i|\boldsymbol{Z}^{k-1}), & m_k > 0 \end{cases} \tag{2.30}$$

$f(\boldsymbol{z}_k^i|\boldsymbol{Z}^{k-1})$ in (2.30) is found using (2.26). Since a parametric IPDA with a clutter map has specific estimates for λ, these estimated densities can be used directly to calculate δ_k. The estimated number of false alarms $\hat{m}_k$ from (2.27) can then be calculated for each measurement according to its area of origin, and we get

$$\delta_k = \begin{cases} P_D P_G, & m_k = 0 \\ P_D P_G - P_D \sum_{i=1}^{m_k} f(\boldsymbol{z}_k^i|\boldsymbol{Z}^{k-1})/\lambda_i, & m_k > 0 \end{cases} \tag{2.31}$$

This altered δ_k readjusts the contribution of each measurement depending on the associated clutter density. Measurements from low clutter areas will have a bigger impact on the existence probability than other high clutter measurements on the same scan. For instance, measurement z_k^2 will contribute more positively to the existence probability than measurement z_k^1 in Figure 2.6.

2.4 Clutter map

A clutter map is useful when there are variations in the clutter density distribution across the scanned area. It can be combined with a multitude of tracking methods such as PDA, IPDA [17], M/N, SPRT [19], to improve their performance. As listed in Section 2.2.2, one of the assumptions for the PDA is; "False measurements are independently and uniformly spatially distributed and independent across time", which is an invalid assumption in many maritime scenarios. Therefore, a register of specific clutter densities connected to smaller regions of the map can make the tracking method applicable in a larger scope of scenarios. Figure 2.7 shows how including a clutter map with the IPDA can assist in existence calculation.

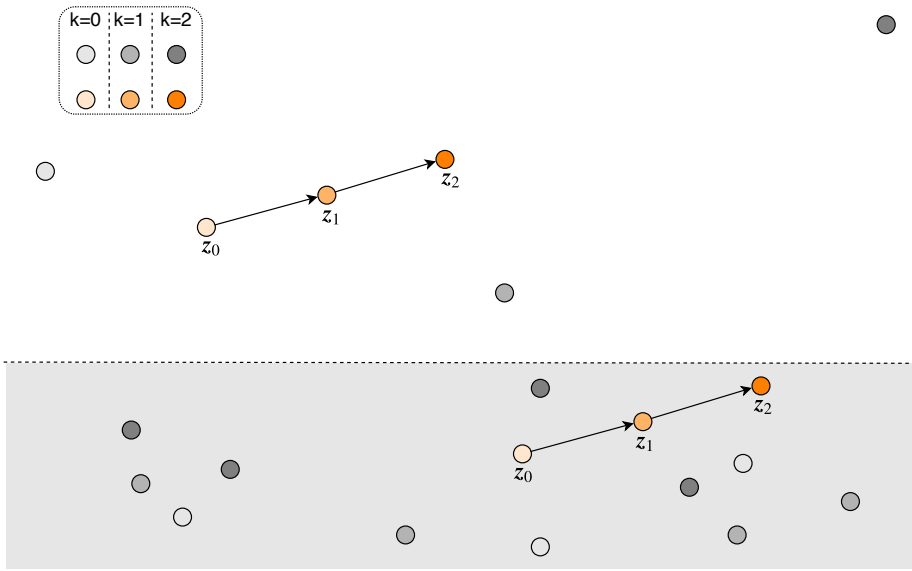

Figure 2.7: Illustration of a target tracking scenario with a clutter map. The upper target will have a higher existence probability for a parametric IPDA.

Without a priori knowledge of the clutter density in an area, a parametric method has to estimate clutter density over a number of scans [17]. This can be done globally for the whole search area or locally in smaller regions in the form of a clutter map. We will examine three different clutter map estimators originally proposed in [17];

- Classic clutter map estimator
- Spatial clutter map estimator
- Temporal clutter map estimator

The three estimators share some fundamental principles and equations. The surveillance volume is divided into clutter map cells c with volume $V(c)$, which may or may not be equal. The chosen cell sizes may depend on clutter density characteristics, sensor resolution and other system requirements. System runtime may for instance be of importance, and can be affected negatively by having too many and too small cells. It is assumed that the clutter statistics are stationary. In the case that they are not stationary, they must not change faster than the clutter map is updated. In other words, the clutter densities must not change with a time constant which is lower than the time constant to update the map.

The clutter map cells are updated using moving average for averaging. It is also possible to use an auto regressive filter [17], but only the moving average will be

used in this thesis. The moving average is calculated as follows

$$N_k(c) = \frac{1}{L} \sum_{j=k-L+1}^{k} \mu_j(c) \tag{2.32}$$

where $\mu_j(c)$ is the clutter density measurement to be averaged in cell c. For both the Classic and the Spatial method, k is the scan index. For Temporal clutter map, k is the index of measurement arrival in clutter cell c. $N_k(c)$ denotes the averaged value and the averaging time constant is denoted by L.

As previously stated, (2.32) is a fundamental equation used by all three estimators. The following subsections addresses how the estimators update a single clutter map cell c.

2.4.1 Classic clutter map estimator

The Classic estimator averages the number of measurements which have fallen into the cell over time. $\mu_k(c)$ in (2.32) is in this case the number of measurements in the cell at scan k. The average found in (2.32) is used to calculate the clutter density in the cell. The clutter density is calculated as follows

$$\hat{\lambda}_k(c) = \frac{N_k(c)}{V(c)} \tag{2.33}$$

$$\widehat{\left(\frac{1}{\lambda_k(c)}\right)} = \frac{V(c)}{N_k(c)} \tag{2.34}$$

To avoid division by zero, the estimator using (2.32) is slightly modified

$$\begin{aligned} N_k^1(c) &= \begin{cases} N_k(c), & N_k(c) > 0 \\ \frac{1}{L}, & N_k(c) = 0 \end{cases} \\ \widehat{\left(\frac{1}{\lambda_k(c)}\right)} &\propto \frac{V(c)}{N_k^1(c)} \end{aligned} \tag{2.35}$$

This modification is necessary when no measurements have appeared in the cell during the averaging window L. The clutter is then estimated as though it was one measurement in the cell during the averaging window.

2.4.2 Spatial clutter map estimator

The Spatial estimator uses the volume of a shape to estimate distance from the center of the cell. The shape can be a circle, a square or another type of polygon. The shape expands in size, while the center is retained, until it touches the nearest measurement position. The size expansion is not limited by the cell area, and may overlap other cells to touch the closest measurement. However, the volume will need to be adjusted if parts of the shape ends out of the surveillance area, by discarding the portion outside of the surveillance area. An illustration of the shape

is shown in Figure 2.8 where the red box is an illustrated cell on the map, the blue box is the shape associated with the cell and the orange circles are measurements.

The volumes found in each cell are stored as $\mu_k(c)$ and are averaged by (2.32) to find $N_k(c)$. The clutter density estimate is

$$\widehat{\left(\frac{1}{\lambda_k(c)}\right)} = N_k(c) \tag{2.36}$$

An effective implementation of the Spatial estimator is somewhat more difficult to conduct than the other two estimators, since the cell is not limited to measurements within its own borders. Using a nearest-neighbour method through the use of a k-d tree [5] can help minimize runtime when implementing the estimator.

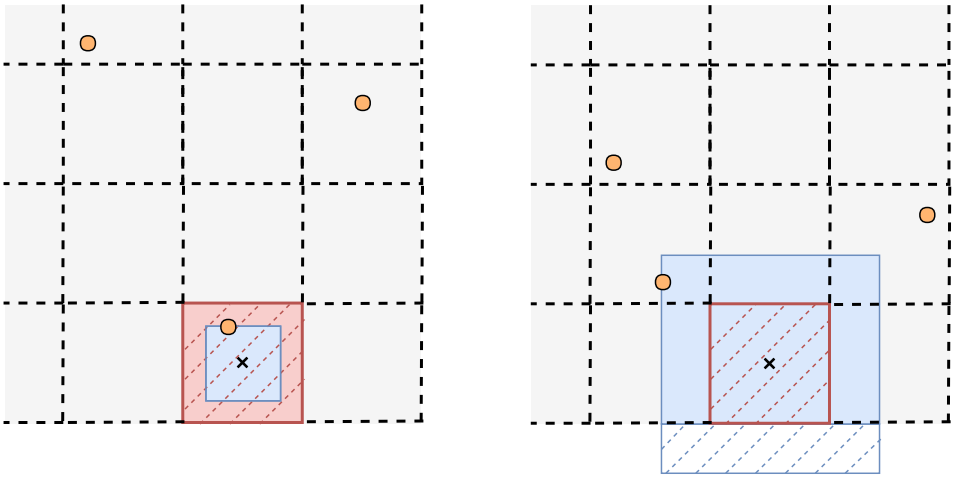

Figure 2.8: Illustration of the shape expansion for the Spatial estimator for two different scans. The blue solid colored areas are the volumes used as $\mu_k(c)$.

2.4.3 Temporal clutter map estimator

The Temporal estimator uses inter-arrival time $\tau(c)$ of two Poisson measurements arriving in the same map cell, to calculate the mathematical distance between the two. For instance, if one measurement arrives three scans later than the previous measurement in the cell, its inter-arrival time $\tau(c)$ is three. If more than one measurement, $n_k(c) > 1$, arrives in the cell at scan k, a number of $n_k(c)$ equal values $\tau_k(c)/n_k(c)$ is used as input to (2.32). If one were to check the density of a cell c prior to any measurements falling in the cell, an inter-arrival time of the current scan index k is used to estimate the clutter density. The average time between measurements arriving in the cell combined with the volume of the cell provides an estimate of the clutter density within that volume. The clutter density

of the Temporal estimator is

$$\widehat{\left(\frac{1}{\lambda_k(c)}\right)} = V(c) \cdot \widehat{(\overline{\tau}_k(c))} \tag{2.37}$$

2.4.4 Biases of the estimators

The three clutter map estimators provide good approximations to the real clutter density, yet they have some biases and other disadvantages connected to their estimation techniques. Their difference in methodology makes them more or less applicable for distinct scenarios and setups. Knowing their advantages and disadvantages can be valuable when choosing which type of estimator to use. In this section we will refer to the relative bias of the estimators, which is given as

$$\text{relative bias} = \frac{\widehat{\frac{1}{\lambda(c)}} - \frac{1}{\lambda(c)}}{\frac{1}{\lambda(c)}} \tag{2.38}$$

The clutter estimation of the Classic and the Temporal estimator is constrained by the measurements falling within the area of one cell c. Thus we will discuss their biases according to the average number of measurements falling in a cell during one averaging length of L scans.

The Classic estimator has a bias due to the fact that there may be zero measurements within the averaging window at any given time. Without measurements, the estimator has no base to say anything about the clutter density in the cell c. When there are no measurements, the modified moving average $N_k^1(c)$ in (2.35) sets a limit for the lowest estimated density in the gate, as it estimates the density as though it was one measurement within the window L. Hence, when the real clutter density is lower than the lowest limit of the cell, it gets a negative bias. A negative bias means a higher estimated density than the actual density. The theoretical bias of the Classic estimator is shown in Figure 2.9.

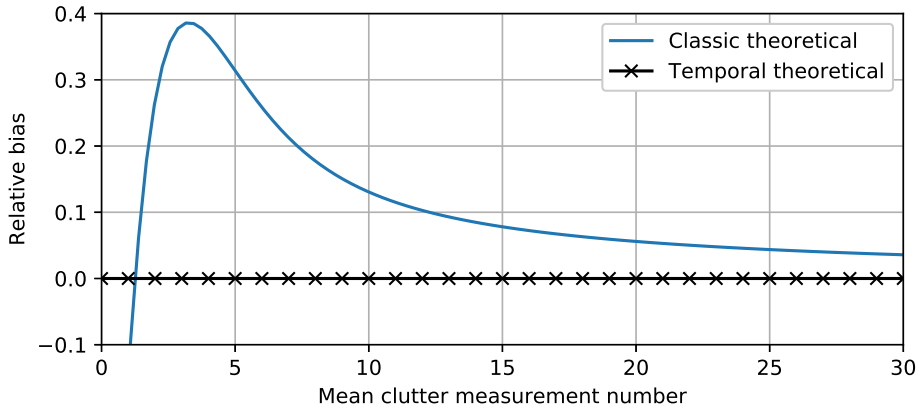

Figure 2.9: Theoretical values for classic and temporal clutter map bias.

The reason why the Classic estimator has a positive bias converging to zero as the mean number gets higher, may be because of the limited data set within the averaging window. When there are fewer measurements than average inside the window, the estimated clutter density is relatively lower than the true value than when there are more measurements than average inside the window. This results in a lower estimated density on average than the true density.

While the Temporal estimator does not have any theoretical bias, as shown in Figure 2.9, it does have some amount of bias during initiation, as well as some bias caused by the length of the averaging window. Its theoretical bias of zero is only achieved as $t \to \infty$ and $L \to \infty$. This will be further explained and tested in Section 4.1.

The relative variance of the Classic and the Temporal estimators converge slower than the Spatial estimator, as illustrated and discussed in [17]. Temporal has a variance diverging towards infinity as the mean clutter measurement number approaches zero. However, for mean numbers of clutter higher than approximately 2.5 measurements, Classic has the most relative variance.

The Spatial estimator does not have any bias of the clutter density in uniform clutter, but it does have a bias when there is heavy clutter close to the cell. Meaning that if a cell in a low clutter area is located near a high clutter area, its estimated density will be higher than otherwise, and vice versa for cells in high clutter areas. When there exists a discontinuity in real clutter density, the estimator will have a transition between regions of different density. The other two estimators does not suffer from this bias. This is shown in Figure 2.10.

Nevertheless, even though the Spatial estimator has a bias for real clutter density discontinuities, it has the fastest convergence of the three estimators. The Spatial estimator averages in every scan even when there are no measurements in a cell, as shown in Figure 2.8. The Temporal estimator on the other hand, has no theoretical bias, but can have especially slow convergence in low clutter. Considering that it only accumulates information in scans where there are measurements falling into the cell, it can take many scans to fill up the averaging window in low clutter areas. Spatial and Classic will have filled up their averaging window when scan index k equals the averaging length, and on average they will not produce more accurate estimates beyond that point.

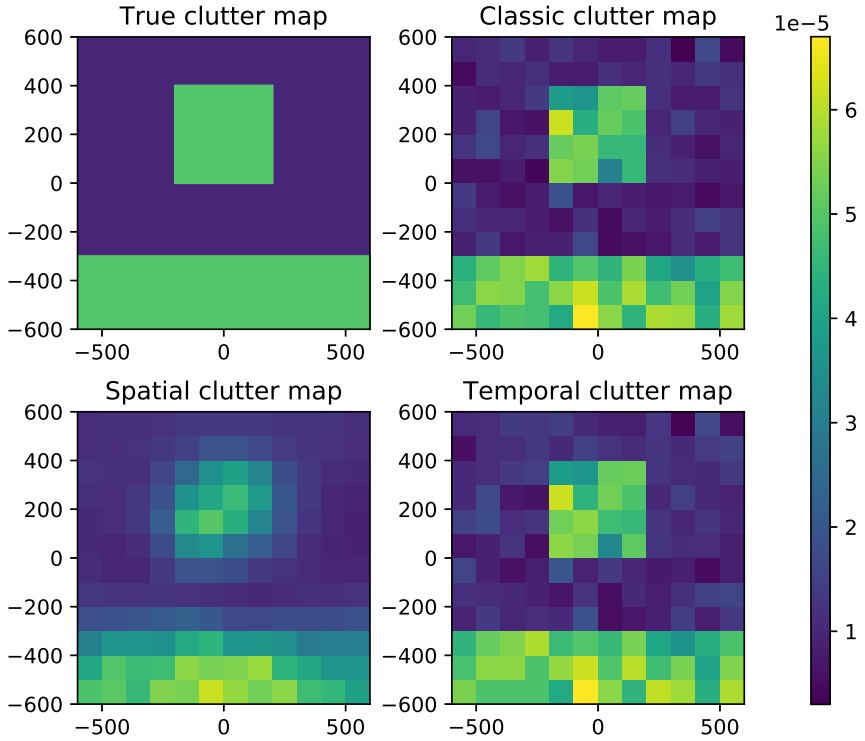

Figure 2.10: All estimators after 100 scans and an averaging length of 100.

2.4.5 Clutter map dilemmas

There are a few minor dilemmas with the clutter map estimators in [17] which are examined in this subsection. These are:

- How to exclude target measurements from clutter density calculation?
- Should a clutter map method have time to initialize before any tracking begins?

How to exclude target measurements from clutter density calculation?

The first question applies to all three methods and has a number of possible solutions. The easiest solution is to ignore the issue entirely and use all measurements in the calculation of the clutter map. This can however lead to a negative bias when calculating existence probability of a preliminary track. If a target measurement is included in clutter calculation, the measurement itself will appear less reliable and thus lessen the existence probability of the target.

An idea is to exclude the most likely target originated measurements from the estimation of the clutter map. Selecting which measurement to exclude can be achieved in several ways. The simplest way would be to pick the measurement closest to the estimated position of the preliminary track. Another way is to pick the measurement with the highest association probability from (2.16), though this may be slightly more complicated to implement efficiently. In the end there is no simple way of choosing the target measurement with certainty. The issue is not addressed in Musicki's article [17].

If one were to exclude the measurements most likely originate from targets from the clutter calculation, it would be most essential for confirmed tracks. Though it can also be applied to preliminary tracks. The reason for applying this to preliminary tracks is to hopefully improve the track initiation phase. However, excluding certain measurements associated with preliminary tracks from the clutter estimation may contradict its purpose and lead to a less accurate clutter map. Preliminary tracks which are subsequently terminated would erroneously lower the estimated clutter density in their area.

Should a clutter map method have time to initialize before any tracking begins?

The estimated values of a clutter map cell are improved for each incoming scan until L scans after startup. Consequently, given an initialization time before the tracking starts, the estimated clutter densities may be more precise and lead to more accurately estimated existence probabilities in the IPDA algorithm. A possible complication with having an initialization phase is that there might be ships or other objects in the search area which adds additional false clutter to the map, as with the first listed dilemma in this section. Combining an initialization phase with tracking, to exclude target measurements could result in better initiated clutter maps. The tracking process could then be reset once the clutter map is initiated. A proper way of excluding target measurements would thus solve more than one problem.

Chapter 3

Simulator Development

All simulations in this thesis were run in Python 2.7. The simulator and test environment are based on the work done in [1] and done in a collaboration with the Autosea project.

3.1 Motion model

The system is modeled with the same assumptions for system process and noise as in the Kalman filter in section 2.2.1. The state is given as $\boldsymbol{x} = [N, V_N, E, V_E]$ with N and E being North and East position in the North East Down (NED) reference frame. V_N and V_E are velocities in north and east direction.

$$\boldsymbol{F} = \begin{bmatrix} 1 & T & 0 & 0 \\ 0 & 1 & 0 & 0 \\ 0 & 0 & 1 & T \\ 0 & 0 & 0 & 1 \end{bmatrix}, \quad \boldsymbol{H} = \begin{bmatrix} 1 & 0 & 0 & 0 \\ 0 & 0 & 1 & 0 \end{bmatrix} \tag{3.1}$$

The covariances for the process and measurement noise are set as follows

$$\boldsymbol{Q} = \sigma_q^2 \begin{bmatrix} T^4/4 & T^3/2 & 0 & 0 \\ T^3/2 & T^2 & 0 & 0 \\ 0 & 0 & T^4/4 & T^3/2 \\ 0 & 0 & T^3/2 & T^2 \end{bmatrix}, \quad \boldsymbol{R} = \sigma_r^2 \begin{bmatrix} 1 & 0 \\ 0 & 1 \end{bmatrix} \tag{3.2}$$

$\boldsymbol{Q}$ is based on equation (15) in [13], but expanded for two dimensions.

3.2 Tracking filter and initialization parameters

The parameters for the tracking model were tuned during testing to achieve desired results. The parameters are shown in Table 3.1 and Table 3.2. The values for Markov chain coefficients in Table 3.2 were based on the parameters used in the simulations in [15]

P_G	0.99
γ_G	9.21

Table 3.1: PDA parameters

p_{11}	0.98	p_{12}	0.02
p_{21}	0.00	p_{22}	1.00
Terminate threshold			0.15
Initial probability			0.5

Table 3.2: IPDA parameters

The initiation threshold of the IPDA was varied to test the initiation accuracy of various thresholds.

3.3 Clutter map parameters

There are several parameters and implementation choices one can make that will affect the runtime, accuracy and desired results of a clutter map. The most apparent choices are listed below.

- Map size
- Clutter cell size
- Reference frame
- Initialization time
- Averaging length

Map size will obviously affect runtime, but also clutter cell size will have an impact on runtime, as a unique clutter density will have to be calculated for each cell. Thus there is a trade off between accuracy and runtime depending on the chosen size of each cell. In scenarios where certain predictable areas of the map contain more clutter variation than others, one might consider having smaller cells in these areas. Near-shore areas might bring more clutter returns, and may thus be a sensible choice for applying smaller cells. A problem with not having a fixed clutter cell size throughout the map, is that one will need additional logic for selecting desired size for each area. Running additional logic could potentially slow down the system, meaning a smaller fixed cell size for the whole map may be just as efficient.

In this thesis, the map size has been chosen to be $1200m \times 1200m$. Real radar data recorded by the Autosea project using the ship Telemetron have a significantly wider range, so the simulated map size in this thesis is just scaled down to reduce testing runtime. The clutter cell size was originally thought to be chosen to

something similar to $20m \times 20m$ which was used in [17], however, with the chosen map size this equals 3600 individual cells and took far too long to simulate. Hence, the cell size was chosen to $100m \times 100m$ which gave acceptable simulation times. No attempt was made to implement any logic for selecting unique cell sizes on different areas on the map.

Making the clutter cells earth-fixed or fixed to the ownship's reference frame can be justified for different reasons. If the radar itself is more prone to clutter in certain directions of the scan or specific areas relative to the ownship, a clutter map with a body-fixed reference frame could prove beneficial. If the clutter is more prone in certain areas of the sea, for instance close to shore, the reference frame is ideally an earth-fixed coordinate frame. In this thesis we have chosen to use an earth-fixed reference frame, since our main focus is geographical clutter and not due to radar or other ship equipment.

The number of scans one uses for initialization impacts the accuracy of the clutter map upon the start of the tracking phase, as discussed in Section 2.4.5. Appropriate initialization length depends on a number of factors such as clutter cell size and clutter density. Different initialization lengths are tested based on the test setup in Section 3.4 and discussed in Section 3.5. The averaging length L is in each test case set to the same length as the selected initiation time.

The clutter map was updated prior to the IPDA calculations in each scan, meaning that the clutter densities used in the IPDA were based on N_k and not N_{k-1} as in [17]. This was done to accumulate more data prior to IPDA estimation, but it could also be argued that it alters gate measurement contributions negatively by weighting the measurements twice in the same scan.

All parameters used in conjunction with the clutter map implementation are listed in Table 3.3.

Map size	$1200m \times 1200m$
Cell size	$100m \times 100m$
Number of cells	144
Initialization time	20, 40 or 80
Averaging length	Same as initialization time
Reference frame	Earth-fixed

Table 3.3: Clutter map parameters

To test the real measurement data presented in Chapter 5, the clutter map implementation had to be partially modified. The real data sets are recorded from radar data on board a moving vessel, so the measurement area is not stationary. Clutter cells were therefore only initiated and updated when they were inside the maximum range of the radar.

3.4 Test setup

To test the clutter map methods properly we need a set of scenarios with various target setups and clutter layouts on the map. We want to imitate real tracking challenges in marine environments to find the best possible methods and solutions these difficulties. Since this thesis is focused on track initialization near shore, a test scenario was made to mimic situations where vessels move close to or out of high clutter areas. The final setup can be seen in Figure 3.1, which is the same setup which was used in [19]. The initial state of the two spawned targets follow Table 3.4.

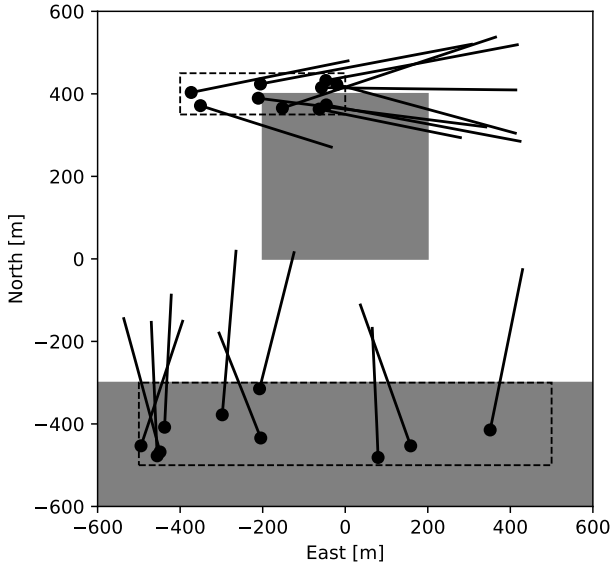

Figure 3.1: Test scenario. The gray areas are high clutter areas. One target is spawned in each dashed region in one test run. The black lines are sample trajectories of spawned targets.

Parameter	Lower target	Upper target
N_0	$\mathcal{U}(-500m, -300m)$	$\mathcal{U}(350m, 450m)$
E_0	$\mathcal{U}(-500m, 500m)$	$\mathcal{U}(-400m, 0m)$
V_0	$\mathcal{U}(5ms^{-1}, 10ms^{-1})$	$\mathcal{U}(5ms^{-1}, 12ms^{-1})$
χ_0	$\mathcal{U}(-30°, 30°)$	$\mathcal{U}(70°, 110°)$

Table 3.4: Target parameters for the test scenario

After an initial state is set, the simulated targets move according to a nearly constant velocity (NCV) model [13]. The values for σ_q^2 and σ_r^2 are shown in Table 3.5.

Sampling time	T	$2.5s$
Process noise covariance	σ_q^2	$(0.05ms^{-2})^2$
Measurement noise covariance	σ_r^2	$6m^2$
Probability of detection	P_D	0.9
Low clutter density	λ_{low}	$1 \times 10^{-5}m^{-2}$
High clutter density	λ_{high}	$5 \times 10^{-5}m^{-2}$

Table 3.5: Tracking system parameters

To test falsely initiated tracks, the clutter measurements were used exclusively at first to check if any false targets were initiated. Then the same set of measurements was simulated along with the target measurements to check if the real targets were discovered.

A successful scenario of all measurements where both targets were found is shown in Figure 3.2. The blue dots are all the measurements during runtime. The blue color shading of the measurements represent their timestamps. The brightest shades are from the beginning of the simulation, while a darker shade of blue means a measurement originating from a scan closer to the end of the simulation. The red lines show confirmed tracks.

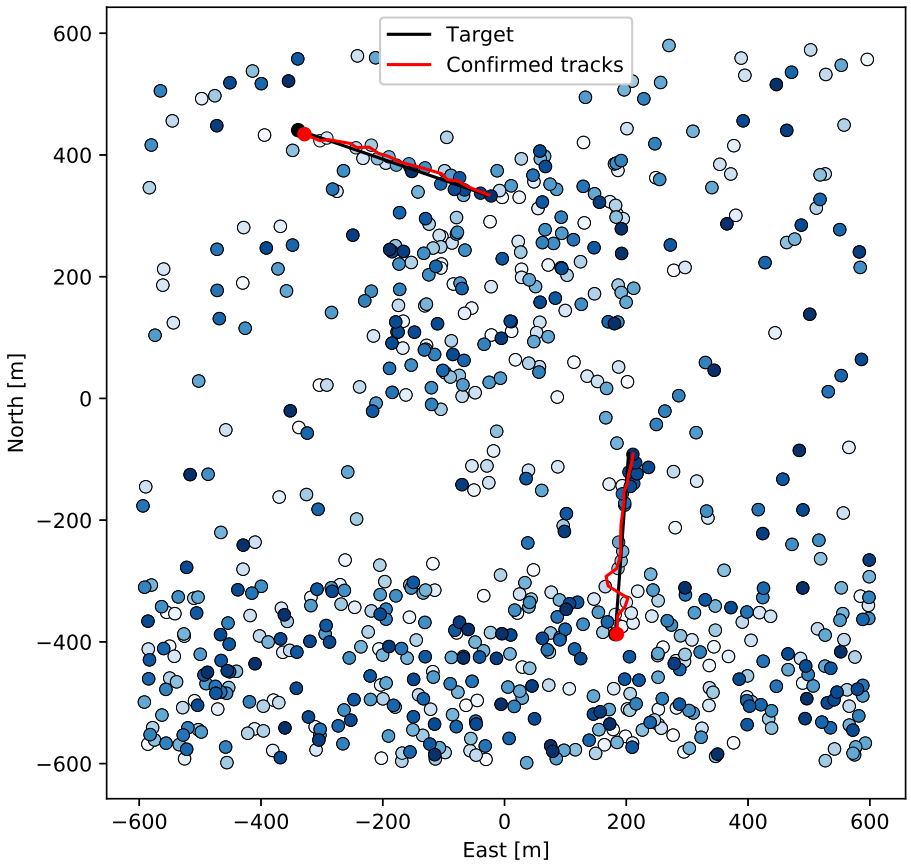

Figure 3.2: All measurements from an example run of the test setup plotted along with target trajectories and confirmed tracks.

3.5 Initialization time

As mentioned in Section 2.4.5, an initialization time might be beneficial for the performance of tracking using a clutter map. An appropriate number of initialization scans can vary from method to method. Different scan lengths were tested in [17] where it was suggested that approximately 200 scans is sufficient for all methods to reach convergence. Comparing the test setup and lowest density of $\lambda_{low} = 1 \times 10^{-5}$ for this thesis, the same number of scans would be 80. This is however not necessary for all three estimators, as some converge quicker than others.

The original test setup described in Section 3.4 was run using only clutter measurements to initialize the clutter maps. After initialization, the regular scenario was run with spawned targets.

Initial testing of initialized clutter maps was done using a relatively short initialization time and averaging length of 20 scans. The Spatial estimator was tested first, since it has the fastest convergence of the three estimators [17]. The initialized map is shown in Figure 3.3 along with a comparison to an uninitialized map. Comparing the gray high clutter areas in Figure 3.1 with the estimated clutter densities in Figure 3.3 we see that the initialized map is much closer to the ground truth than the uninitialized. This confirms that an initialization period will provide more accurate estimated clutter values, though it gives no further indication of what an optimal number of scans may be. Figure 3.4 shows that the two clutter maps are quite similar after ten scans, meaning that a total of 40 scans may be somewhat overly cautious for the Spatial estimator. Though, since the averaging length is also 40, it may still be beneficial to have a longer averaging window for better estimates.

The three estimators were tested for three chosen lengths of initialization and averaging lengths. 20 scans was picked as the shortest length which should favor the Spatial estimator, while 80 scans were chosen as the longest length to compare the estimators when they are all close to convergence values. 40 scans was picked as a middle ground between the two and most tests were conducted with this length. Unique initialized maps based on random clutter were created for each run of the simulator, and one example of each tested length is shown in Figure 3.5.

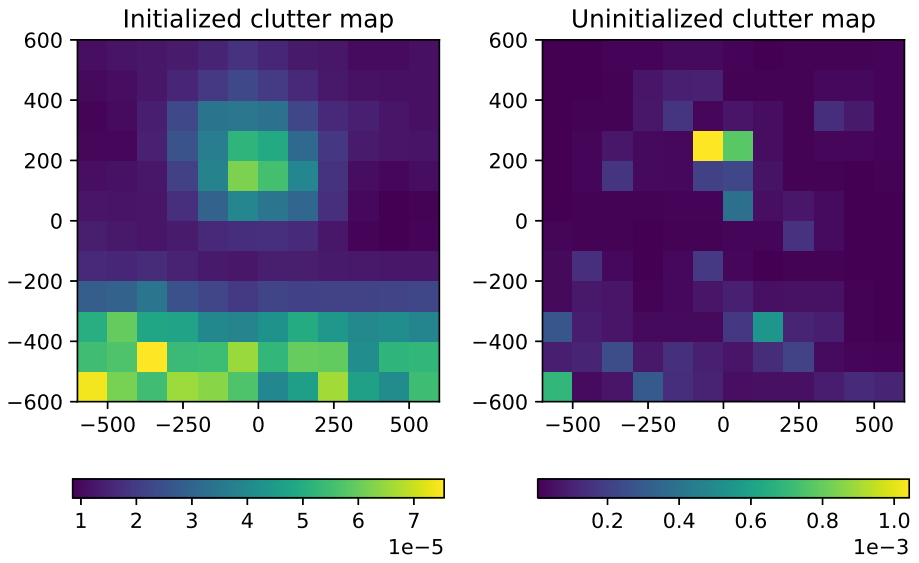

Figure 3.3: Spatial clutter map after one scan with averaging length $L = 40$. The initialization time for the initialized map was 40 scans.

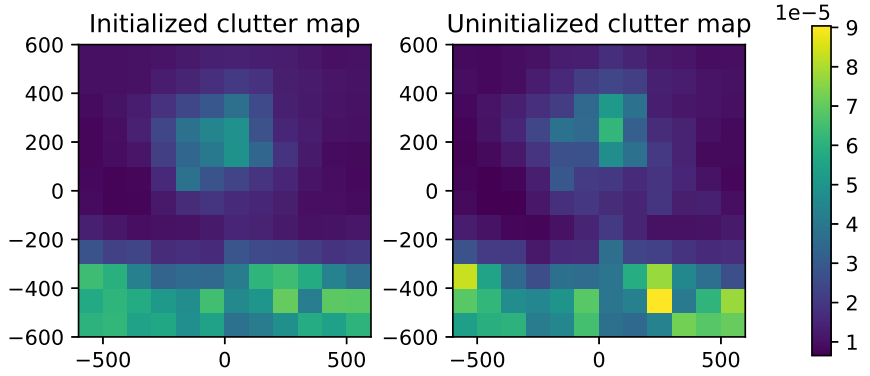

Figure 3.4: Spatial clutter map after ten scans with averaging length $L = 40$. The initialization time for the initialized map was 40 scans.

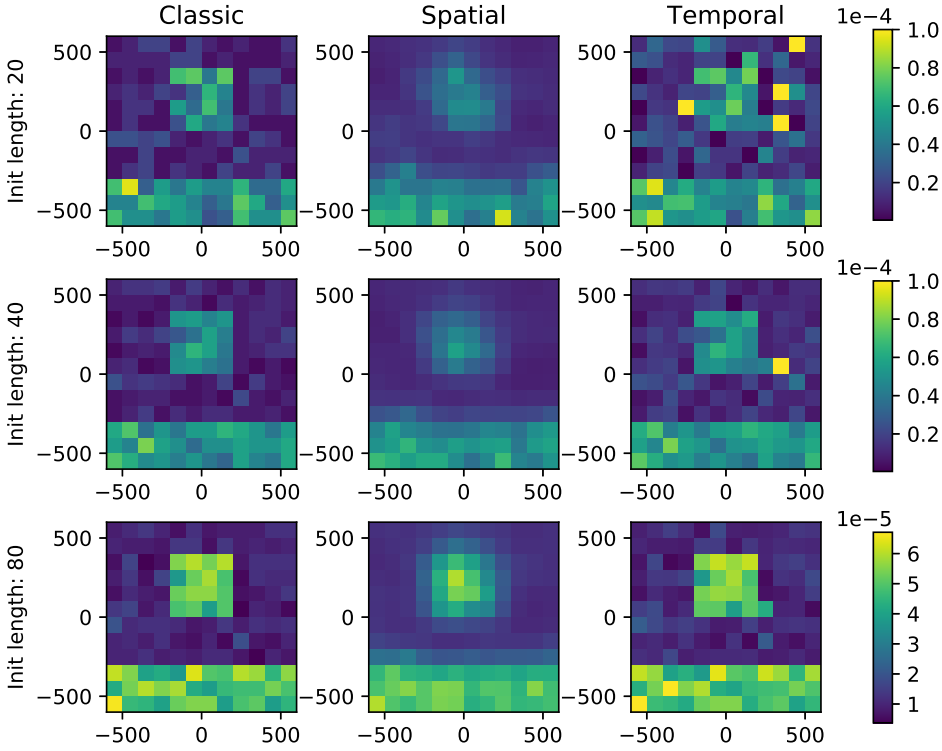

Figure 3.5: Initialized clutter maps of all tested initialization lengths.

3.6 Excluding target measurement from clutter estimation

As described in Section 2.4.5 excluding target measurements from the calculation of the clutter map, should provide a more accurate estimation of the real clutter density. The problem however, is that it may be difficult to be certain of which measurements are in fact from a real target in a tracking scenario. To test whether removing estimated target measurements would have any impact on initialization results, estimated target measurements from preliminary and confirmed tracks were removed from clutter estimation. The gated measurement with the highest β_k^i was assumed to be from a target, as long as it was higher than the likelihood of no target measurements in the gate, β_k^0. A proposed equation is presented as follows

$$\text{Target measurement} = \boldsymbol{z}_k^i, \quad i = \arg\max_i \beta_k^i, \quad \text{if}\Big(m_k > 0 \quad \text{and} \quad \beta_k^i > \beta_k^0\Big) \tag{3.3}$$

This exclusion method is done for all preliminary and confirmed tracks prior to updating the clutter map. The estimated target measurements are then discarded from the set of measurements sent to the clutter map for the current clutter estimation.

Chapter 4

Simulation Results

4.1 Clutter map bias

To test the biases from Section 2.4.4, a single cell of each type of estimator was simulated 10000 times with different number of mean clutter measurements falling into the cell. The idea was to compare the results with Figure 2.9 for the Classic and Temporal estimators, as well as examining simulated bias of the Spatial estimator.

From (2.27) we have that the expected number of clutter measurements in an area at a scan k depends on the clutter density λ and the area V_k. The mean number of measurements over L scans is gives as

$$x \triangleq V(c) \cdot L \cdot \lambda(c) \tag{4.1}$$

By setting $\lambda(c) = 1 \times 10^{-5}$ and $V(c) = 100m \times 100m$, we get an average of 1 measurement over 10 scans. The clutter density was adjusted to get the necessary results to reproduce the figures from [17]. The relative bias was calculated according to (2.38), and the results are shown in Figure 4.1. Each point on the graph represents an explicit test setup simulated 10000 times to get the resulting bias for a specific mean clutter measurement number.

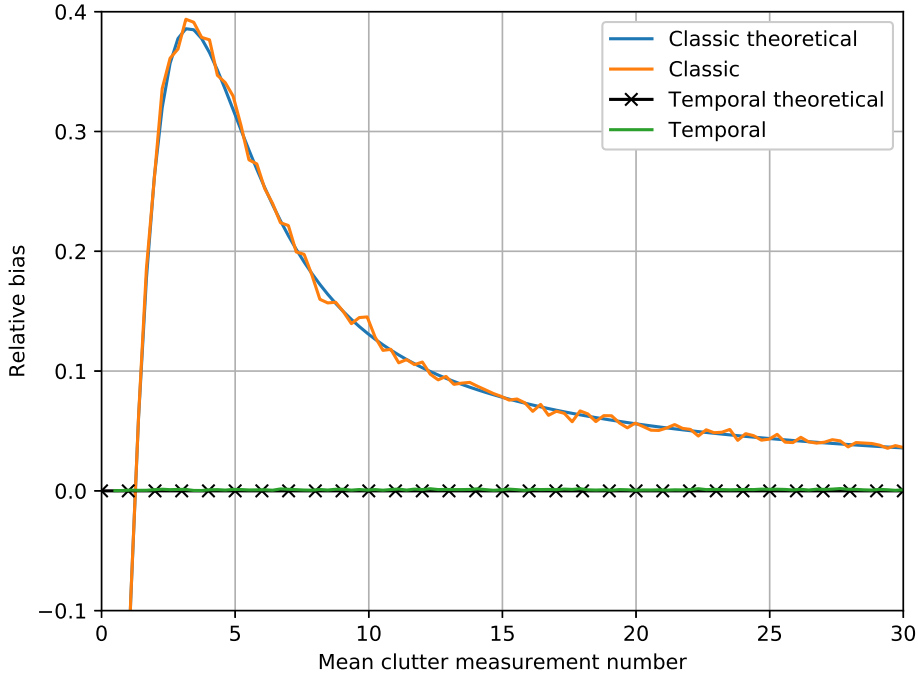

Figure 4.1: Classic and Temporal clutter map bias.

Both estimators show expected bias results concurring with the theoretical values. However, as stated in Section 2.4.4, the Temporal estimator does have some bias during initiation phase, which is not visible in Figure 4.1. The averaging length L also impacts the Temporal estimator. The estimator was run for 1000 times more timesteps than the Classic estimator in each run and an averaging window of $L = 1000$ to achieve the close to zero bias.

The amount of initialization bias from the Temporal estimator depends on the length of the initialization and the mean clutter measurement number. A variety of initialization lengths have been tested and are plotted in Figure 4.2.

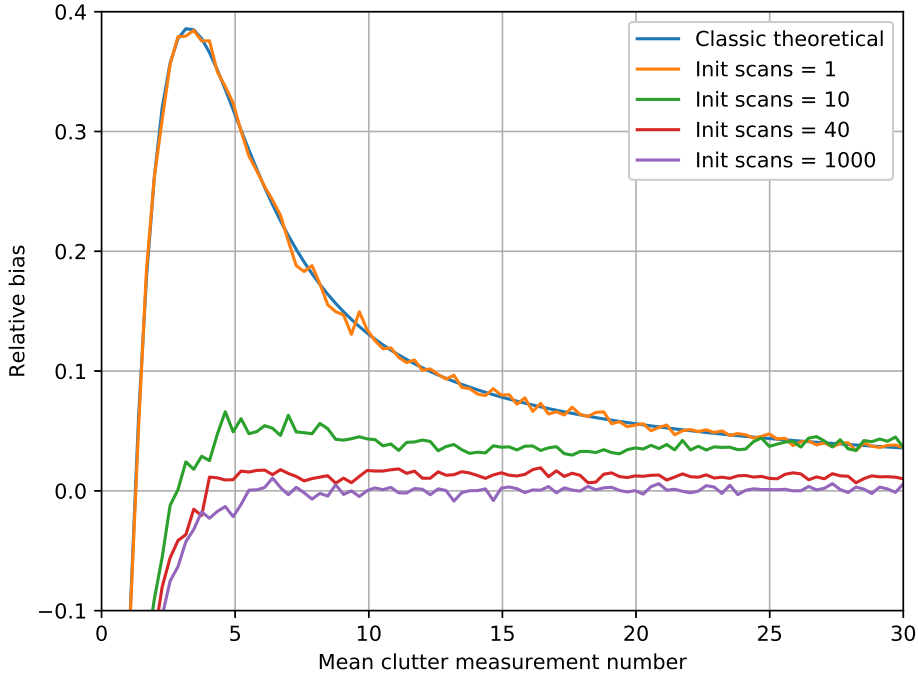

Figure 4.2: Temporal bias with different lengths of initilization.

Figure 4.2 shows that longer init time leads to a lower bias for high mean clutter measurement numbers. The bias seems to approach zero as the initialization time increases. Keep in mind that the clutter density is adjusted for each time length to get the correct mean number of measurements within the time window. This means that a mean number of 5 measurements for 'init time 10' is simulated with the same clutter density as a mean number of 20 measurements for 'init time 40'. An interesting observation with Figure 4.2 is that an initialization time of 1 scan equals the same amount of bias as the Classic estimator. In other words, the Temporal estimator already has a less biased clutter estimate than the Classic estimator at the second scan. The reason why all three lines goes towards negative infinity is due to the event that if no measurements have fallen in the cell, the estimator uses the current scan as the inter-arrival time to estimate the density. The estimated density value is then heavily biased to output a higher clutter density than ground truth for low mean number of measurements, just as the Classic estimator.

Negative bias for small clutter measurement numbers in the Temporal estimator will only occur during the beginning of a scenario and will deteriorate over time. To show the remaining bias of the averaging window size L as time progresses, the bias simulation was run for a number of time steps equal to $1000 \times L/(\text{mean measurements})$. The result is shown in Figure 4.3.

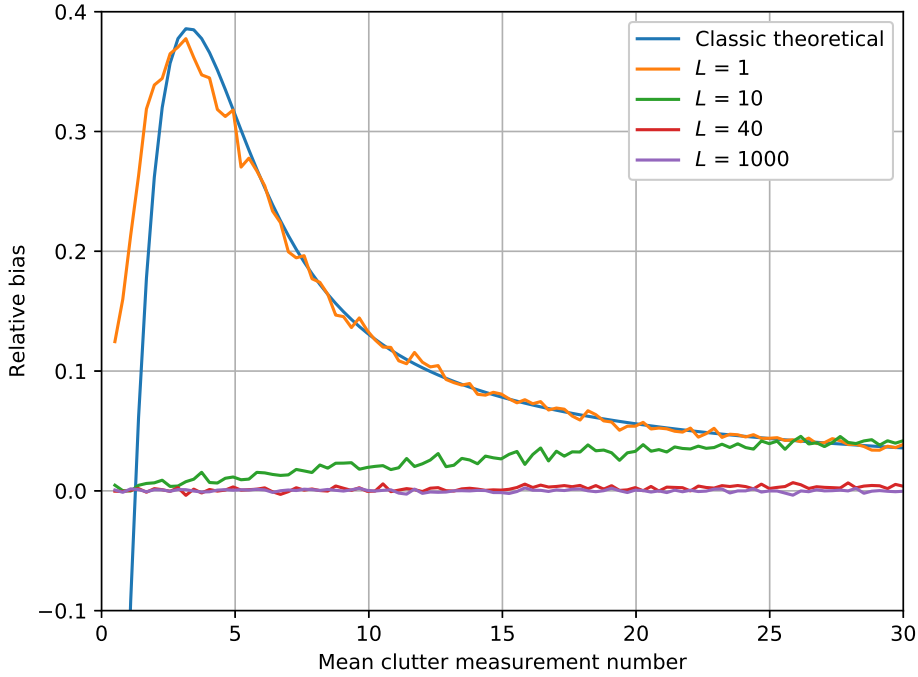

Figure 4.3: Temporal bias of different window lengths L for a longer simulation time.

Figure 4.3 clearly demonstrates that even as the mean clutter measurement number increases, there is still a bias in the Temporal estimator due to the averaging length L. Yet with a length of $L = 40$ the bias is already close to zero.

The Spatial estimator only has bias of the clutter density when cells are located near a discontinuity in real clutter density, as stated in Section 2.4.4. To demonstrate this bias, three cells from the test setup in Section 3.4 were simulated for 10000 scans with a range of window lengths L to get different number of mean clutter measurements (Figure 4.4).

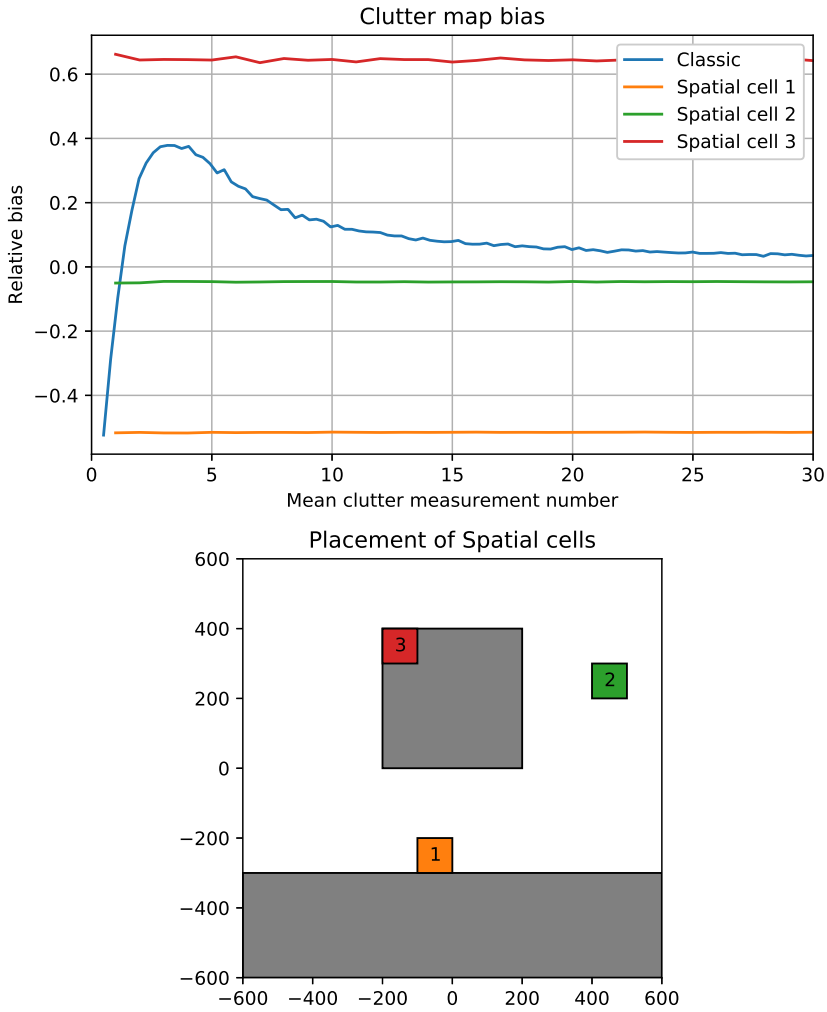

Figure 4.4: Comparison of spatial biases for three different clutter cells.

As we can see in Figure 4.4, the bias of a Spatial clutter map cell is heavily influenced by its proximity to an area of another clutter density. 'cell 1' and 'cell 3' have a significant amount of bias, since they are located on the border between two densities of $\lambda_{low} = 1 \times 10^{-5}$ and $\lambda_{high} = 5 \times 10^{-5}$. The negative bias for 'cell 1' in Figure 4.4 means it estimates higher clutter values than the real values, and the positive bias for 'cell 3' means lower values than the ground truth. This transition between regions of different density coincides with the theory from Section 2.4.4. Areas further away from such discontinuities in density, such as 'cell 2', are less influenced by these variations.

4.2 Performance measures

An important question when comparing techniques in a field of research is how to measure their quality against each other, as well as finding advantages and disadvantages with the application areas of the techniques. One should attempt to adjust performance measures to desired real world applications. Thus, test simulation should attempt to mimic real situations, and the findings should reflect suitability and reliability of a set of methods according to the problem we want to solve.

First, we want to test reliability of track initiation using different methods. Thus, we need clear criteria for classifying true and false initiated tracks.

A typical way of presenting performance results of a given component or a system is with Receiver Operating Characteristic (ROC), which will be the main way of presenting results in this thesis.

4.2.1 SOC analysis

ROC analysis shows the trade off between positive and negative detections in a binary classification problem as the discrimination threshold is varied. When testing a complete system instead of a single system unit, it can also be referred to as System Operating Characteristic (SOC). In terms of track initiation this can be used to show the trade off between true and false tracks detected. For the SOC tests in this thesis, the true positive rate is denoted P_{DT} for true track detection probability and the false positive rate is denoted P_{FT} for false track probability. This notation follows [2].

The SOC estimation was conducted the same way as in [19].

$$P_{DT} = \frac{\text{Number of targets detected}}{\text{Number of total targets}} \tag{4.2}$$

$$P_{FT} = \frac{\text{Number of false confirmed tracks}}{\text{Number of confirmed and terminated tracks}} \tag{4.3}$$

The false and true track probabilities were calculated individually from simulations with and without targets. The false track probability P_{FT} was calculated after running simulations with only clutter in the test setup. Therefore, the simulated targets had no impact on this value.

To test how the different clutter map methods compares against a regular non parametric IPDA, the test scenario was run 1000 times for each method with each given set of parameters. A specific random seed was chosen to create all measurement data for the tests, to make the results easier to compare. The initiation thresholds were set in the range $[0.85, 1.0)$, though most tests were done with thresholds above 0.98.

All SOC plots in this chapter are plotted with logarithmic axes. This is due to

the fact that a COLAV system needs to be as close to a fail-safe state as possible, and improving a value by a decade is thus highly relevant as a criteria for a secure COLAV system. Also, most values in for P_{DT} in this section are in the upper range of $[0.99, 1.0)$, and most values of P_{FT} are in the lower range of $[0.0, 0.01]$.

Since the clutter maps have been tested for a set of initialization lengths, the number of initialization scans is added in the label name. For instance, 'Classic 40' means the Classic estimator initiated with an initialized map from 40 scans of clutter measurements.

4.3 SOC results

In theory, a tracking system with prior knowledge of clutter density should perform better than a non parametric method. This is due to the fact that the non parametric method assumes a uniformly distributed number of clutter measurements, as stated in Section 2.3.2, and may therefore be more inclined to initiate false tracks in high clutter areas. It may also be slower to initiate true targets in low clutter areas compared to parametric solutions. Furthermore, a clutter map with knowledge of ground truth clutter densities would ideally perform better than an estimated clutter map. The results shown in Figure 4.5 corresponds somewhat with these theoretical assumptions, though there are some surprises.

Figure 4.5 and Figure 4.6 shows estimator performance after an initialization period of 20 scans. As expected, the non parametric IPDA performed poorer than the parametric ones. Spatial shows best results of the three estimators both in terms of SOC scores and initiation time. Spatial is also the estimator with the fastest convergence which could explain why it outperforms the other two estimators for an initiation time of only 20 scans. Classic and Temporal has considerable amounts of bias for the low clutter area after only 20 scans which was discussed and demonstrated in Section 4.1. Temporal appears to be slightly better at track initiation time than Classic.

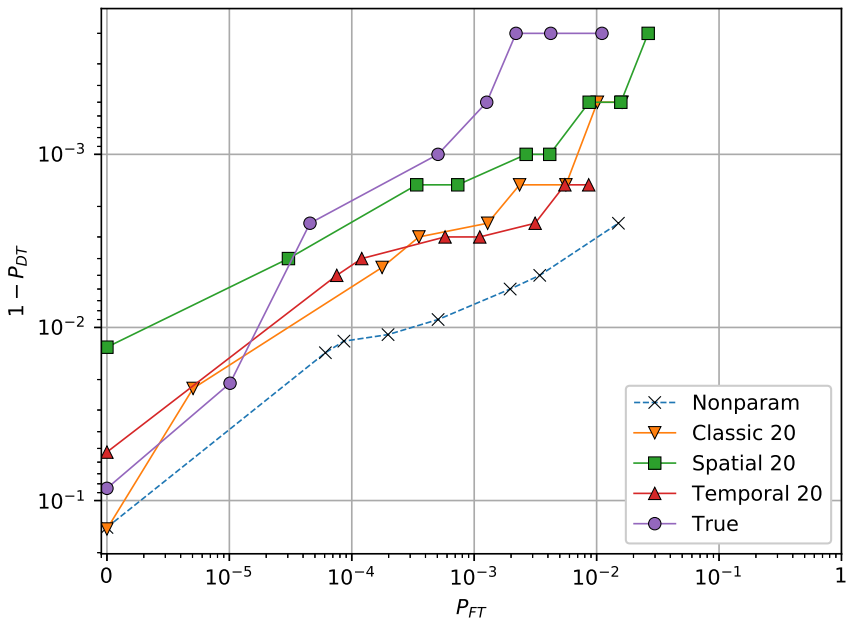

Figure 4.5: SOC performance of initialized clutter maps. $L = 20$.

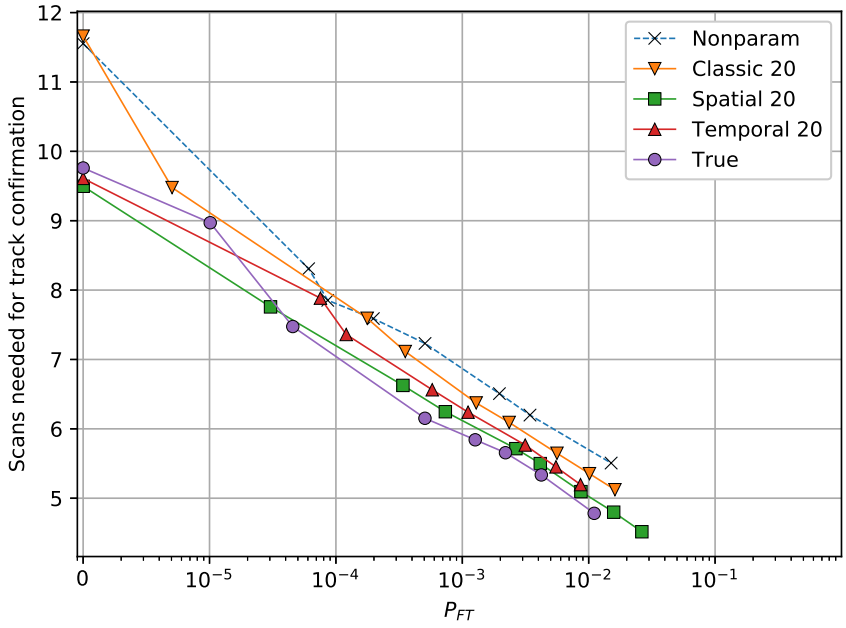

Figure 4.6: Track initiation time. $L = 20$.

The IPDA with a true clutter map provides mostly the best results, but quite surprisingly it has a sudden drop below other methods for the two highest test thresholds. One of the false initiated tracks for the True clutter map was examined to discover why this occurred. One of the missed true targets was also examined. The false initiated track scenario is shown in Figure 4.7, while the missed true target scenario is shown in Figure 4.8.

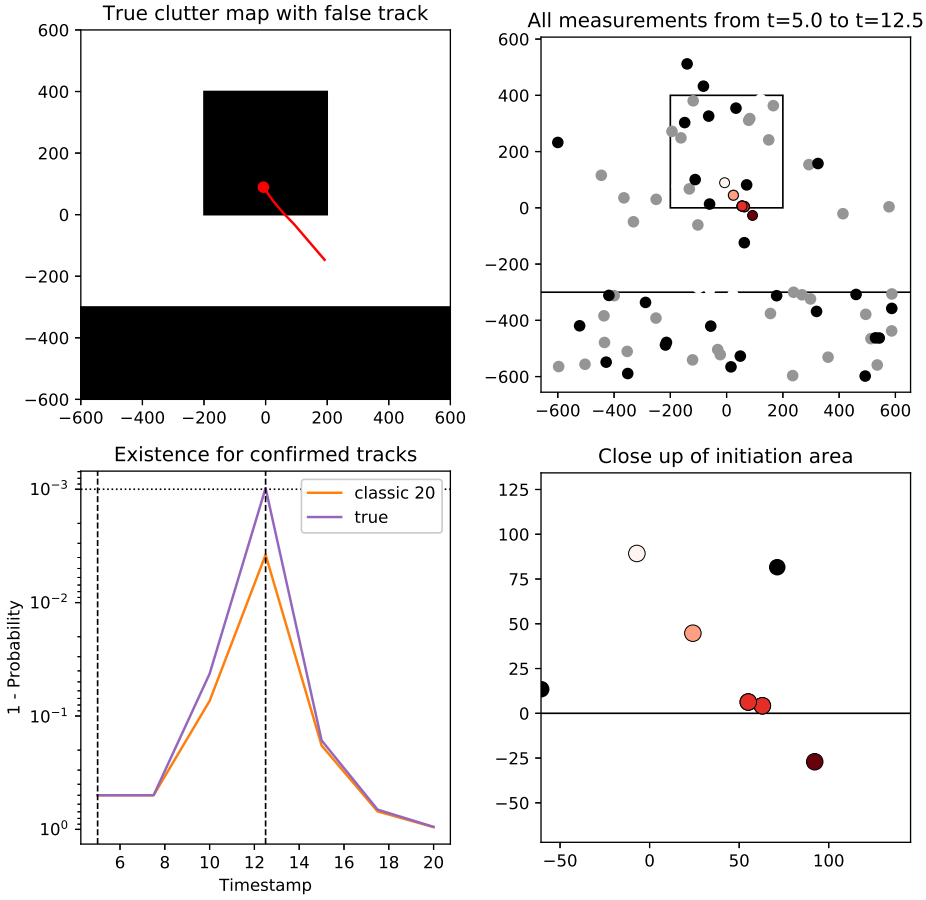

Figure 4.7: A false track initiation for True clutter map.

As we see in Figure 4.7, the IPDA with a True clutter map estimates a higher existence probability at timestamp = 12.5 than with a Classic estimator, and thus it confirms the track even at a higher initiation threshold. The Spatial and Temporal estimators performed similarly to the Classic estimator for this scenario, and did not initiate the track at the threshold of 0.999 which is shown with a dotted line in the figure. The circled measurements with a red color scale on the rightmost subplots show which measurements were gated in the preliminary track. In the time period $t = 5.0$ to $t = 12.5$ it appears as though there may be a target with course

towards the bottom right corner of the map, but as we see from the existence plot, the existence drops quickly in the following scans, and the false (preliminary or confirmed) track is terminated by all methods by $t = 20$. As we can see, the first three gates of the IPDA caught four measurements from the high clutter area, while the last measurement is from the low clutter area. Since the clutter measurement at $t = 12.5$ happens to align with the estimate of the preliminary track, the parametric solutions are primed to trust it more than the previous measurements due to a lower assumed clutter density. However, to understand why the True clutter map gets a higher spike than the other parametric solutions, we have to look into their estimated clutter density in that area. The estimated clutter densities are listed in Table 4.1 along with the existence probabilities.

Estimator	Clutter density t=10	Clutter density t=12.5	Existence Probability t=12.5
True	1.00×10^{-5}	1.00×10^{-5}	0.999
Classic 20	1.50×10^{-5}	2.50×10^{-5}	0.996
Spatial 20	2.34×10^{-5}	3.39×10^{-5}	0.997
Temporal 20	2.22×10^{-5}	2.31×10^{-5}	0.996

Table 4.1: Clutter and existence estimates for the false track from Figure 4.7 at low clutter cell area N={-100, 0}, E={0, 100}

We see from Table 4.1 that all the estimated clutter map methods have a higher estimated density than the true density in the area. Higher clutter density λ explains the lower existence probabilities, due to (2.31), which proves to be beneficial in this specific scenario. The reason for the estimators' high clutter estimates can be explained by their individual biases discussed in Section 2.4.4 and Section 4.1. The Spatial estimator estimates higher density since it is located close to a high density region, as can be seen for 'Spatial cell 1' in Figure 4.4. The Classic estimator estimator should not have any considerable bias based on a mean clutter measurement number of 2 in Figure 4.1, but does have a large variance for low clutter measurement numbers, which may explain its poor estimation in this scenario.

We also see from Table 4.1 that all three estimators estimate higher clutter density in $t = 12.5$ than in $t = 10$, especially Classic and Spatial. This shows that they still have high variance in estimated densities at this point. Variance will for all estimators be higher in low clutter areas since they have less information to base their estimates on. It also demonstrates that without any way of filtering measurements from a real target, the estimators will measure higher clutter amounts in areas where a target is moving. In this case the preliminary track was in fact false, so the higher estimates worked in their advantage.

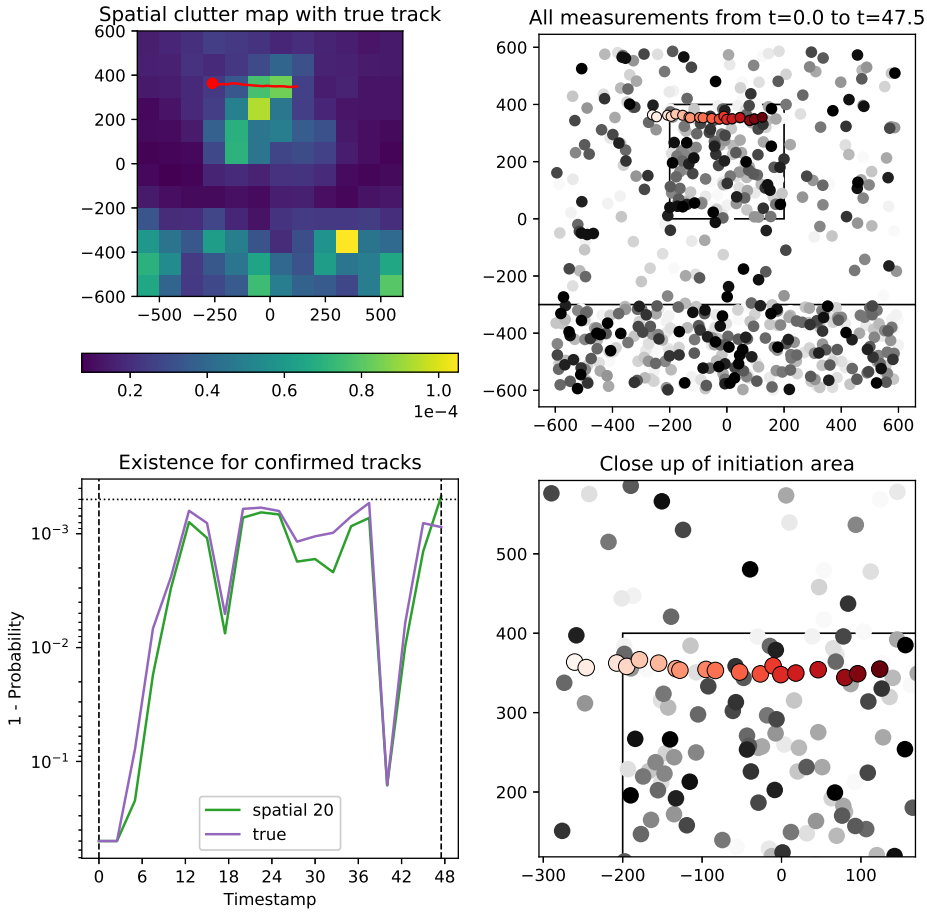

Figure 4.8: A true track initiation where Spatial clutter map is more certain than True clutter map. Only clutter measurements and upper target measurements are included.

Figure 4.8 shows a scenario where the Spatial estimator is able to initiate the track at a threshold as high as 0.9995, while the True clutter map is not. Scenarios in which the real target is moving just inside of the high clutter area for nearly the whole simulation, are the scenarios which the True clutter map struggles the most. These scenarios are mostly what is causing the sudden drop in the P_{DT} values for True clutter map for high thresholds in Figure 4.5. The clutter map estimators are able to initiate more of these targets due to their inaccuracies in clutter estimation, causing sudden spikes in existence probability. This is especially apparent for the Spatial estimator which has significant bias in corners of the high clutter area, demonstrated in Figure 4.4. We can see that the last existence probability spikes over the threshold in the last scan in Figure 4.8 due to entering a lower estimated clutter cell.

Other initialization lengths

The test scenario was also simulated with initialized clutter maps of 40 and 80 scans. Results from 40 scans of initialization are shown in Figure 4.9 and Figure 4.10. Both Classic and Temporal have improved their SOC values from the results from 20 scans of initialization (Figure 4.5), while Spatial show similar values as before. Otherwise the three estimators produce comparable SOC results for this initialization length. The track initiation time for Classic and Temporal has also improved, although Temporal is still slightly better, as shown in Figure 4.10.

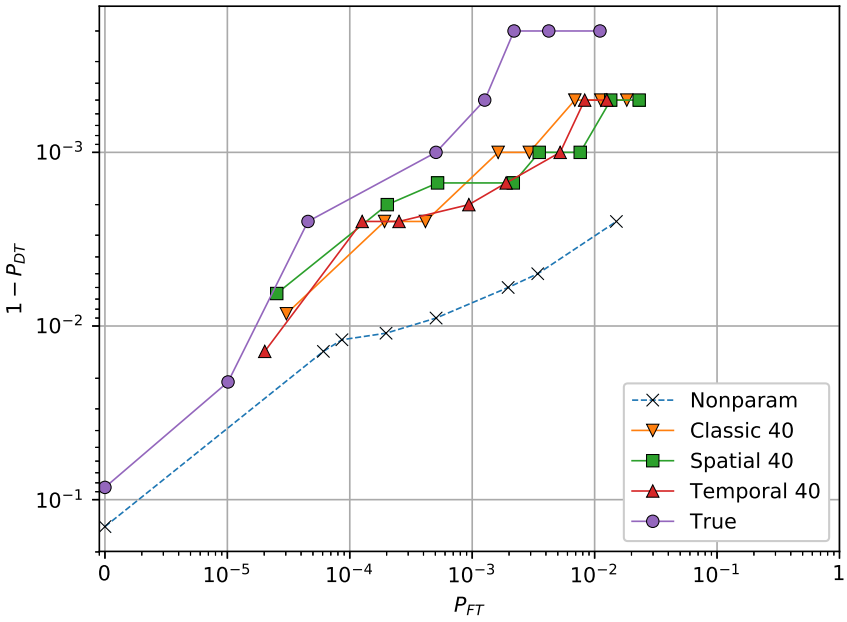

Figure 4.9: SOC performance of initialized clutter maps. $L = 40$.

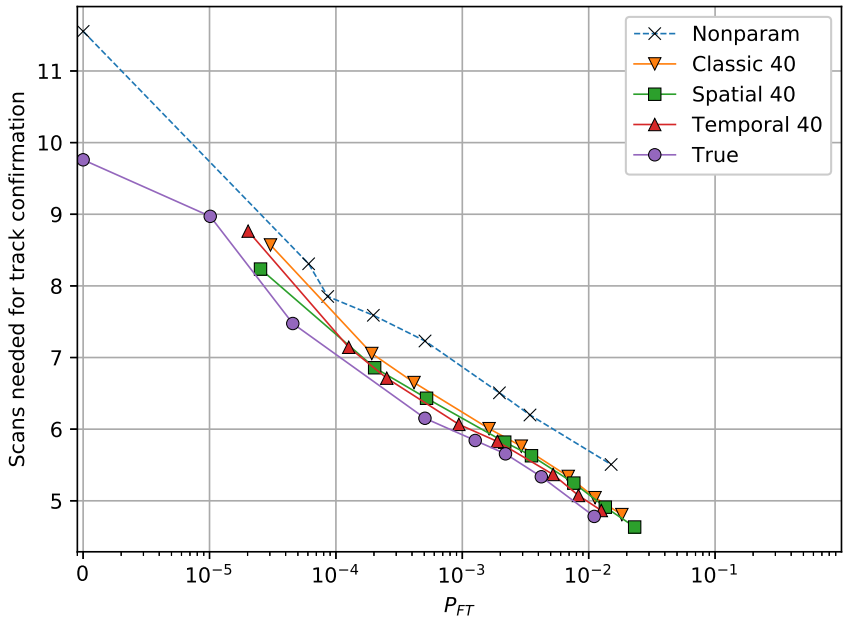

Figure 4.10: Track initiation time. $L = 40$.

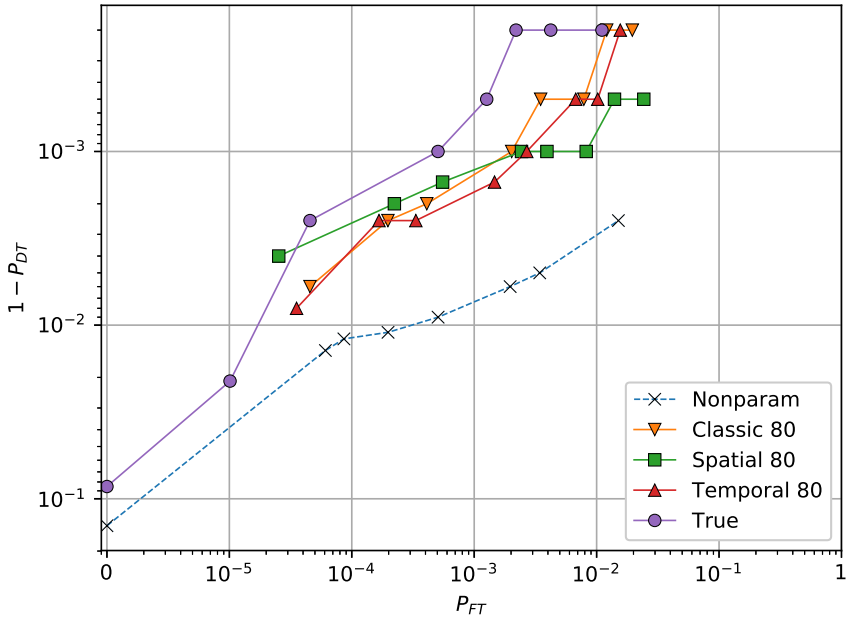

Figure 4.11: SOC performance of initialized clutter maps. $L = 80$.

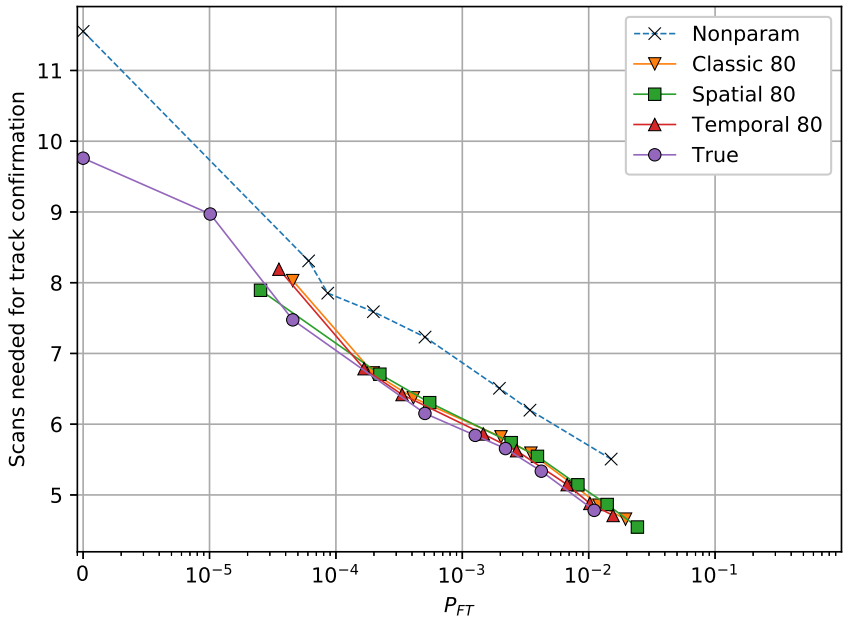

Figure 4.12: Track initiation time. $L = 80$.

Results from 80 scans of initialization are shown in Figure 4.11 and Figure 4.12. The Classic and Temporal estimator have slightly improved their SOC results from 40 scans (Figure 4.11) while Spatial remains approximately at the same values as for 40 scans. Spatial seems to have a flatter curve than the other estimators, being able to detect more targets at lower false track rates, but fewer targets at higher P_{FT} acceptance rates. Initiation times for all estimators seem to be almost identical to the True clutter map for P_{FT} values higher than 10^{-4}. At this length of initialization, all three estimators should be close to convergence and increasing initialization length further would not likely yield improved results.

4.3.1 Without initialization time

While initialization time can be helpful to provide initial estimates prior to starting a vessel at sea, this extra time may not always be available. Hence, the clutter map estimators were also simulated without initialization time. The Classic estimator shows the best results of the three estimators in both Figure 4.13 and Figure 4.14. Temporal has just barely better SOC results than the non parametric IPDA. As for initiation time, all three estimators are outperformed by the non parametric method.

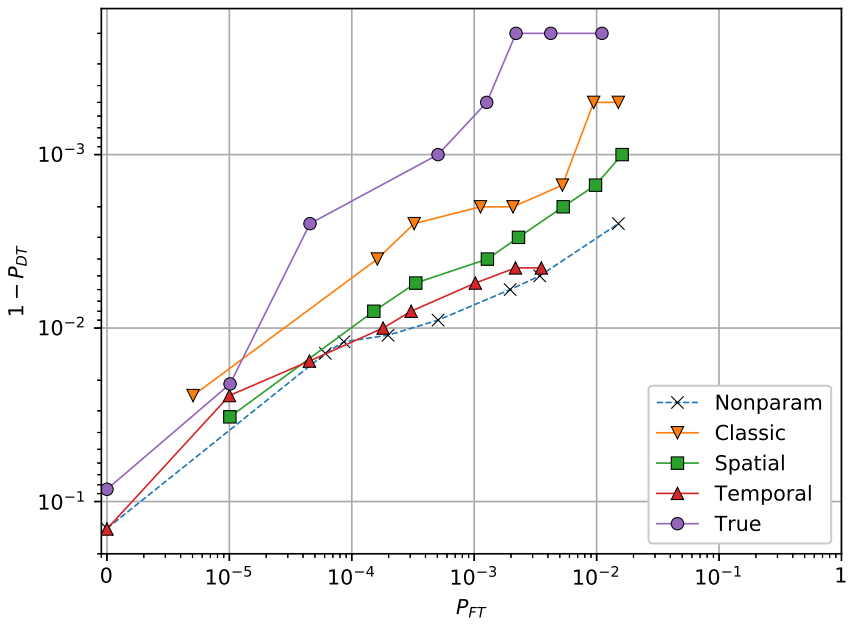

Figure 4.13: SOC performance without initiation time.

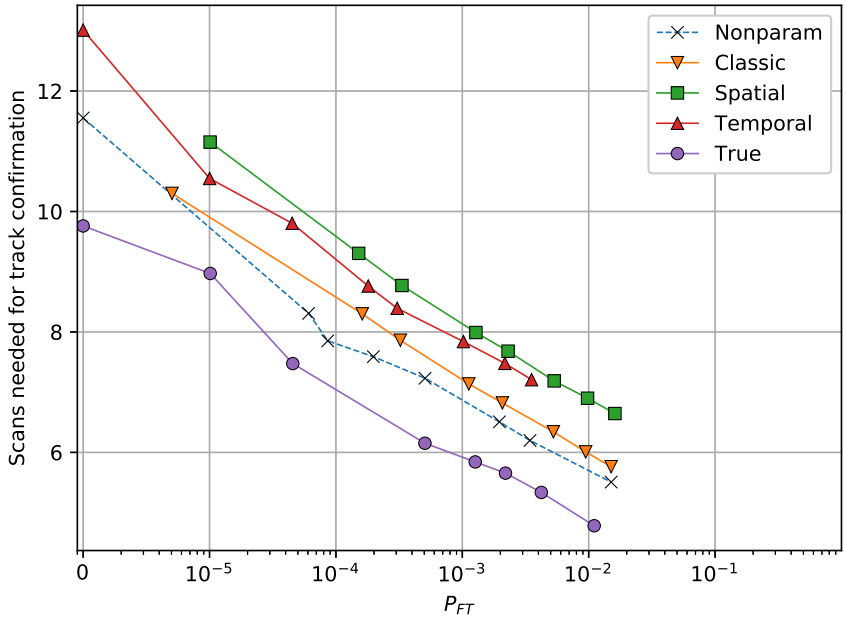

Figure 4.14: Track initiation time. No initialization for clutter maps.

The results from varying the initialization lengths is shown for each estimator individually in Appendix A demonstrating the points made about initialization time in another way.

4.3.2 Excluding target measurements

As stated in Section 2.4.5, excluding the target measurements from clutter estimation might be beneficial for the estimated clutter densities. In Section 3.6 a suggested method for removing target measurements was proposed. This method has been implemented and tested using the test scenario from Section 3.4. The results from the clutter maps estimated without the real target measurements are labeled 'wo target', while the maps estimated without the proposed target removal technique are labeled 'wo est target'. The results from 'wo target' are only achievable through simulation, as there is no way to be completely sure of what are target measurements or not in a real scenario. If one could pinpoint the target measurements with certainty there would be no need for clutter maps in target tracking.

In Figure 4.15 we see that 'wo target' is closer to the True clutter map results than the other results, which is expected. The target removal technique performs slightly worse than the estimator with all measurements except for some results with higher P_{FT} values.

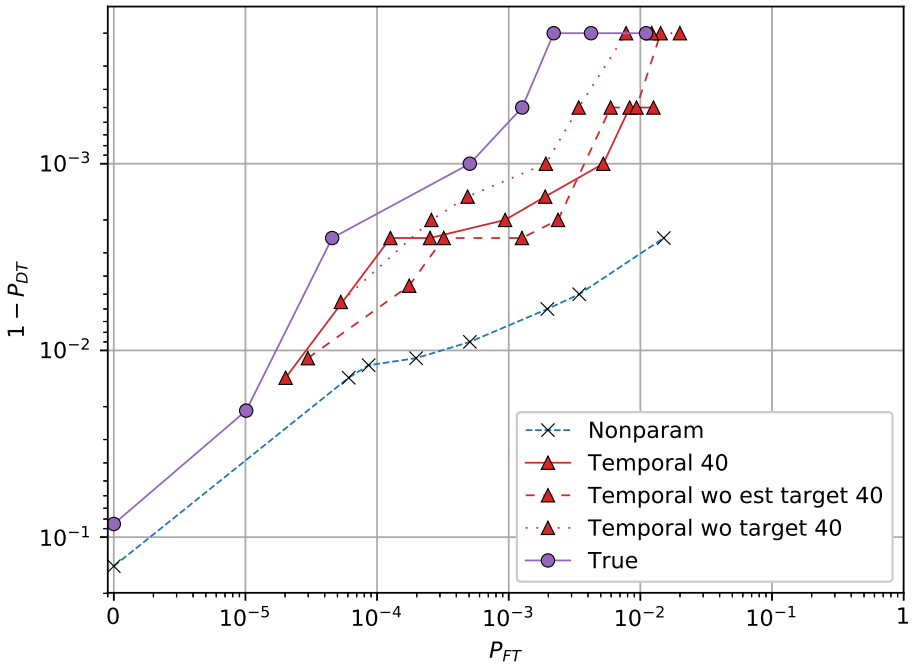

Figure 4.15: SOC performance for Temporal estimator with and without target measurements.

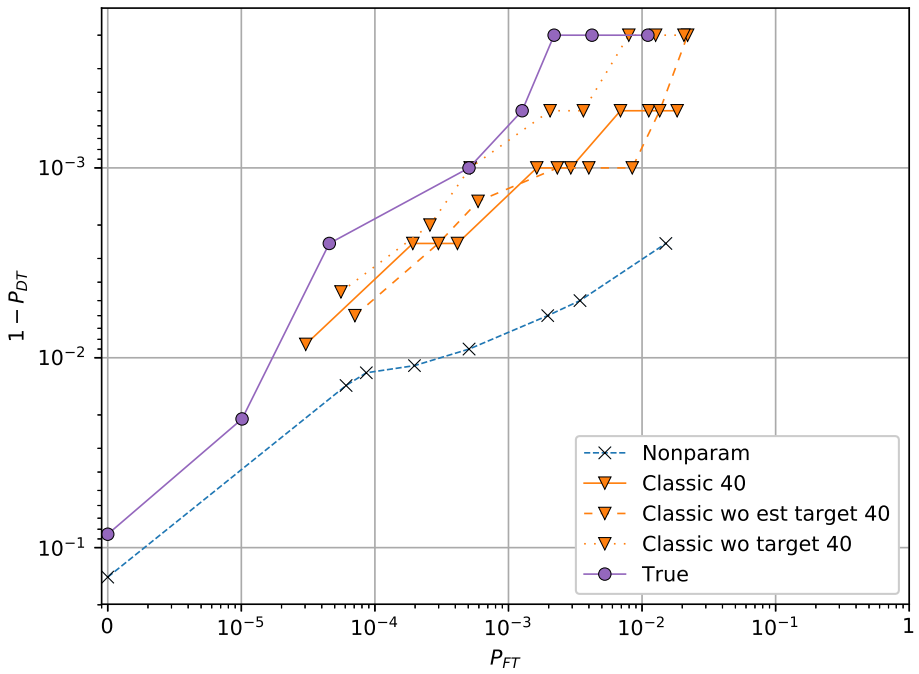

Figure 4.16: SOC performance for Classic estimator with and without target measurements.

The removal of estimated target measurements performs better for Classic, shown in Figure 4.16, but still shows reduced performance for some tests.

4.3.3 Uniform clutter performance

The test setup (Section 3.4) is set up to imitate the challenges of clutter variations in the map, which favors favors parametric methods over non parametric. Therefore the same setup was simulated, but with uniform clutter across the whole map. The uniform clutter was set to λ_{high} from the original test setup. This simulation was done to see if the parametric methods are still better than the non parametric method in uniform clutter.

This test was only done for initialized maps of length $L = 40$. The initialized maps are shown in Figure 4.17. Initialized maps of both high uniform with λ_{high} and low uniform with λ_{low} were created to compare the differences.

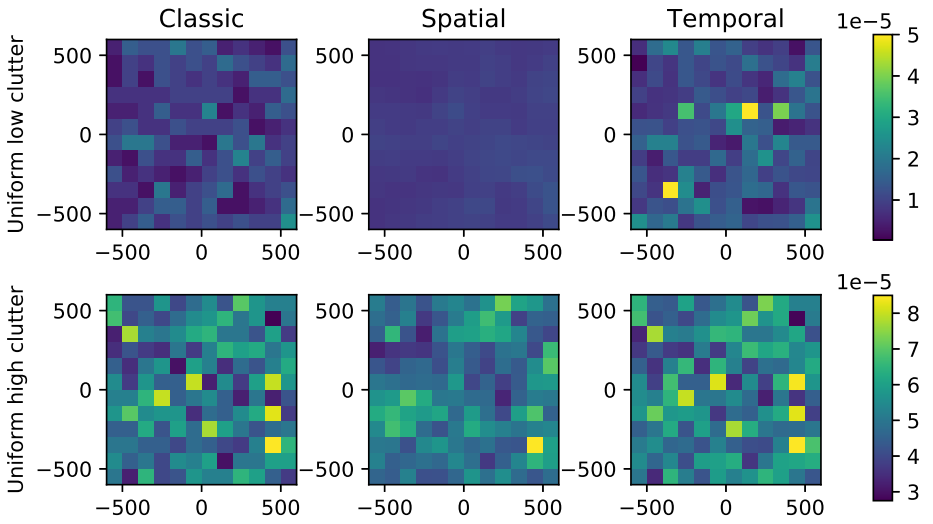

Figure 4.17: Initialized clutter maps in uniform clutter. $L = 40$. High clutter: 5×10^{-5}. Low clutter: 1×10^{-5}.

We can see in Figure 4.17 that the Spatial estimator produce the best results in uniform clutter, which is consistent with the theory from Section 2.4.4. The low clutter density equals a mean clutter measurement number of 4, which means a high positive bias for the Classic estimator in Figure 4.2. That explains the dark areas which are lower than the true density value. The Temporal estimator has a huge negative bias for values for mean numbers close to zero, which is why it estimates some very high density values in a few cells on the map. Nonetheless, Temporal is the only estimator which will on average improve estimate accuracy in subsequent scans, since the averaging windows of Spatial and Classic are already full.

The clutter maps based on high clutter show significantly better results, since it means a higher mean clutter measurement number and thus a lower bias.

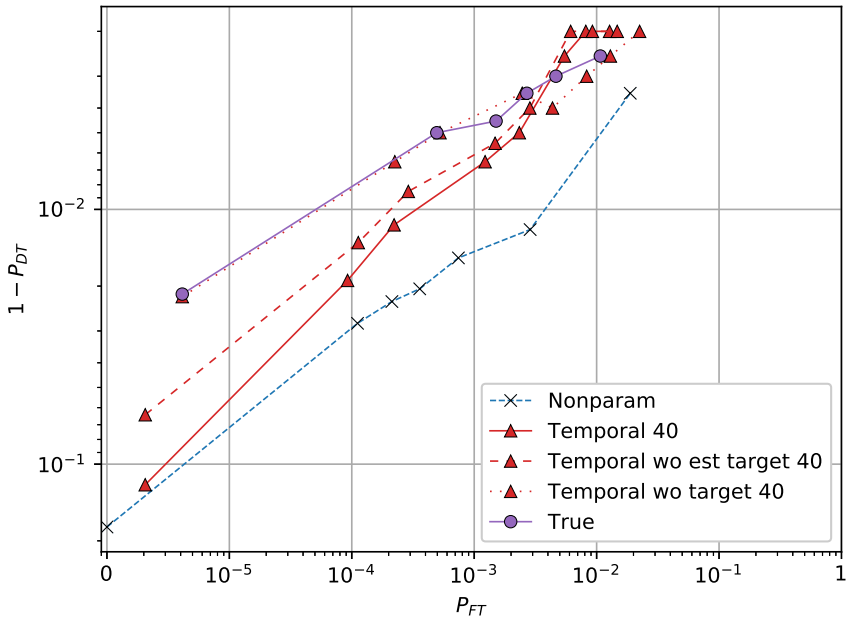

Figure 4.18: SOC performance in uniform clutter for Temporal.

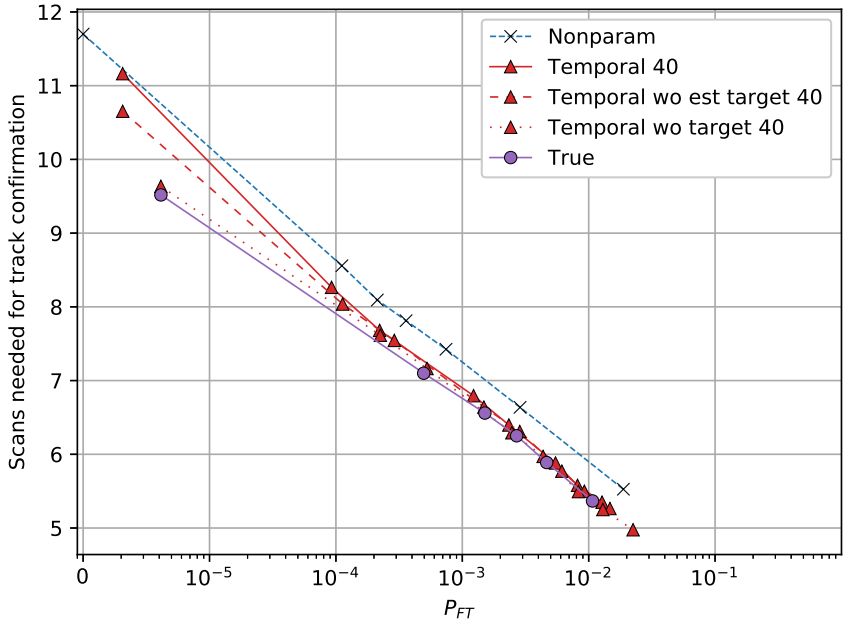

Figure 4.19: Track initiation time in uniform clutter for Temporal.

Figure 4.18 shows that the non parametric IPDA is still inferior to the parametric IPDAs. Nonetheless, it is the only method to achieve zero false targets for the highest tested initiation threshold. We see that the temporal estimator 'wo target' performs nearly equally to the True clutter map. This conforms with the bias results in Figure 4.2, where the bias is close to zero for this density. In this scenario the mean clutter measurement number equals 20 when $L = 40$.

Track initiation time results in Figure 4.19 seem to show the same order from best to worst, although it is difficult to tell them apart for higher P_{FT} values.

The Spatial SOC results in Figure 4.20 show even better results than the Temporal estimator, as expected for uniform clutter. The fact that they outperform the True clutter map may be because the clutter in the current scan are included in clutter density calculation, making them slightly adaptable to dynamic changes in clutter from one scan to the next.

The proposed method of excluding target measurements shows promising results for both the Spatial and the Temporal estimator. This suggests that using such a method may be beneficial in other scenarios as well.

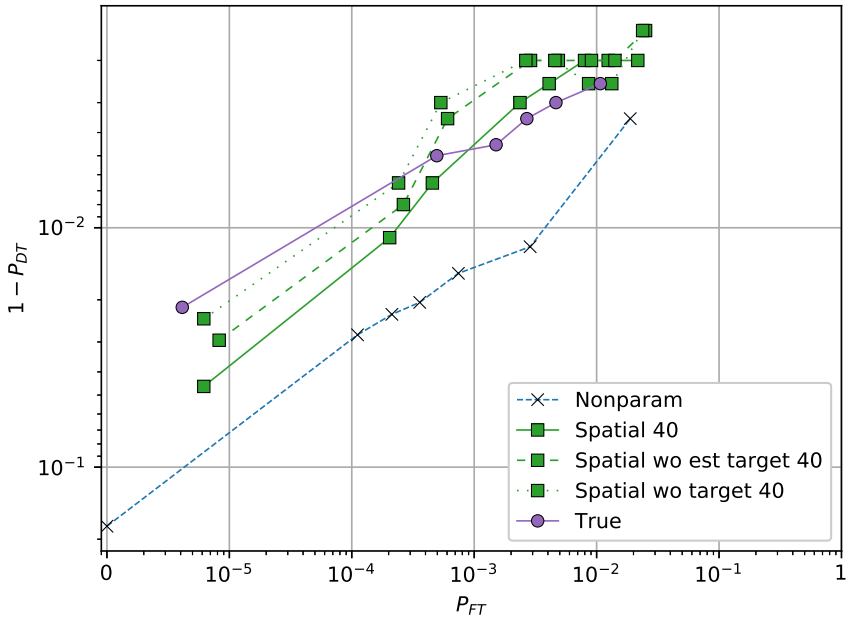

Figure 4.20: SOC performance in uniform clutter for Spatial.

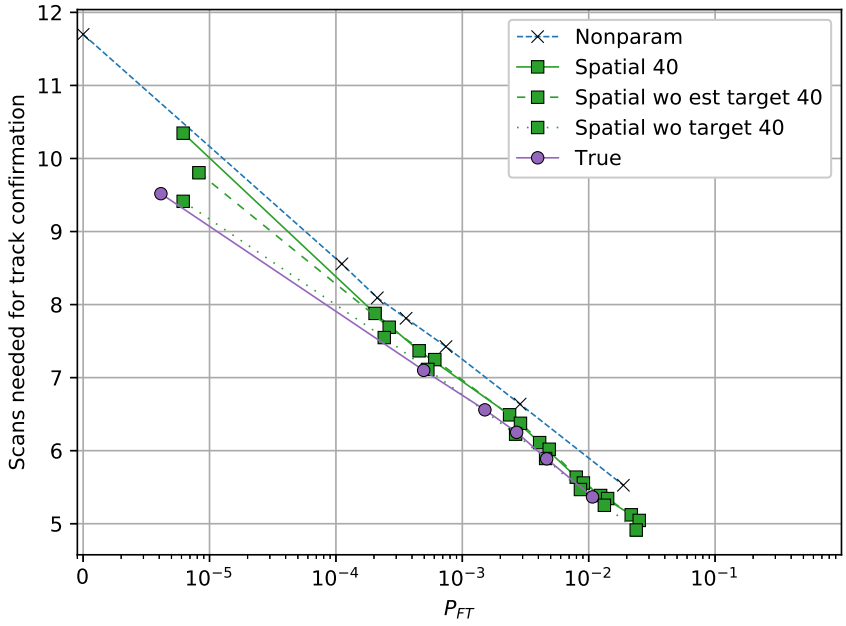

Figure 4.21: Track initiation time in uniform clutter for Spatial.

The initiation time results in Figure 4.21 are almost identical to the Temporal results in Figure 4.19.

Chapter 5

Results from real data

The Autosea project provided two real data sets of scenarios tested and recorded in Trondheimsfjorden:

- Seatex Drone
- Near Munkholmen and Approach

The two scenarios vary in terms of length, number of vessels, and clutter amount and distribution. The radar used in the experiments, shown in Figure 5.1, has a radar scan rate of approximately 2.5 seconds per scan. The autonomous vessel 'Telemetron', shown in Figure 5.2, is used as the ownship in both scenarios with the radar on board. It should be noted that the position of the radar is therefore non stationary in an earth-fixed reference frame in contrast to the the simulator which would be more comparable to a stationary radar. The measurement detection probability in the scenarios was seemingly poorer than in the simulator, so the detection probability in the IPDA was readjusted to $P_D = 0.8$. Some targets in the 'Near Munkholmen and Approach' scenario had more maneuverability than the NCV model used for targets in the simulator, so target noise covariance was readjusted up to $\sigma_q^2 = (0.5ms^{-2})^2$ in that scenario. This noise covariance value was set according to testing done in [20].

Figure 5.1: The radar used to obtain measurements for the scenarios 'Seatex Drone' and 'Near Munkholmen and Approach'.

Figure 5.2: The ownship, labeled 'Telemetron', in the scenarios 'Seatex Drone' and 'Near Munkholmen and Approach' with the radar from Figure 5.1 on board.

Since there was only one set of test data in each scenario, the results are not shown with SOC plots and initiation time plots. Instead, each tracking setup and parameters were tuned according to each scenario to get the best results where they found the known targets.

The scenarios were also tested with clutter estimation posterior to updating track estimates, meaning that the clutter estimates used in estimate update was based on N_{k-1} as in Musicki's theory [17]. The results did not change noticeably by this algorithm adjustment, and was thus not included in the results in this chapter.

5.1 Scenario: Seatex Drone

This scenario was done as a COLAV experiment with 'Telemetron' as the ownship trying to avoid Kongsberg's autonomous drone 'KSX_0SD1' on its path southwards. Radar measurement data and AIS data from the relevant targets are shown in Figure 5.3. The two targets are shown in Figure 5.4.

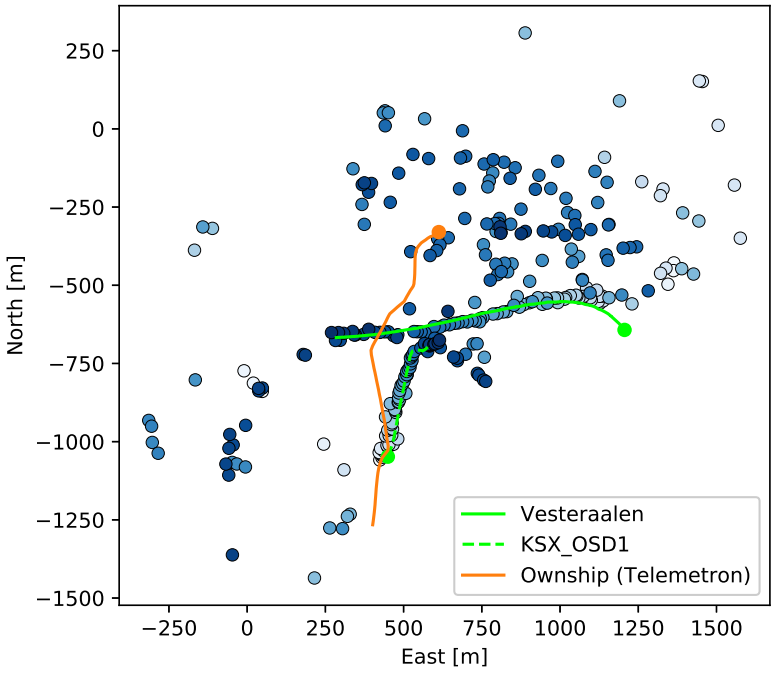

Figure 5.3: Raw radar and AIS data from the real scenario referred to as 'Seatex Drone' recorded with AIS and radar.

Figure 5.4: The two targets present in the scenario. The top is labeled 'KSX_0SD1' and was the target of a coordinated COLAV experiment, the bottom is a passing ferry labeled 'Vesteraalen'.

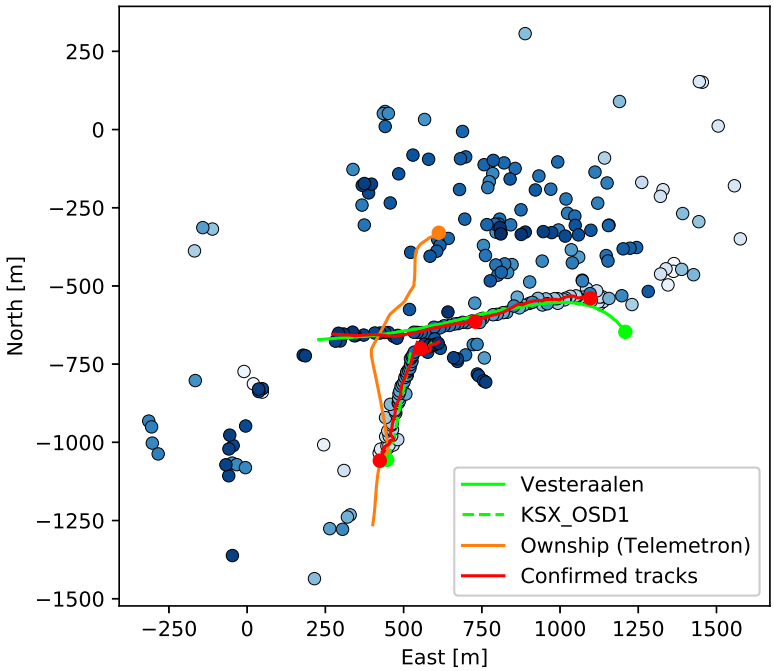

Figure 5.5: Confirmed tracks from a non parametric IPDA running the 'Seatex Drone' scenario shown in Figure 5.3.

Figure 5.5 shows a successful run of a non parametric IPDA initiating only true tracks. The reason why tracks are dropped only to be re initiated can be explained by the fact that the implemented IPDA in this thesis uses a Markov chain one, which does not take into account that the target may not be observable for parts of the tracking process. Using a Markov chain two after initiating a track could assist in this problem (Section 2.3.2) but since we are focusing on track initiation in this thesis, we will continue to use Markov chain one. The process noise covariance following a NCV model used in this scenario might also contribute to premature track termination, if the targets show a sudden change in course or velocity. Setting a higher value for process noise covariance could solve this issue, but would also make the track trajectories more exposed to measurement noise. It could also lead to more false initiated tracks.

The results are presented in Table 5.1. The initiation times are listed for the detected targets in this order: [KSX_0SD1, Vesteraalen].

Method	Init threshold	N	Cell size	Number of false tracks	True track conf. time (scan nr.)
Non parametric	0.90	-	-	2	[4, 15]
Non parametric	0.94	-	-	**0**	**[4, 16]**
Non parametric	0.99	-	-	0	[5, 17]
Classic	0.90	20	$50m \times 50m$	3	[5, 16]
Classic	0.94	20	$50m \times 50m$	**0**	**[5, 16]**
Classic	0.99	20	$50m \times 50m$	0	[9, 26]
Classic wo est target	0.90	20	$50m \times 50m$	6	[4, 15]
Classic wo est target	0.99	20	$50m \times 50m$	0	[5, 17]
Spatial	0.99	20	$50m \times 50m$	6	[14, 14]
Spatial	0.999	20	$50m \times 50m$	1	**[15, 15]**
Spatial wo est target	0.99	20	$50m \times 50m$	8	[4, 14]
Spatial wo est target	0.999	20	$50m \times 50m$	1	**[4, 15]**
Spatial wo est target	0.9999	20	$50m \times 50m$	0	[5, 15]
Temporal	0.99	20	$50m \times 50m$	**1**	**[16, 25]**
Temporal	0.995	20	$50m \times 50m$	0	[24, 26]
Temporal wo est target	0.99	20	$50m \times 50m$	**1**	**[9, 17]**

Table 5.1: Tracking results from the scenario 'Seatex Drone'

The non parametric IPDA performs very well in this scenario, and it seems that the addition of a clutter map does not improve the results. In fact, the parametric methods actually perform worse in most cases, and need higher initiation thresholds to filter the false tracks. Still, the Classic clutter map is able to filter all false tracks with a threshold of 0.94 and still has nearly identical confirmation times for the true tracks. Temporal only initiates one false track with a threshold of 0.99, but its initiation time for 'KSX_0SD1' is not comparable to that of Classic and non parametric.

Increasing the averaging length N or using larger cell sizes than $50m \times 50m$ did not show improvement for any of the estimators.

Removing estimated target measurements shows an improvement for the Spatial estimator and the Temporal estimator. They do not initiate fewer false tracks, but the confirmation times show huge improvement. 'Spatial wo est target' reduce the initiation time for the target 'KSX_0SD1' from 15 to 4 scans for an initiation threshold of 0.999. It is also able to maintain this improved initiation time for a threshold as high as 0.9999 where it has no false tracks. The Temporal estimator significantly reduce initiation time for both targets while still only initiating one false track.

The reason why the non parametric IPDA performs so well compared to the parametric IPDA's, can be somewhat explained by the change in clutter densities throughout the scenario. By comparing Figure 5.6 and Figure 5.7 we see that a field of clutter appears in the top center of the map in the second half of the scenario. This results in low estimated clutter densities in this area from the first scans, which is only slowly changed to higher estimated densities as the clutter appears in the second half of the scan set. False tracks may easier initiate when these low estimates from the first half are still apparent in the averaging windows of the estimators. Rapid changes in clutter densities are therefore not suitable for the clutter map estimators.

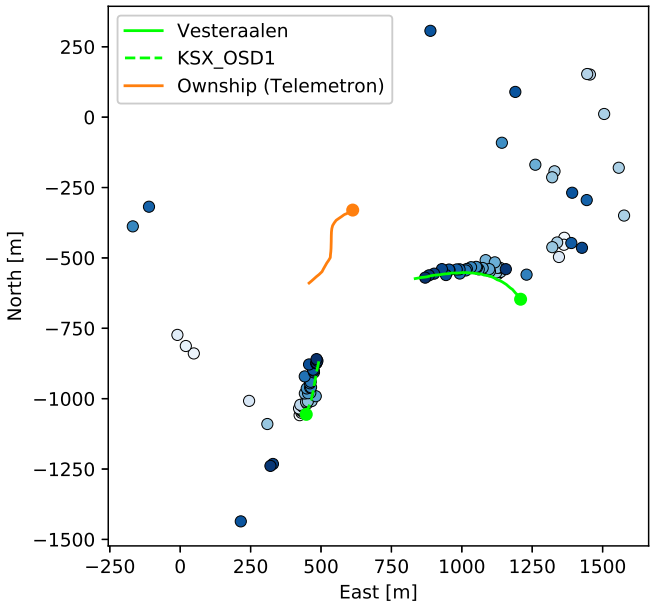

Figure 5.6: First 37 scans of the real scenario referred to as 'Seatex Drone'.

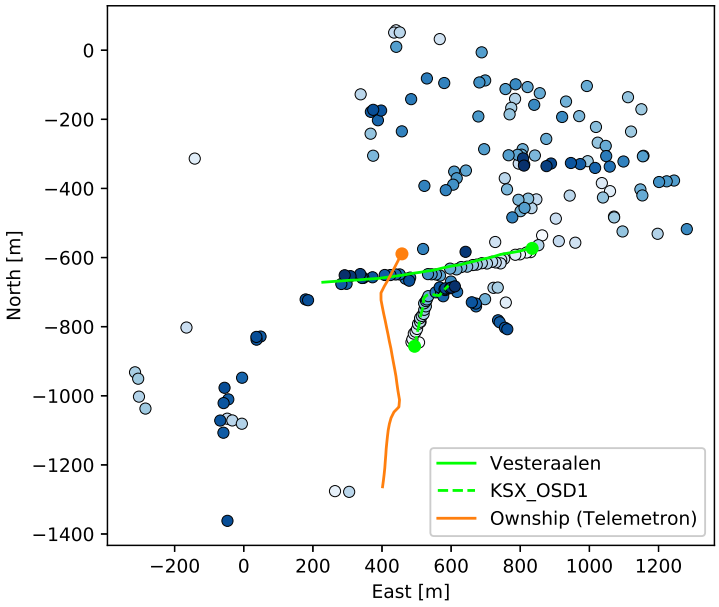

Figure 5.7: Last 37 scans of the real scenario referred to as 'Seatex Drone'.

5.2 Scenario: Near Munkholmen and Approach

This scenario is almost 10 minutes long and covers a larger area than the 'Seatex Drone' scenario. The data set consists of 234 radar scans along with AIS data for vessels in the measurement area. The measurement set and AIS data are plotted in Figure 5.10. Since the set consists of a significant amount of scans, the first and second half of the set are shown in Figure 5.11 and Figure 5.12 for a clearer view.

The vessel 'Autosea target vessel' was a relatively small private pleasure craft used to test the tracking and COLAV capabilities of 'Telemetron', and is shown in Figure 5.8. There were other vessels with AIS present in this scenario which were not affiliated with the Autosea experiment. These vessels are shown in Figure 5.9.

Figure 5.8: The vessel labeled 'Autosea target vessel' in the 'Near Munkholmen and Approach' scenario.

Figure 5.9: Three of the vessels present in the AIS data for the 'Near Munkholmen and Approach' scenario. The top is a rowing boat labeled 'Aerora', the middle is a high speed craft labeled 'Ladejarl' and the bottom is a passenger vessel labeled 'Nidarholm' (© Joakim Eide).

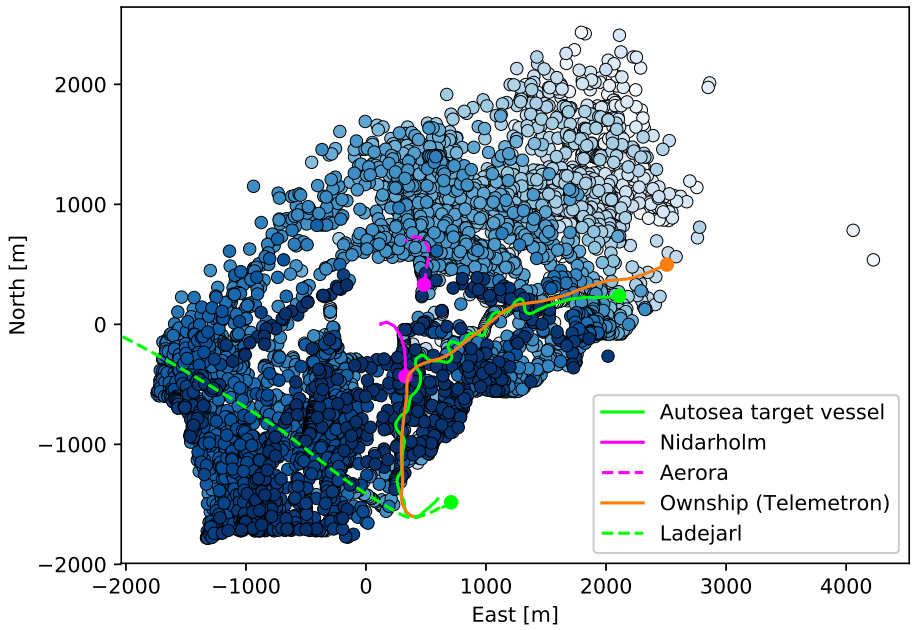

Figure 5.10: Raw radar and AIS data from the real scenario referred to as 'Near Munkholmen and Approach'.

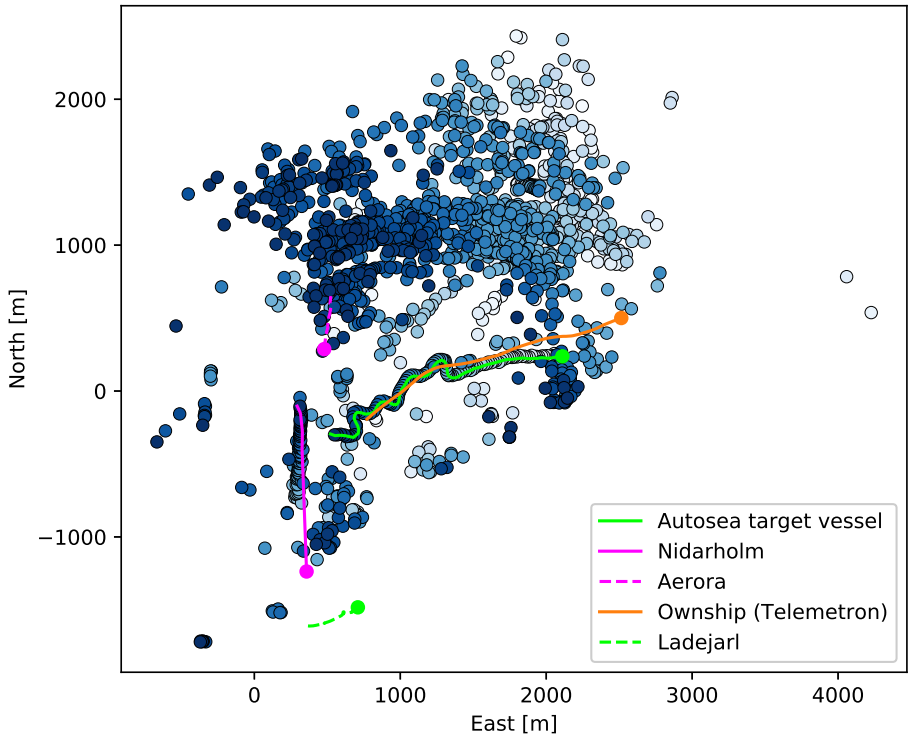

Figure 5.11: First 117 scans of the real scenario referred to as 'Near Munkholmen and Approach'.

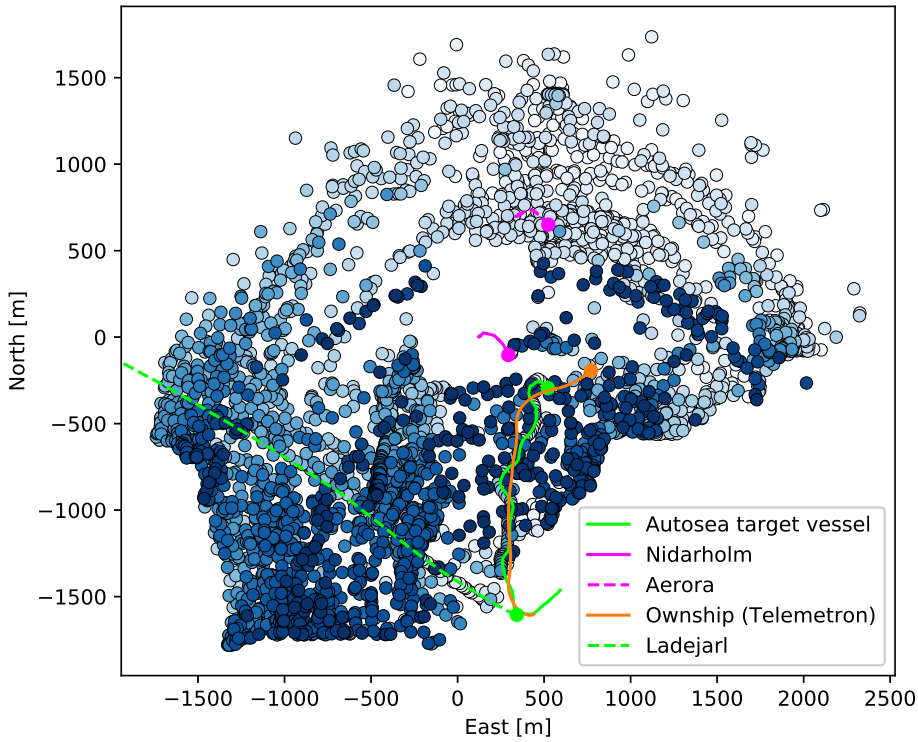

Figure 5.12: Last 117 scans of the real scenario referred to as 'Near Munkholmen and Approach'.

The AIS target 'Aerora' was not found for any parameter setup, and does not seem to be detected in much of the radar scans. This is only a small rowing boat, as we can see from Figure 5.9. In addition, it moves close to shore and more than $500m$ away from the ownship (Figure 5.11), making it harder for the radar to detect it. The setup was therefore tuned only enough to find the other targets in the scenario.

Measurements from land has been filtered out in this scenario by using a nautical map, in the way discussed in Section 2.1.4. The ship 'Nidarholm' can be seen arriving at the island 'Munkholmen' in Figure 5.12, and the ownship and 'Autosea target vessel' arrive at shore in the end of the scenario. Unfortunately, there was not sufficient time to implement the clutter map in combination with the nautical map, so the clutter estimators along the coast may be biased by this fact. The volume $V(c)$ used in the Classic and Temporal estimator should ideally be adjusted for cells containing masked land. Without any readjustment, these cells are biased to estimate lower clutter values. However, the lack of a nautical map may affect the Spatial estimator the most as it should ideally remove volume expansion in the masked areas (Figure 2.8). All cells with volume expansion which sometimes occur in masked areas are therefore affected by this issue.

One track which was usually initiated by all tested tracking methods appear to the naked eye to be a real target. However, since we have no AIS data from this target and no other data to confirm whether it was a vessel or just clutter, the tracks initiated for these measurements are neither counted as false or true. The potential target measurements are shown in Figure 5.13. This also demonstrates the difficulty of determining real targets from radar measurements alone.

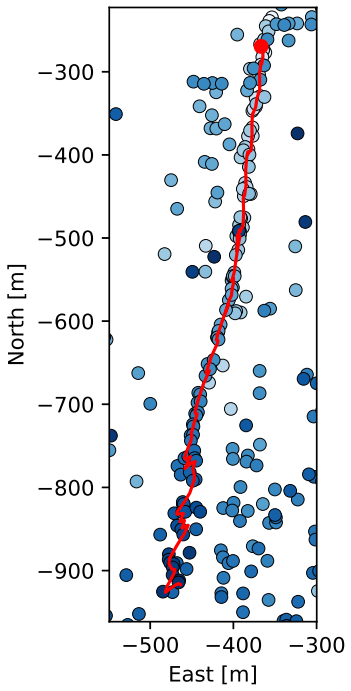

Figure 5.13: Measurements from scan 100 to scan 234 showing a potential target without any associated AIS data.

The results from the testing of non parametric vs parametric IPDA's show promising results for some of the estimators. The results are presented in Table 5.2. The initiation times are listed for the detected targets in this order: [Autosea target vessel, Nidarholm, Ladejarl].

Method	Init threshold	N	Cell size	Number of false tracks	True track conf. time (scan nr.)
Non parametric	0.99	-	-	19	[6, 69, 122]
Non parametric	**0.998**	-	-	6	[9, 70, 123]
Classic	0.985	20	$50m \times 50m$	**12**	**[10, 69, 124]**
Classic	**0.998**	20	$50m \times 50m$	**1**	[10, 70, 128]
Classic	**0.998**	40	$50m \times 50m$	**6**	[10, 70, 126]
Classic	**0.998**	40	$100m \times 100m$	**26**	[15, 65, 122]
Classic wo est target	0.998	20	$50m \times 50m$	2	[10, 70, 128]
Spatial	0.9995	20	$50m \times 50m$	**25**	**[5, 65, 121]**
Spatial wo est target	0.9999	20	$50m \times 50m$	**10**	**[6, 65, 121]**
Temporal	0.9995	20	$50m \times 50m$	**7**	**[38, 96, 131]**
Temporal	0.9995	40	$100m \times 100m$	19	[21, 75, 128]
Temporal	0.9999	40	$100m \times 100m$	**0**	[53, 96, 140]
Temporal wo est target	0.9995	20	$50m \times 50m$	20	[24, 70, 128]

Table 5.2: Tracking results from the scenario 'Near Munkholmen and Approach'

Even though the non parametric IPDA was outperformed by the clutter map estimators in Chapter 4, the results are not as clear in this real scenario. The non parametric IPDA initiated six false tracks for a threshold of 0.998, and increasing the threshold lead to the real targets not being found. Spatial performed arguably the worst in terms of number of false tracks. The estimator lead to falsely initiated tracks across the whole map. Nonetheless, some of its poor estimation may be due to the lack of a nautical map implementation, as mentioned earlier in this section.

The Temporal estimator manages to filter all false preliminary tracks by setting the threshold to 0.9999. Yet the high threshold affects initiation time considerably. 53 scans is in most scenarios not an acceptable initiation time. At a lower threshold of 0.9995 it does considerably better in terms of initiation time, but also initiates several false tracks. The tracks initiated by the Temporal estimator are mostly close to the areas by the border where the nautical map has masked out land, as we see in Figure 5.14. Since clutter cells overlapping this border has a bias, they may contribute to higher existence estimation and cause the initiation of false tracks. There is in general more clutter near shore, so the implementation faults are not necessarily entirely to blame, but they may have a negative impact on the results.

The Classic estimator performs best out of the three estimators, and is the only one of the three to clearly outperform the non parametric method. With cells of size $V(c) = 50 \times 50m^2$ and an averaging length of $N = 20$, it is able to initiate only one false track. The initiation times are still nearly as good as the non parametric

method. The initiated tracks are shown in Figure 5.15. When comparing it to Figure 5.16 we see that it does not initiate the false tracks that 'Non parametric' does in some of the high clutter areas on the map.

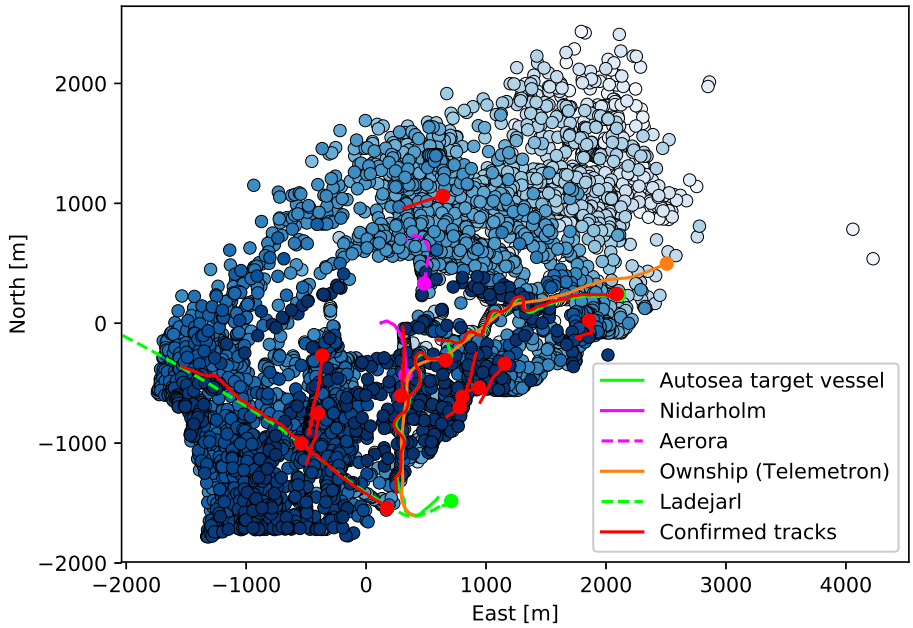

Figure 5.14: 'Near Munkholmen and Approach' scenario with confirmed tracks from Temporal estimator. $N = 20, V(c) = 50 \times 50m^2$.

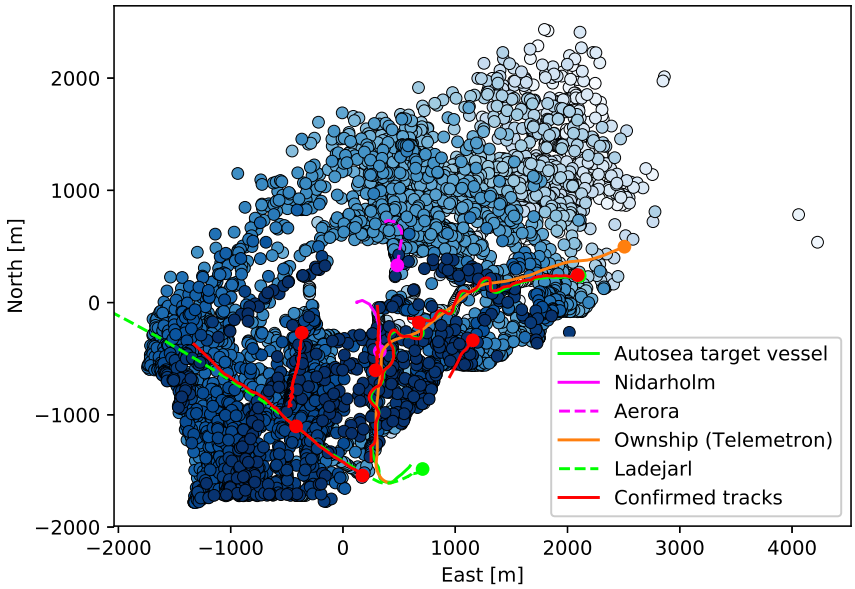

Figure 5.15: 'Near Munkholmen and Approach' scenario with confirmed tracks from Classic estimator. $N = 20, V(c) = 50 \times 50m^2$.

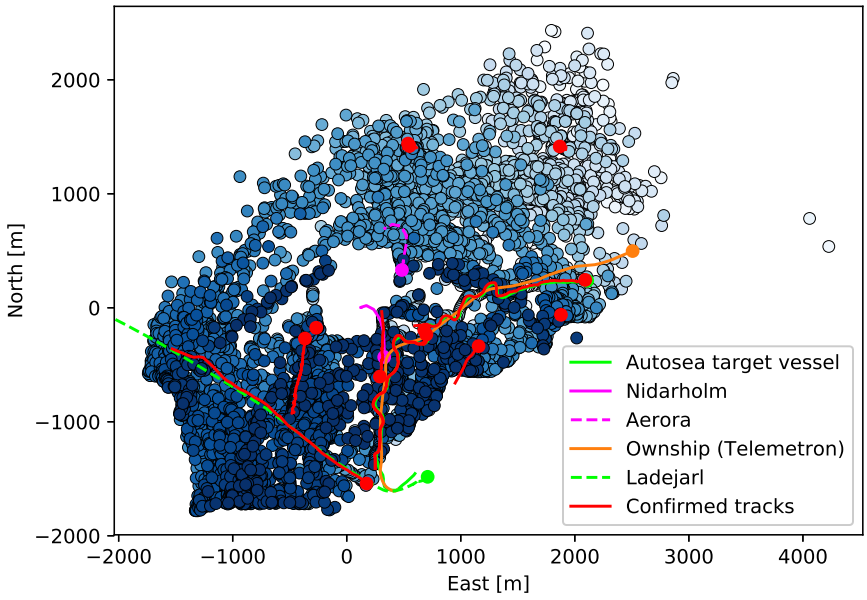

Figure 5.16: 'Near Munkholmen and Approach' scenario with confirmed tracks from a non parametric IPDA.

It is worth noticing that both Spatial and Temporal have to use very high initiation thresholds to filter out false tracks. This may be due to poor clutter estimation as the real clutter densities change when the ownship changes location, which was also mentioned for the 'Seatex Drone' scenario in Section 5.1.

Excluding estimated target measurements from clutter estimation was also tested and proved very beneficial in some tests. The Spatial estimator was able to reduce the number of false tracks from 25 to 10 and retained the same initiation times. Excluding the estimated target measurements in the clutter estimator lead to higher existence values for the true tracks, meaning the threshold could be increased without loosing true confirmed tracks. Figure 5.17 demonstrates how the estimated clutter densities in the area around a real target are lower when the exclusion of target measurements method is applied. However, as this may happen for false preliminary tracks as well, some false tracks can get higher existence values. This was the case for Classic and Temporal, which did not benefit from removing the estimated target measurements in terms of number false tracks. Temporal got lower track initiation times, but at the cost of almost the triple amount of false tracks.

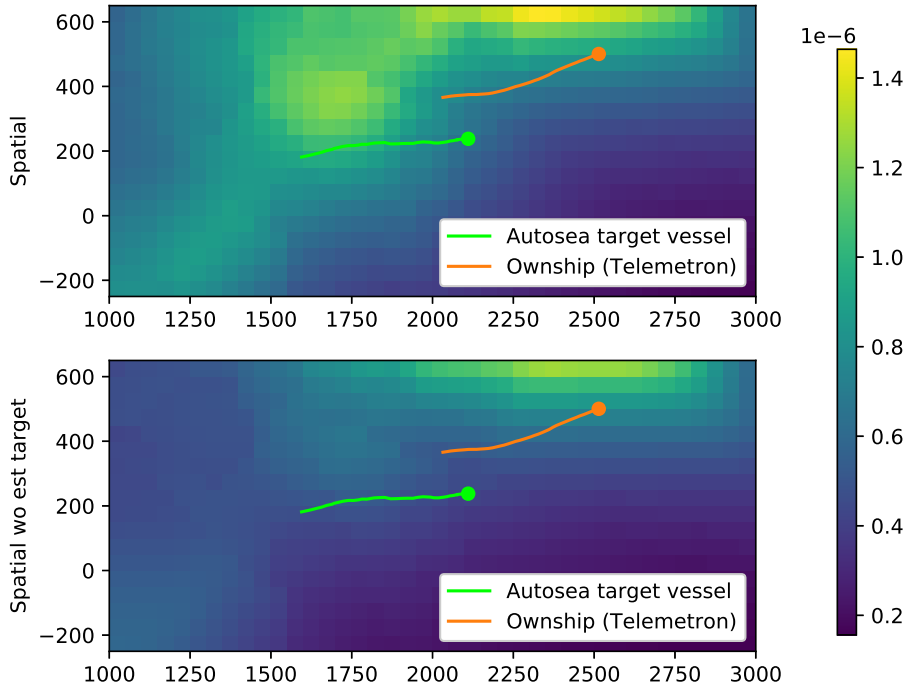

Figure 5.17: A comparison of a Spatial estimator with and without the estimated target measurements removal method (Section 3.6) in an area around a real target. The clutter maps are based on the first 40 scans of the 'Near Munkholmen and Approach' scenario. $N = 20, V(c) = 50 \times 50m^2$.

In general it was found that a smaller cell size ($50m \times 50m$) and a shorter averaging length (20) was beneficial for the results. As we see in Table 5.2, Classic was tested with the same threshold for $N = 40$ and a cell size of $100m \times 100m$, and both affected the number of false tracks negatively. That being said, smaller cell size also means more cells on the map and will thus affect runtime negatively.

Visual inspection of the entire clutter map did not contribute to much analysis. The clutter map for the Classic estimator with $N = 20$ can be seen in Figure 5.18. The cells with a higher clutter density forming an ellipse in the outer regions of the map, show the the outer rim of the circular radar range as the ownship has moved. These cells were not within radar range long enough to fill up the averaging window and therefore estimate a higher clutter density than the lowest bound of the estimator.

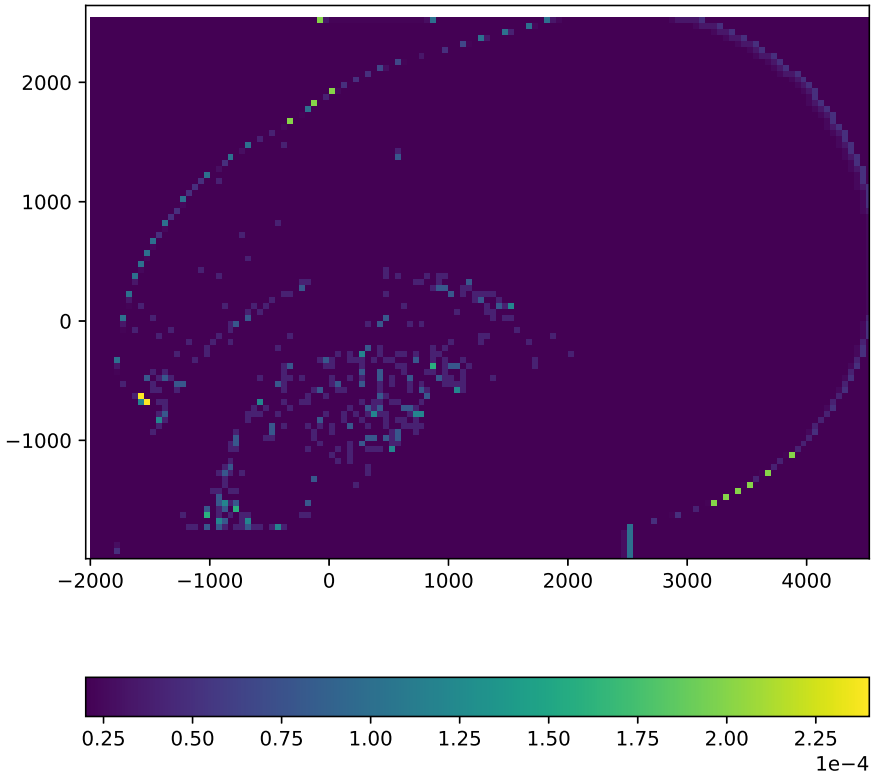

Figure 5.18: Clutter map using Classic estimator at the end of the scenario 'Near Munkholmen and Approach'. $N = 20, V(c) = 50 \times 50m^2$.

Chapter 6

Discussion

The efficiency and quality of a track initiation process is affected by choice of equipment and initiation techniques, and depends on the desired outcome and quality measures for the field of application. An effective COLAV system at sea relies on quick initiation of tracks in order to avoid dangerous situations when navigating through congested areas. On the other hand, setting initiation thresholds too low for faster initiation can lead to false tracks and erroneous trajectory behaviour of the ownship.

The requested maximum number of scans to initiate a track may differ according to the radar equipment and details about ownship location and the specific scenario. The ratio between radar scan rate and vicinity to new potential targets has a direct impact on the requested number of scans. When moving close to shore, the time to initiate targets might be lower than when further out on the ocean due to the proximity to vessels which may suddenly leave harbor. Areas near shore are also prone to producing more clutter, making the process of filtering clutter measurements from target measurements critical for a reliable initiation system. Parametric editions of tracking methods attempt to assist in clutter calculation by taking previous or additional information into account when evaluating the current situation. Using a clutter map is one way to estimate clutter in different regions of a scanned area, with advantages and disadvantages accompanied with the chosen clutter estimation method.

Musicki's three clutter estimators featured in [17] each have certain estimation biases connected to parameter choices and real clutter distribution. Theoretical biases of each estimator were presented in Section 2.4.4 and were tested in Section 4.1. The performance results of each estimator in simulations was presented in Section 4.3 where most results corresponded with the theory. The real results on the other hand, showed that parametric IPDA's does not necessarily provide superior results. One significant difference of the real data sets compared to the simulations is that the radar is on board a moving vessel, meaning the radar position is not stationary. One prominent effect from the moving radar in the real

scenarios was a non stationary clutter distribution over time. By comparing figures from the first and second half of both scenarios (Figure 5.6, Figure 5.7, Figure 5.11, Figure 5.12), we see that new high clutter areas have formed, while others have disappeared. In addition, new clutter cells are also spawned as the radar range reaches outside the current outer clutter cells in the 'Near Munkholmen and Approach' scenario (Section 5.2). Another effect influencing the results comes from the filtered land areas in the measurement set in 'Near Munkholmen and Approach'. Both of these effects are further discussed for each clutter map estimator below.

6.1 Classic clutter estimator

The Classic estimator is theoretically the least preferred estimator of the three due to its slow convergence time and bias for low mean clutter measurement values. This corresponds with some of the results from the simulations, yet not as much as one might suggest from only knowing the theoretical bias. It was only slightly less favorable to the Temporal estimator for some initialization lengths, but when tested without initialization time, Classic actually showed the best results of the three estimators (Figure 4.13, Figure 4.14). This may be because the other two estimators are more affected by the target measurements when they have no prior initialization period. Classic has a positive relative bias for few measurements in the averaging window meaning that the estimated value is lower than it should be, while the other two estimators are more prone to estimating higher densities in this situation. Higher estimated densities around real targets can cause lower existence values, meaning the other two estimators need lower thresholds to initiate the true tracks.

When tested for real measurement data, the Classic estimator outperforms the other estimators in both tested scenarios (Chapter 5). It shows similar results to the non parametric IPDA in the 'Seatex Drone' scenario in Section 5.1 and shows the best results of all tested IPDA setups in the 'Near Munkholmen and Approach' scenario in Section 5.2. As stated earlier, the non stationary radar in these scenarios may contribute to these results. The Classic estimator has a lower bound for the clutter density estimate when the averaging window is filled up, which may actually be beneficial when the clutter distribution is non stationary. With a clutter area of $V(c) = 50m \times 50m$ and averaging length $N = 20$, the lower bound for clutter density is $\lambda = 2 \times 10^{-5}$. Smaller cells and shorter averaging lengths leads to an increased lower bound. The lower bound may be slightly in favor of the Classic estimator for new high clutter areas as the IPDA will estimate lower existence values for false tracks in these areas. We also see from its relative bias in Figure 2.9 that it is otherwise biased to estimate lower clutter densities than real values. This implies that the estimator in general has a smaller range for most of its estimates, than the other two estimators.

6.2 Spatial clutter estimator

The spatial estimator does not have any bias in uniform clutter which is clearly shown in Figure 4.17 where it outperforms the other two estimators in both accuracy and convergence speed. Hence, in real scenarios where the clutter is known to be close to uniformly spread across the map, the Spatial estimator would be the most effective. Whether the assumption that clutter is uniformly spread at sea can be further discussed, and if so it would be somewhat unnecessary to have a number of clutter cells in the map. Nonetheless, the Spatial estimator outperformed the non parametric IPDA in uniform high clutter by a significant amount, shown in Figure 4.20 and Figure 4.21.

The simulated advantages of the Spatial estimator from Chapter 4 were not as easy to come across in the results from real data sets in Chapter 5. Spatial showed the poorest results in both tested scenarios in terms of number of false tracks. Still, track initiation time for the real targets was comparable to the best results in both tested scenarios. The moving radar and non stationary clutter distribution can explain why the estimator operated so poorly. Since the cell size does not constrain the lowest estimated clutter density for the estimator, as it does to some degree for the other two, it can obtain very low estimated values in areas of the map. If these areas suddenly return more clutter, the new density will use some time to propagate in the averaging window, meaning that the very low estimated value will be retained for some time. False preliminary tracks in such areas will get too high estimated existence values before the new density has propagated in the window, and may thus result in the initiation of false tracks. The lack of a nautical map in the clutter map implementation also contributes to a lower estimated clutter effect, as discussed in Section 5.2. The bias of clutter cells near the edge of a mapped out area can be compared to cells being close to cells close to a lower clutter density region, demonstrated in Figure 4.4.

The results from the Spatial estimator improved significantly when it was combined with the removal of estimated target measurements technique proposed in Section 3.6. The measurements from targets moving in low clutter areas can affect many surrounding Spatial clutter cells by raising their estimated density values, which then lowers the estimated existence of the target in the parametric IPDA algorithm. Hence, removing estimated target measurements made the IPDA able to initiate these tracks at a higher threshold and simultaneously filter other false tracks initiated at the previous threshold. Improving the target measurement removal technique further might therefore improve the estimator results even more.

The runtime of the Spatial estimator was by far the slowest of the three estimators. Including a KD-tree for the measurement set, as mentioned in Section 2.4.2, accelerated the simulations in Chapter 4, but the estimator was still noticeably slower than the other two in the real data simulations in Chapter 5.

6.3 Temporal clutter estimator

Even though the Temporal estimator converges towards no bias given long enough time and a big enough averaging window, the bias demonstrated in Section 4.1 can have a considerable impact on the performance of a parametric IPDA. The results from simulations did not show Temporal as favorable compared to the Classic estimator except for slightly better track initiation time in some tests. Furthermore, it was only better than Spatial for high P_{FT} values in the scenario with 80 scans of map initialization (Figure 4.11). Without initialization it did not show much improvement to the non parametric IPDA, which could be explained by the estimator's high variance and bias during initial scans of a simulation. In Figure 4.17 Temporal is the only estimator to have several cells with significantly higher clutter estimation than the real clutter density after 40 scans of initialization.

The high bias and variance of the Temporal estimator during initial estimation of a clutter area, can explain why the estimator shows so poor results compared to the non parametric IPDA in the real scenarios in Chapter 5. The changes in clutter densities as the ownship moves around on the map, creates a situation where Temporal does not have sufficient time to converge towards an unbiased estimate before a new clutter density is present in the cell. This effect would especially be problematic if a target moves into a previously high clutter area which now has a lower clutter density, as old inter-arrival times τ_k remains in the averaging window and are not propagated out due to little or no following clutter. Spatial and Classic are more adaptable in this sense as they add and remove data in the averaging window in each step.

6.4 Overall discussion

All simulated test scenarios in this thesis clearly represent a cleaner and more uniform clutter distribution than the real world scenarios presented in Chapter 5. Comparing Figure 3.2 with Figure 5.3 we see that the squared clutter density regions in the simulated test scenario are not completely comparable to the oddly shaped and distributed clutter in the 'Seatex Drone' scenario. The implementation of the clutter map in this thesis is based on square regions making it more fit to handle the clutter from the test scenario. Even from one real scenario to another, false measurements are unequally distributed and appear more dynamic over time and it is hard to predict any pattern beforehand. Consequently, we get significantly better results from the simulated testing than from the recorded data from Trondheimsfjorden. When the estimated clutter densities in the clutter map are too far from the truth, they misguide the tracking system more than they assist it, leading to poorer results from the parametric IPDA than the non parametric IPDA.

As previously stated, the relative variance of all three estimators and the bias of the Temporal estimator decrease over time if the averaging window is not full. The downside of having too long averaging window is that all estimates become nearly static. If any real changes occur in the clutter density some time after startup,

they will propagate slower in the averaging window and thus affect the density estimates slower. In other words, there is some trade off between the amount of bias and variance in the estimated map, and responsiveness to real clutter density changes. The final choice depends on how static or shifting one expects the real clutter densities to be.

Since many ships and vessels move close to shore only as they enter or leave a harbor, one could argue that a stationary radar could be set up in such locations to provide measurement data for passing ships. The stationary radar would hopefully have a more stationary and slowly changing clutter distribution than a radar on board a ship, making clutter maps more fit for clutter estimation. In such a radar set up, the Temporal estimator may be preferable since it would have time to converge towards its full potential of zero bias. Otherwise, for radar on board a moving vessel, the Classic estimator appears to be the most suitable clutter map estimator based on the real measurement sets in this thesis.

The parametric IPDA's could beneficially be tested with a nautical map to see their true potential. As a lot of false tracks were initiated near the mapped out coast line, it would be interesting to see if such tracks are filtered when a nautical map implementation is in place.

Although the focus in this thesis has been on a parametric IPDA with clutter map, the non parametric IPDA performs very well on its own in the real scenarios. It also shows promising results compared to other track initiation techniques in [19] which is added in Appendix B.

Chapter 7

Conclusion

Track initiation is an important part of a tracking system for autonomous vehicles. There are many ways of implementing the initiation process, including M/N [2], IPDA [15] and SPRT [9]. Clutter among radar measurements may lead a tracking system to initiate false tracks, and in many real world scenarios, certain areas of the map are more likely to return clutter than others, e.g. near-shore areas [20]. Track initiation methods can therefore beneficially be combined with a clutter map for improved initiation time and accuracy.

The three clutter estimators for clutter maps featured in [17] improve the results of an IPDA in simulations. The Spatial estimator is preferred for short initialization time and in general for relatively uniform clutter. Though, its heavy bias for cells near discontinuities in real clutter density make it unfavorable in many situations. The real data sets tested in this thesis had nonuniform clutter, and Spatial showed the poorest results of the three estimators. Its shape expansion is also more difficult to implement correctly for real scenarios where the radar has masked out measurements from land. The runtime was also significantly slower for the Spatial estimator compared to the other two, making it less desirable.

For a non stationary radar and thus possibly a non stationary clutter distribution, the Classic estimator is to prefer and the only one to perform just as good or better than the non parametric IPDA. In spite of its relative bias, Classic shows almost the same results as Temporal in simulations, and outperforms all other IPDA editions in the 'Near Munkholmen and Approach' scenario (Section 5.2). Still, the Temporal estimator can be preferable for a stationary radar setup when the clutter distribution is more stable.

The suggested method for excluding estimated target measurements showed to be beneficial for track initiation time and somewhat for filtering false tracks.

7.1 Future work

Real world testing

Clutter maps have so far had limited documented real world testing. Even though some scenarios were tested in this thesis, more experiments are needed to draw any full conclusions. This entails testing in various locations, noise conditions, maritime traffic conditions and so on.

Other tracking methods

This thesis has focused on using a clutter map in combination with the IPDA to improve single-target tracking, but clutter maps can also be paired with a range of other tracking methods. It would be especially interesting to see a comparison of parametric methods of single-target tracking against multi-target tracking methods both in simulations and real world testing. The multi-target edition JIPDA [14] of IPDA would be an interesting place to start.

Excluding target measurements

A difficulty with the implementation of clutter maps was how to avoid target measurements from clutter estimation. One suggestion was presented in Section 3.6, yet there might exist other solutions and approaches to this challenge. Additional testing of techniques for estimating and removing target measurements would be useful for future implementations of clutter maps. A suggestion for improving the presented technique in this thesis, is to set an acceptance threshold for the conditional probability β_k^i of the selected target measurement. Furthermore, the threshold could depend on the existence value $P\{\boldsymbol{x}_k|\boldsymbol{Z}^k\}$ of the track.

Map layout and reference frame

The clutter map implementation in this thesis used an earth-fixed reference frame, though it is also possible to use a body-fixed reference frame. This may introduce different implementational challenges than the earth-fixed solution, but could provide improvements to other issues in the field of target tracking. There was also no attempt to relocate the clutter map position during the simulated tests, which would be a more equivalent approach to the real scenarios presented in Chapter 5.

Other clutter dispersal

One can also attempt to simulate other dispersal of clutter density with more randomly created shapes and sizes of high clutter density regions. This could perhaps be a closer approximation of real clutter than what was presented in the simulation scenario for this thesis. It could also be more suitable for testing a clutter map with dynamic decision of clutter cell sizes and positions.

Appendices

Appendix A

Additional SOC results

Figure A.1, Figure A.2 and Figure A.3 demonstrate further how initialization time impacts each clutter estimator, which was discussed in Section 4.3. Spatial does not improve noticeably after 20 scans, but Classic and Temporal do.

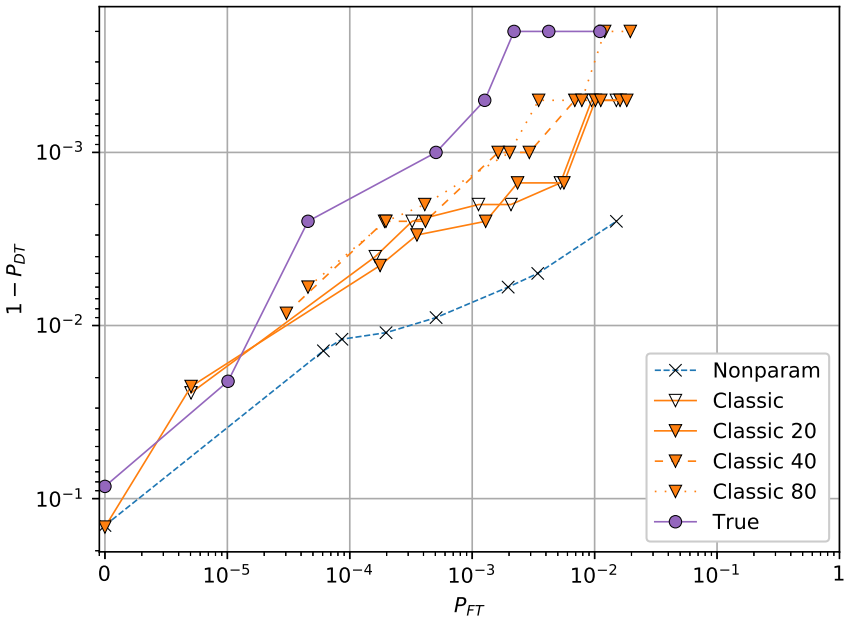

Figure A.1: A comparison of SOC results for all tested initialization lengths for the Classic estimator running the test scenario (Section 3.4).

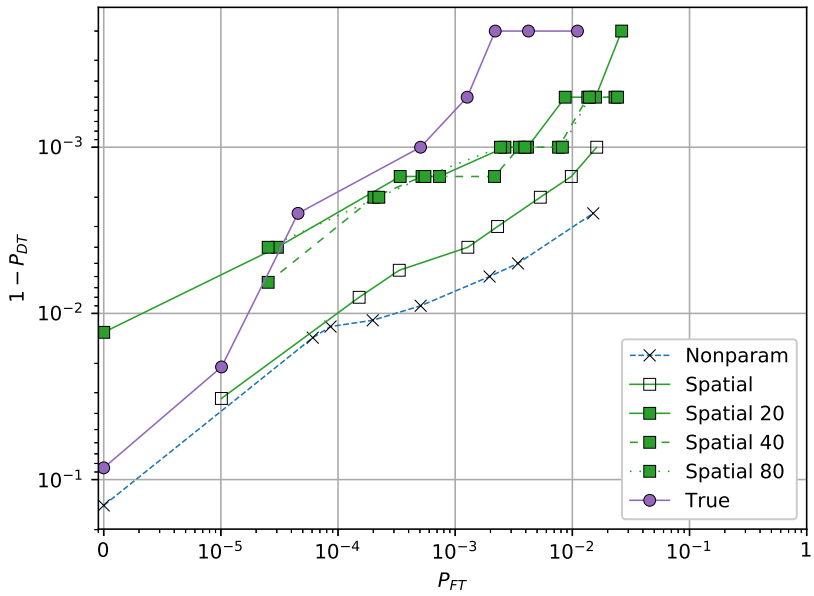

Figure A.2: A comparison of SOC results for all tested initialization lengths for the Spatial estimator running the test scenario (Section 3.4).

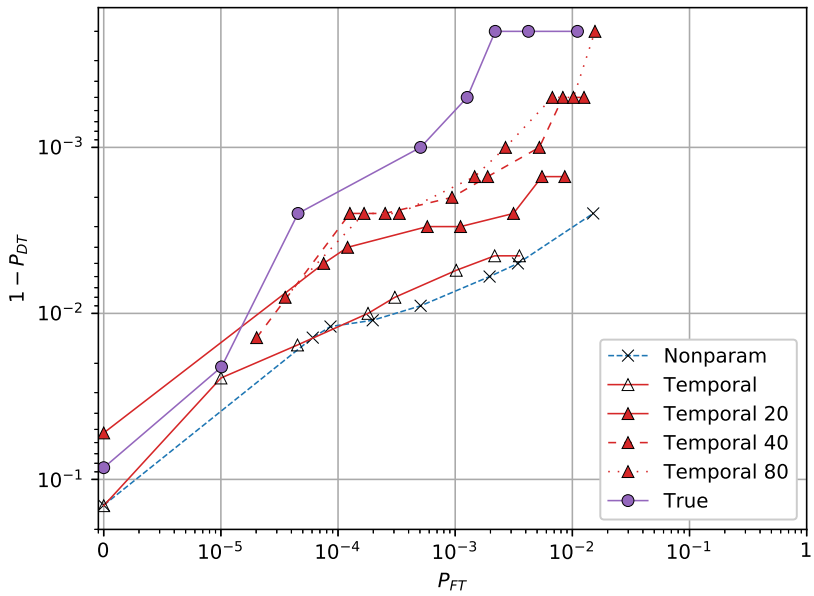

Figure A.3: A comparison of SOC results for all tested initialization lengths for the Temporal estimator running the test scenario (Section 3.4).

Appendix B

Track Initiation for Maritime Radar Tracking with and without Prior Information

Track Initiation for Maritime Radar Tracking with and without Prior Information

Erik F. Wilthil, Edmund Brekke and Oskar B. Asplin
Department of Engineering Cybernetics
NTNU
Trondheim, Norway
Email: erik.wilthil@ntnu.no, edmund.brekke@ntnu.no, oskarba@stud.ntnu.no

***Abstract*—Reliable track initiation is an important component of a tracking system, especially when it is used as part of a more general collision avoidance (COLAV) system. Some tracking methods (e.g., IPDA) come with an in-built track initiation capability, while other methods (e.g., JPDA) lack this capability, which in such cases typically is taken care of by heuristic rules such as the M/N logic. Although Reid's multiple hypothesis tracker (MHT) is capable of track initiation, many implementations do not include track initiation in the MHT framework due to the increased complexity. While MHT is fundamentally Bayesian, the non-Bayesian sequential probability ratio test (SPRT) of Van Keuk is often used for track initiation.**

In this paper we derive a Bayesian SPRT for track initiation based on Reid's MHT. The approach is compared with the classical SPRT, both from a theoretical perspective and using simulations. Furthermore, the paper provides a comparison between the two SPRT versions, the IPDA and the M/N logic in terms of system operating characteristic (SOC) curves and track initiation time. The initiation methods are also tested on real radar data recorded during full-scale maritime COLAV experiments.

I. Introduction

Target tracking is a key ingredient in collision avoidance (COLAV) for autonomous surface vehicles (ASVs). Commercial vessels may use the automatic identification system (AIS) system, but exteroceptive sensors such as radars are still needed to track objects without AIS and in case of transmission failure. Track management is a vital part of the target tracking system, and track initiation in particular is an important piece. The main priority is to initiate tracks such that collision can be avoided. However, if a false track is present, it may cause unnecessary maneuvers which confuse other ships, or in worst case cause a collision with a real target. On the other hand, if the target is detected too late, collision may be imminent. Thus, the false alarm rate and track initiation time must be weighted against each other.

There are two main approaches to track initiation. Soft methods, such as the integrated probabilistic data association (IPDA) may in principle maintain a track based on a measure such as existence probability without making a decision on whether a target is present or not. Another example of such soft methods are trackers designed from the related concept of perceivability [7]. Hard methods such as the M/N logic-based method [1] and the sequential probability ratio test (SPRT) [6] will declare a target present or not. When hard decisions have to be made with basis in the tracking results (e.g., to start a COLAV maneuver), the soft methods are thresholded to declare tracks either present or absent. Choosing this threshold is neither straightforward nor inconsequential. To the authors' knowledge, there are few documented real-world results for the IPDA, but it was successfully applied to tracking divers in [3]. Methods based on finite set statistics (FISST) [10] are also considered soft, as they also evaluate the probability that a target is present.

Experimental validation of target tracking systems in maritime environments have been reported previously in e.g. [16], [5], [4]. In [16], the M/N method was used, which resulted in several false tracks from returns close to land. In [5], the SPRT was used to track static objects with an automotive radar, from a high-speed vehicle. The SPRT has also been successfully applied in the P-MHT framework [14], and in the perceivability framework [8].

The remainder of the paper is organized as follows: In section II, we summarize the existing methods for track initiation. In section III, we introduce a Bayesian SPRT based on the equations of Reid's multiple hypothesis tracker (MHT) [12]. In section IV, we present simulation results, and experimental results are presented in section V.

II. Summary of existing methods

A. Assumptions

Assumption 1: Clutter is modelled as a Poisson process, such that the number of clutter measurements m^c in an area V for a single scan has the distribution

$$P_c(m^c) = \frac{(\lambda(z)V)^{m^c}}{m^c!} e^{-\lambda(z)V} \tag{1}$$

where $\lambda(z)$ is the clutter density in V, in general spatially varying. The concept of "new targets" can be further separated as newborn targets, and undiscovered targets [15]. However, they are highly related [2], and the distinction will not be discussed in this work.

Assumption 2: The number of new targets m^b in an area V has the distribution

$$P_b(m^b) = \frac{(\beta(z)V)^{m^b}}{m^b!} e^{-\beta(z)V} \tag{2}$$

where the parameter $\beta(z)$ is the new target density in V. As the clutter density, this may also be spatially varying.

Assumption 3: The target motion model is assumed to be linear-Gaussian, i.e. if the target exist, it moves according to

$$p(x_{k+1}|x_k) = \mathcal{N}(x_{k+1}; F_k x_k, Q_k) \tag{3}$$

where x_k is the state of the system at time k, and F_k and Q_k are the state transition and process noise covariance, respectively.

Assumption 4: If the target exist, it is observed independently over scans with constant detection probability P_D, according to

$$p(z_k|x_k) = \mathcal{N}(z_k; Hx_k, R_k) \tag{4}$$

The prior and posterior estimates of the target state at time k are denoted $\hat{x}_{k|k-1}$ and $\hat{x}_{k|k}$, respectively. The data, which is the sequence of scans from time k_0 to k, are denoted Z^k:

$$Z^k = (Z_{k_0}, Z_{k_0+1}, \ldots, Z_k), \qquad Z_i = \left\{z_i^j\right\}_{j=1}^{m_i} \tag{5}$$

where the scan Z_i is the set of measurements at time i.

B. Single-target tracking

Given an initial target estimate, the target can be tracked by a probabilistic data association filter (PDAF). First, the measurements in the surveillance region are gated, according to

$$\nu_k^{jT} S_k^{-1} \nu_k^j < \gamma_G \tag{6}$$

where $\nu_k^j = H\hat{x}_{k|k-1} - z_k^j$ is the innovation of measurement j, S_k is the innovation covariance and the threshold γ_G can be found from the χ^2-distribution by specifying the gate probability P_G [1, p. 96]. Gating ensures only measurements close to the existing track will be considered in estimating $\hat{x}_{k|k}$. Not only does this save computational resources, it also lets us track several targets in parallel, as long as they remain well separated.

After gating measurements by using the prior estimate, the posterior estimate $\hat{x}_{k|k}$ is calculated based on a weighted sum of the measurement innovations. These weights, called association probabilities, are given by

$$P(\theta_k^j|Z^k) \propto \begin{cases} p(z_k^j|\theta_k^j, Z^{k-1}) & j \in (1, \ldots m_k) \\ \frac{2(1-P_D P_G)}{P_D \gamma_G} \lambda_j V_k & j = 0 \end{cases} \tag{7}$$

where θ_k^j is the event that the measurement z_k^j is target-originated, and $p(z_k^j|\theta_k^j, Z^{k-1})$ is the likelihood of measurement z_k^j being target-originated, given the prior estimate $\hat{x}_{k|k-1}$. The event θ_k^0 represents a misdetection, i.e. that all the measurements in the gate are clutter. If the clutter density $\lambda_j = \lambda(z_k^j)$ is not available, the nonparametric version of the PDAF can be obtained by setting $\lambda_j = m_k/V_k$, i.e. the number of measurements in the gate [1].

C. M/N-based track initiation

The simplest form of track initiation is arguably M/N-based track initiation [1]. The initiation test starts with any measurement not associated with any existing tracks, which are set up as tentative tracks. In the following scan, the measurements within a certain range of the tentative track are used to initiate a preliminary track. The range is determined by the assumed maximum velocity of the targets and measurement covariance. The preliminary track is then tracked using a PDAF for maximally the next N scans. If the PDAF gates measurements in at least M of these scans, the target is confirmed. If not, the track is terminated.

D. Integrated PDA

The IPDA was introduced in [11]. It extends the PDAF by a discrete existence variable E_k, which can attain two values, O_k: The target exists at time k, and N_k: The target does not exist at time k. Some versions of the IPDA let E_k attain a third value, I_k: target exists, but cannot be detected at time k. This is more suited for track management than track initiation [11], and will not be used here. This means O_k and N_k are mutually exclusive and exhaustive, thus $P(N_k) = 1 - P(O_k)$. The target existence is assumed to follow a Markov chain according to

$$P(O_k) = \pi_{11} P(O_{k-1}) \tag{8}$$

where π_{11} is often called the *survival probability* of the target. After gating measurements, the probability of target existence is updated according to

$$P(O_{k|k}) = \frac{(1-\zeta_k)P(O_{k|k-1})}{1 - \zeta_k P(O_{k|k-1})} \tag{9}$$

where $P(O_{k|k-1})$ is short-hand for $P(E_k = O_k|Z^{k-1})$, and

$$\zeta_k = \begin{cases} P_D P_G & m_k = 0 \\ P_D P_G \left(1 - \sum_{j=1}^{m_k} \frac{1}{\lambda_j} p(z_k^j|Z^{k-1})\right) & m_k > 0 \end{cases} \tag{10}$$

The nonparametric IPDA is obtained by setting

$$\lambda_j = \frac{m_k - P_D P_G P(O_{k|k-1})}{V_k} \tag{11}$$

which means the expected number of clutter measurements are the total number of measurements in the gate, compensated for the existence probability of the target. In order to compare the IPDA with the other methods, the tracks are confirmed if the existence probability exceeds a threshold P_C, and terminated when it is below P_T.

E. Sequential tests

A sequential probability ratio test (SPRT) is a statistical test where the sample size is not fixed in advance. Instead, data are accumulated, and a likelihood ratio (LR) of the form

$$\mathrm{LR}(k) = \frac{p(Z^k|\mathcal{H}_1)}{p(Z^k|\mathcal{H}_0)} \tag{12}$$

is evaluated [6]. For track initiation, the hypotheses are defined as

- $\mathcal{H}_0$: The data Z^k consists only of clutter
- $\mathcal{H}_1$: The data Z^k contains measurements from a target.

At a given time step, there are three possible outcomes of the SPRT,

- Accept $\mathcal{H}_0$ if $\mathrm{LR}(k) < \mathrm{LR}_{\min}$
- Accept $\mathcal{H}_1$ if $\mathrm{LR}(k) > \mathrm{LR}_{\max}$
- Make no decision if $\mathrm{LR}_{\min} \leq \mathrm{LR}(k) \leq \mathrm{LR}_{\max}$.

In the last case, the test waits for additional measurements at scan $k+1$, and calculates the LR based on the new measurements. The limits $\mathrm{LR}_{\min}$ and $\mathrm{LR}_{\max}$ are related to the error probabilities of the test by

$$\mathrm{LR}_{\min} \geq \frac{1-P_1}{1-P_0} \qquad \frac{1}{\mathrm{LR}_{\max}} \geq \frac{P_0}{P_1} \tag{13}$$

where P_0 and P_1 are the probabilities of rejecting $\mathcal{H}_0$ when it is true, and accepting $\mathcal{H}_1$ when it is true, respectively [13].

In [6], an SPRT is derived for track initiation using (12). The main idea is that, under $\mathcal{H}_1$, there is a sequence of measurements in the data that are target-originated. The likelihood of each of these sequences, called data interpretations in [6], can be summed to obtain the total likelihood of $\mathcal{H}_1$:

$$p(Z^k|\mathcal{H}_1) = \sum_{\Omega_k^j \in \mathcal{H}_1} p(Z^k|\Omega_k^j) P(\Omega_k^j) \tag{14}$$

where Ω_k^j is a measurement association hypothesis, i.e. it defines what measurements in Z^k that are target-originated, and what measurements are clutter. It is similar to the hypothesis structure of Reid [12]. The likelihood ratio then takes the form

$$\mathrm{LR}(k) = \frac{\sum_{\Omega_k^j \in \mathcal{H}_1} p(Z^k|\Omega_k^j) P(\Omega_k^j)}{p(Z^k|\mathcal{H}_0)} \tag{15}$$

which is initialized by selecting a single measurement, which is to be tested whether it is target-originated or not. The likelihood of the first scan is calculated by "assuming that the target exist anywhere in V" [6], and leads to the initial LR

$$\mathrm{LR}(k_0) = \frac{P_D}{\lambda V_{k_0}} \tag{16}$$

where V_{k_0} is the area of the initial search area. A small search area means that the initial likelihood ratio will be large. The reason for this is that if it is assumed that a target exists "anywhere in V_{k_0}", and then a measurement is found in that very region, the likelihood that it is target-originated is very large. Although practical aspects are discussed in [6], selecting the gate area for the SPRT is not in one of them. Our approach for this gate will be presented in the following section.

F. Choosing the data set area

As mentioned, the only way of choosing the gate of an SPRT in [6] is to select the region containing all the measurements in the surveillance area. For the same reasons as we gate established tracks, we want to restrict the data of the SPRT to a subset of the measurements from the surveillance area, close to the initial measurement $z_{k_0}^j$ of the SPRT. Initial measurements are chosen when they are not associated with any previous gates or confirmed tracks. If $z_{k_0}^j$ is target-originated, the true target position z_k^T is found in a region around $z_{k_0}^j$ with probability P_G according to

$$(z_{k_0}^j - z_k^T)^T R^{-1} (z_{k_0}^j - z_k^T) < \gamma_G \tag{17}$$

where γ_G is found in the same way as for the track gate, namely by specifying the gate probability P_G. This region is in general ellipsoidal, and in the following, we will approximate it by a circle, since in the subsequent scans, this radius will increase according to the maximum velocity and sampling time. A circular gate for selecting the measurements that will be included in the test is simpler than maintaining the direction and length of the ellipsoid major and minor axes. The maximum range r_{k_0} from the measurement to the true target position is given by

$$r_{k_0}^2 \mathrm{eig}(R)_{\max}^{-1} = \gamma_G \Rightarrow r_{k_0} = \sqrt{\gamma_G \mathrm{eig}(R)_{\max}} \tag{18}$$

where $\mathrm{eig}(R)_{\max}$ is the maximum eigenvalue of the measurement covariance matrix R. The corresponding initial area V_{k_0} is given by $\pi r_{k_0}^2$. If R is very ellipsoidal, this approximation may result in very large search regions.

In the following scans, the gate will grow in radius according to

$$r_k = 2r_{k_0} + (t_k - t_{k_0}) v_{\max} \tag{19}$$

where $v_{\max}$ is the presumed maximum velocity of the targets. This gating strategy is illustrated in Fig. 1. The corresponding gate area is given by $V_k = \pi r_k^2$.

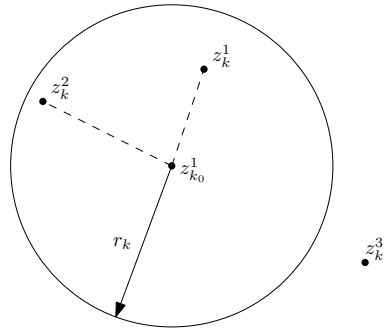

Fig. 1. The SPRT gates measurements according to (19). The measurements z_k^1 and z_k^2 are within range of $z_{k_0}^1$, and will be used in the test started by $z_{k_0}^1$. A new test will be started at z_k^3.

This gating strategy also allows us to develop a nonparametric version of the SPRT. Under $\mathcal{H}_0$, all the measurements in V_k at time k are clutter, and under $\mathcal{H}_1$, all but one of the measurements in V_k are clutter, unless the hypothesis is that the target is not detected. The resulting expression is

$$\lambda = \frac{m_k - \delta_k}{V_k} \tag{20}$$

where $\delta_k = 1$ if Ω_k^j declares that a measurement is target-originated, and $\delta_k = 0$ if the target is not present or undetected.

III. A Bayesian sequential test

The main result in Reid's seminal paper [12] is the recursive calculation of the data association hypothesis probability given the data, denoted by $P(\Omega_k^j|Z^k)$. In general, the data association hypothesis Ω_k^j contains information on what measurements originate from existing targets, new targets and clutter. It is given by

$$P(\Omega_k^j|Z^k) = \frac{1}{c_k} p(Z_k|\Omega_{k-1}^i, \omega_k^j, Z^{k-1}) \cdot P(\omega_k^j|\Omega_{k-1}^i, Z^{k-1}) P(\Omega_{k-1}^i|Z^{k-1}) \quad (21)$$

where Ω_{k-1}^i is the parent hypothesis from the previous scan, and the association hypothesis ω_k^j associates all the measurements in Z_k with either targets from Ω_{k-1}^i, new targets or clutter. The constant $c_k = p(Z_k|Z^{k-1})$ is independent of the hypotheses. Given Assumptions 1 to 4, the recursion is shown in [12] to be

$$P(\Omega_k^j|Z^k) = \frac{m_k^c! m_k^n!}{c_k m_k!} P_D^{m_k^d} (1-P_D)^{m_k^t - m_k^d} P_c(m_k^c) P_b(m_k^n) \cdot \prod_{l=1}^{m_k^d} p(z_k^l|\Omega_k^j, Z^{k-1}) V_k^{-(m_k^c + m_k^n)} P(\Omega_{k-1}^i|Z^{k-1}) \quad (22)$$

where m_k^c and m_k^n is the number of clutter measurements and new targets, m_k^t is the number of targets given by Ω_{k-1}^i, of which m_k^d are detected. This expression is suited to handle an arbitrary number of existing and new targets. The computational burden can be significant, but (22) can be simplified in the context of single-target track initiation, according to the hypotheses $\mathcal{H}_0$ and $\mathcal{H}_1$.

According to $\mathcal{H}_0$, all of the measurements are clutter. This leads to a single data association hypothesis given by

$$p(\mathcal{H}_0|Z^k) = P_c(m_k) P_b(0) V_k^{-m_k} p(\mathcal{H}_0|Z^{k-1}) \quad (23)$$

$$= \frac{1}{c_k^*} \prod_{i=1}^{m_k} \lambda_i \cdot p(\mathcal{H}_0|Z^{k-1}) \quad (24)$$

where the product of clutter densities of the measurements replaces λ^{m_k} due to nonstationary clutter. The normalization constant c_k^* contains all the factors that will cancel in the calculation of the LR, and includes the constant c_k of (22), along with other constants of (23). The recursion is initialized by

$$p(\mathcal{H}_0|Z_{k_0}) = \frac{1}{c_{k_0}} \frac{\prod_{i=1}^{m_{k_0}} \lambda_i}{m_{k_0}!} e^{-(\lambda+\beta)V_{k_0}} \quad (25)$$

For $\mathcal{H}_1$, as in the previous section, there are several possible data association hypotheses that may support the existence of a track. Summing over all of these gives

$$p(\mathcal{H}_1|Z^k) = \sum_{\Omega_k^j \in \mathcal{H}_1} P(\Omega_k^j|Z^k) \quad (26)$$

where each hypothesis may be seen as leaf nodes of a tree structure. In order to calculate the probability $P(\Omega_k^j|Z^k)$, it is useful to separate the cases $k = k_0$ and $k > k_0$. For the first scan under $\mathcal{H}_1$, a new target is present in the data Z_{k_0}. This gives

$$P(\Omega_{k_0}^j|Z_{k_0}) = \frac{1}{c_{k_0}} \frac{1}{m_{k_0}} P_c(m_{k_0}-1) P_b(1) V_k^{-m_{k_0}} \quad (27)$$

$$= \frac{1}{c_{k_0}} \frac{\prod_{i=1}^{m_k-1} \lambda_i}{m_{k_0}!} \beta e^{-(\lambda+\beta)V_{k_0}} p(\mathcal{H}_1) \quad (28)$$

where it is assumed that the measurement $z_{k_0}^{m_{k_0}}$ is the target-originated measurement. For the second scan and onwards, we assume that there is a single existing target, and no other new targets is entering the surveillance region. This leads to

$$P(\Omega_k^j|Z^k) = \frac{1}{c_k} \frac{(m_k - \delta_k)!}{m_k!} P_D^{\delta_k} (1-P_D)^{1-\delta_k} P_c(m_k - \delta_k) \cdot P_b(0) p(z_k^l|\Omega_k^j) V^{m_k - \delta_k} P(\Omega_{k-1}^i|Z^{k-1})$$

$$= \begin{cases} \frac{1}{c_k^*} P_D \prod_{i=1}^{m_k-1} \lambda_i p(z_k^l|\Omega_k^j) P(\Omega_{k-1}^i|Z^{k-1}) & \delta_k = 1 \\ \frac{1}{c_k^*} (1-P_D) \prod_{i=1}^{m_k} \lambda_i P(\Omega_{k-1}^i|Z^{k-1}) & \delta_k = 0 \end{cases} \quad (29)$$

where the first case means that measurement l is target-originated, and the second case means that the target is not detected at time k. Comparing (24) and (29), we see that all the constants concerning the gate area V_k are cancelled in the evaluation of (12), as well as the term c_k^*. The LR at the initial scan k_0 is found from (25) and (28), and is given by

$$\mathrm{LR}(k_0) = \frac{\beta}{\lambda} \quad (30)$$

where the new target density parameter β replaces the ratio P_D/V_{k_0} in the initial likelihood ratio of [6], given by (16).

After the test confirms that a target is present in the data, the most likely hypothesis is extracted as the track. This is similar to how a decision is made in a typical MHT, where the maximum a posterior (MAP)-hypothesis often is chosen when a track is needed. In addition to extracting the MAP estimate, the SPRT additionally requires the LR test to be successful.

IV. Simulation results

A. Test setup

The test scenario is set up to mimic challenges in maritime tracking scenarios. The two main challenges with respect to varying clutter densities comes from land and shallows. The scenario in Fig. 2 is set up to reflect both these challenges. Two targets are generated in the surveillance area. One is representing a target leaving berth, starting in the lower high-clutter region and moving north. The other target represents a passing ship, and moves from west to east along the high-clutter square in the surveillance area. The initial position, velocity and course is generated uniformly in certain intervals according to Table I. After the initial state has been determined, the targets move according to a nearly constant velocity (NCV) model [9], given by Assumption 3 where

$$F_k = \begin{bmatrix} F_1 & 0_{22} \\ 0_{22} & F_1 \end{bmatrix} \qquad F_1 = \begin{bmatrix} 1 & T \\ 0 & 1 \end{bmatrix} \quad (31)$$

$$Q_k = \begin{bmatrix} Q_1 & 0_{22} \\ 0_{22} & Q_1 \end{bmatrix} \qquad Q_1 = q \begin{bmatrix} T^4/4 & T^3/2 \\ T^3/2 & T^2 \end{bmatrix} \quad (32)$$

TABLE I
TARGET PARAMETERS FOR THE TEST SCENARIO

Parameter	Lower target	Upper target
N_0	$\mathcal{U}(-500\,\mathrm{m}, -300\,\mathrm{m})$	$\mathcal{U}(350\,\mathrm{m}, 450\,\mathrm{m})$
E_0	$\mathcal{U}(-500\,\mathrm{m}, 500\,\mathrm{m})$	$\mathcal{U}(-400\,\mathrm{m}, 0\,\mathrm{m})$
V_0	$\mathcal{U}(5\,\mathrm{m\,s^{-1}}, 10\,\mathrm{m\,s^{-1}})$	$\mathcal{U}(5\,\mathrm{m\,s^{-1}}, 12\,\mathrm{m\,s^{-1}})$
χ_0	$\mathcal{U}(-30°, 30°)$	$\mathcal{U}(70°, 110°)$

TABLE II
TRACKING SYSTEM PARAMETERS

Sampling time	T	$2.5\,\mathrm{s}$
Process noise covariance parameter	q	$(0.05\,\mathrm{m\,s^{-2}})^2$
Measurement noise covariance	R	$(6\,\mathrm{m})^2 I_2$
Probability of detection	P_D	0.9
low clutter density	λ_{low}	$10^{-5}\,\mathrm{m}^{-2}$
high clutter density	λ_{high}	$5 \times 10^{-5}\,\mathrm{m}^{-2}$

where $T = t_k - t_{k-1}$ is the time between two scans, and q is the target maneuvering parameter. 0_{22} denotes the 2×2 zero matrix. Test scenario parameters are found in Table II, which are common for both of the targets.

To avoid initiating multiple tracks on the same target, all confirmed targets will be tracked by a regular PDAF. The measurements that are gated by these confirmed tracks are not used in the track initiation methods, and confirmed tracks are terminated if it gates no measurements for the next 5 scans. The life and death of these confirmed tracks have no impact on the track initiation, except for the measurements they are confiscating. In settings where new targets appear close to each other, another tracking method may be applied. All the tests are done on track confirmation, such that the future trajectory of these tracks do not impact the results.

In the following scenarios, we analyse the probability of target detection, P_{DT}, and the probability of a false track, P_{FT}. The probability of target detection is calculated based on how many of the targets described above are detected in total, i.e.

$$P_{DT} = \frac{\text{number of targets detected}}{\text{number of total targets}} \tag{33}$$

over the course of 1000 tests for each target. A target is considered detected if at least one of the measurements used to confirm the track is target-originated. The probability of a false track is calculated based on how often the track initiation method confirms a track that is started on a clutter-originated measurement, and is defined as

$$P_{FT} = \frac{\text{number of false confirmed tracks}}{\text{number of confirmed and terminated tracks}} \tag{34}$$

where the test is conducted without the targets present, such that all confirmed tracks must be false. This is to avoid evaluating ambiguities, for example when tests are started on clutter measurements, but associate target measurements, and vice versa.

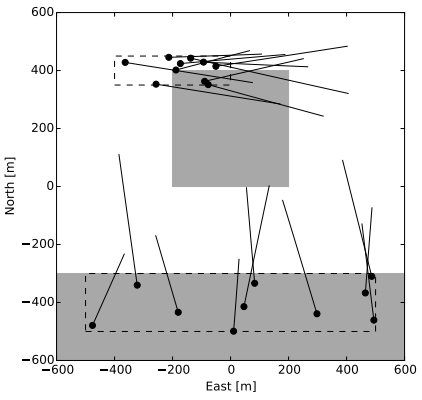

Fig. 2. Scenario setup with 10 sample trajectories of each target. The gray areas are high-clutter areas, while the dashed boxes show the initial positions of each target.

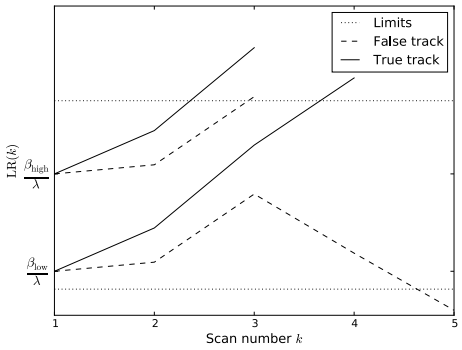

Fig. 3. The likelihood ratio for two SPRTs from the test scenario, one from a target, and one from clutter. The parameter β_{high} corresponds to using P_D/V_k as the new target density parameter.

B. Evaluation of new target density

We investigate the impact of the new target density parameter β, and its effect is only seen in the initial LR as illustrated in (30). The parameter β can thus be seen as a way to adjust the initial likelihood, as seen in Fig. 3. Choosing β high will confirm the true track after three scans, but the false track is also confirmed. If β is chosen lower, but still above the lower limit, the false track is terminated after five scans. This, however, comes at the expense of confirming the true track at scan four.

The new target density was varied logarithmically between the values which places the initial LR on the lower and upper bound, respectively. These bounds are chosen according to (13) by changing the inequalities to equality signs, ensuring the bounds on P_0 and P_1 are satisfied. The test was carried out for three different parameter sets and by spanning over the

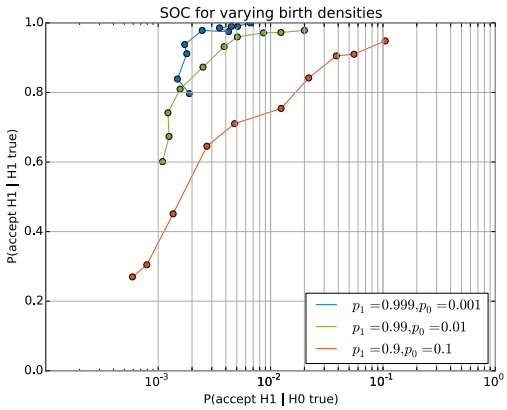

Fig. 4. Performance characteristic for varying new target density.

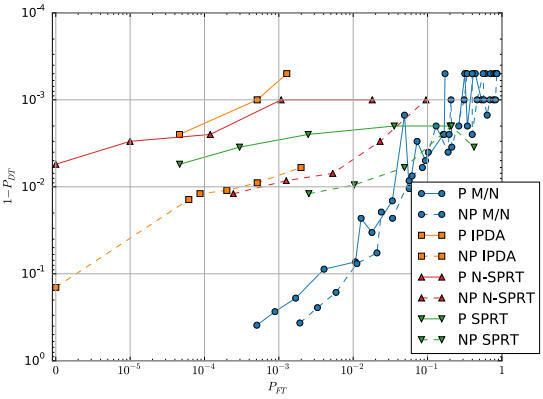

Fig. 5. Track initiation performance. The prefix "P" or "NP" denotes the parametric and nonparametric versions, respectively. N-SPRT is the SPRT in section III.

TABLE III
PARAMETER VARIATIONS FOR THE PERFORMANCE TEST

Method	Parameter	Value
M/N	N, M	1 to 7, 1 to N
IPDA	P_C, P_T	0.95 to 0.999, 0.1
SPRTs	P_0, P_1	10^{-2} to 10^{-5}, 0.99

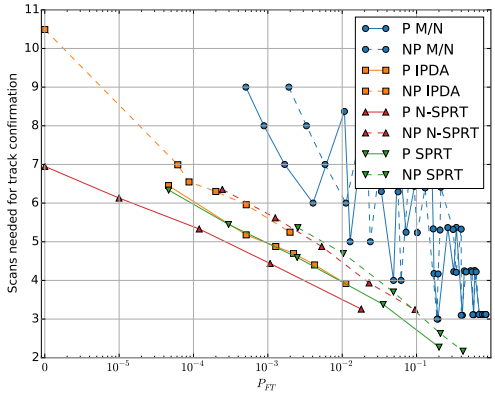

Fig. 6. Track initiation time. The prefix "P" or "NP" denotes the parametric and nonparametric versions, respectively. N-SPRT is the SPRT in section III.

different values of the new target density β. The performance curves can be seen in Fig. 4, which compares the probability of target detection to the probability of a false track. All the curves are rapidly increasing in detection probability, and flattens out as the detection probability approaches unity. The performance increase when P_1 increases and P_0 decreases. From (13), we see that the increasing P_1 or decreasing P_0 increase the upper bound and decrease the lower bound. As the limits move apart from each other, the more certain the SPRT need to be to conclude. This means that the choices of P_0 and P_1 will change the track confirmation or rejection time.

C. Performance comparison of track initiation methods

The tests are conducted as described in section IV-A, where 1000 simulations were run for clutter only, in order to evaluate P_{FT}, and 1000 simulations on each of the targets in order to evaluate P_{DT}. The threshold parameters were varied according to Table III. In the following, we will denote the SPRT with new target density as N-SPRT, to separate it from the traditional SPRT. Both the parametric versions, using the true clutter density, and the nonparametric versions were tested. For the SPRT with new target density, the value $\beta = 10^{-5}\,\mathrm{m}^{-2}$ was chosen.

Fig. 5 shows the probability of target detection and false target for the M/N, IPDA and the two SPRTs. The results have been sorted in increasing P_{FT}, and plotted against $1 - P_{DT}$. The y-axis have been reversed, such that a result higher up on the y-axis means higher track detection rate. Some of the methods failed to give any false tracks, and have been put at zero false track rate. Fig. 6 shows the average number of scans needed to confirm a true track.

The initiation method that sticks out the most is the M/N test. It is the only method that, despite the benign conditions in this test scenario, is able to achieve a target detection probability below 95%. It also have a very high false alarm rate. This is not surprising, as it is the only method that does not consider the possibility of the measurement being clutter. The SPRT with new target density and the IPDA performs similarly in their parametric versions, achieving excellent target detection probabilities for all the test values. The classical parametric SPRT also performs very well on target detection, but has a higher false alarm rate. The main difference can be seen in track initiation time, where the SPRT with new target prior

is approximately half a scan faster than the IPDA and the classical SPRT, which performs equally well.

The nonparametric versions all have slightly lower performance than their parametric counterparts. The difference is the lowest for the IPDA. This is also the only nonparametric method that contains the probability of track existence. The other methods do a hard count of the number of measurements, which makes the clutter density estimate potentially more fluctuating.

Although M/N only considers whether a measurement falls inside the validation gate or not, there is a notable difference between the parametric and nonparametric versions. Although the clutter density is not explicitly part of the track initiation evaluations, it enters through the association probabilities of the PDAF. The association probability of a misdetection, $P(\theta_k^0|Z^k)$, depends on the clutter density λ, while the others do not. In the nonparametric version, the clutter density is replaced by the ratio m_k/V_k. The total number of measurements in the validation gate m_k may contain the target-originated measurement as well as any clutter measurements, which makes the clutter density estimate too high. This increases $P(\theta_k^0|Z^k)$, which in turn increases the posterior covariance $P_{k|k}$ of the estimate through the state update equation. This will make the gate for the next scan larger, and prone to gating more clutter measurements, increasing the probability of reaching the test variable M.

V. Real data results

The data used in the experiments are shown in Fig. 7. In addition to the ownship, shown in grey, it shows two ships that entered the range of the radar during the experiment. The lower target were part of a COLAV experiment, hence the ownship maneuver to starboard. The target moving west was a ferry that exited a channel just as the experiment started. The targets are shown in Fig. 8. The total time of the experiment was three and a half minute. The bottom right half of the surveillance area have been masked out due to being close to land, which is why there are no measurements in that region. The rest of the region shows notable difference in amount of clutter. At the end of the experiment, the lower target does a sharp starboard maneuver to avoid the other target moving west. The ferry moving west also have an extended period when it is undetected, possibly due to obstructions from the other ship.

The nonparametric trackers were tested on the measurements from this scenario. The tracking results of the two SPRTs are show in Figures 9 and 10, and results from all the tracking methods are summarized in Table IV. The tracks from the real targets are similar in all the tests. Since the tracking system assumes a NCV model and has no interacting multiple model (IMM) capabilities, the lower target is lost after doing the hard maneuver. The track is initialized again after this maneuver, which results in two true tracks from the lower target. Additionally, the period of detection of the lower target also results in two true tracks for the ferry, in total four true tracks. All other targets have been labeled false.

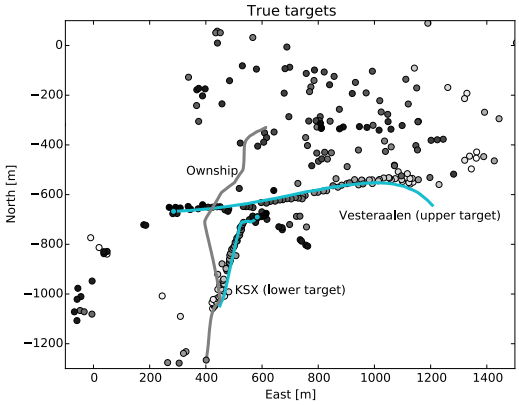

Fig. 7. Radar measurements and AIS tracks from the experiments. The measurements are shaded such that the darker measurements are from later in the experiments. The grey track is the ownship, on which the radar is mounted.

Fig. 8. The two targets present in the scenario. The top is the target of a coordinated COLAV experiment, KSX ocean space drone. The bottom is a passing ferry, Vesteraalen.

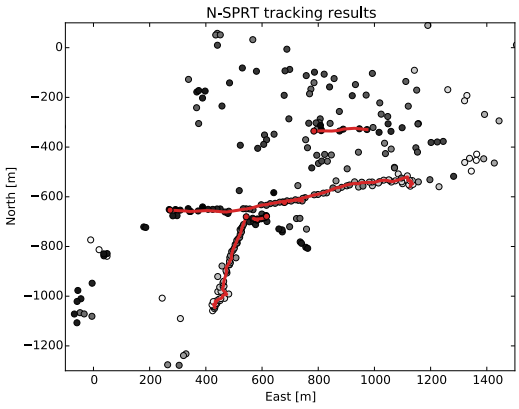

Fig. 9. Tracking results from the SPRT with new target density. The red dots mark the final target position estimate.

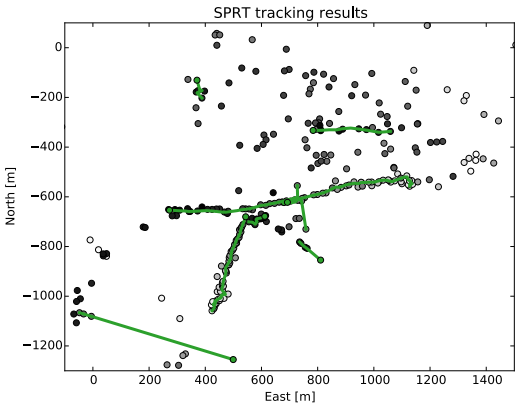

Fig. 10. Tracking results from the SPRT without new target density. The green dots mark the final target position estimate.

TABLE IV

Method	Number of false tracks	true track conf. time (scans)
M/N	3	[5, 5, 5, 5]
IPDA	0	[4, 5, 5, 5]
SPRT	5	[4, 5, 5, 9]
N-SPRT	1	[4, 5, 5, 6]

The IPDA has the best performance of the tested methods, while the classical SPRT has both the highest number of false tracks, and also takes very long to confirm a true track. Between the SPRTs, the main benefit of adding a new target density seems to be mostly beneficial for reducing the amount of false tracks, although track initiation times are also slightly improved.

The reason M/N gives fewer false tracks than the classical SPRT may be that M/N require two measurements close in successive scans in order to initiate a track, the SPRT allows for no measurements in the gate as long as the LR does not go below the termination threshold. Although this allows for more misdetections, it may enable more false tracks. Also recall from Fig. 5 that the nonparametric classical SPRT had a very high P_{FT}.

VI. Conclusion

A Bayesian formulation of the classical SPRT of Van Keuk using a new target density akin to the one used in Reid's MHT can improve track initialization when tracks are started in areas with small initial areas. The M/N, IPDA and two SPRTs have been compared, and all have a high probability of detection overall due to the good conditions in which the target is tracked. However, M/N is vastly outperformed on the amount of false tracks. The classical SPRT does slightly better, and SPRT with a new target density seems to perform comparably to the IPDA. The nonparametric versions of the track initiation methods performs slightly worse than their parametric counterparts, where the IPDA show the lowest difference. This is also reflected in tests on real data, where the IPDA had the lowest number of false tracks.

The non-parametric IPDA does a kind of implicit clutter estimation and the investigation of similar clutter estimation techniques for the MHT is a topic for future research. Also important is the success rate of track extraction from an SPRT, which has not been considered in this paper.

Acknowledgment

This work was supported by the Research Council of Norway through project number 244116 and the Centres of Excellence funding scheme with project number 223254. The authors would like to express great gratitude to Kongsberg Maritime and Maritime Robotics for providing the equipment and platforms used in the experiments. The professorship of the second author is funded by DNV GL.

References

[1] Y. Bar-Shalom and X.-R. Li, *Multitarget-Multisensor Tracking: Principles and Techniques*. YBS Publishing, 1995.

[2] E. Brekke and M. Chitre, "Success rates and posterior probabilities in multiple hypothesis tracking," in *Proceedings of the 21st International Conference on Information Fusion (FUSION) (accepted)*, 2018.

[3] E. Brekke, O. Hallingstad, and J. Glattetre, "The signal-to-noise ratio of human divers," in *OCEANS 2010 IEEE-Sydney*. IEEE, 2010.

[4] L. Elkins, D. Sellers, and W. R. Monach, "The autonomous maritime navigation (AMN) project: Field tests, autonomous and cooperative behaviors, data fusion, sensors, and vehicles," *Journal of Field Robotics*, vol. 27, no. 6, pp. 790–818, 2010.

[5] D. Hermann, R. Galeazzi, J. C. Andersen, and M. Blanke, "Smart sensor based obstacle detection for high-speed unmanned surface vehicle," *IFAC-PapersOnLine*, vol. 48, no. 16, pp. 190–197, 2015.

[6] G. V. Keuk, "Sequential track extraction," *IEEE Transactions on Aerospace and Electronic Systems*, vol. 34, no. 4, pp. 1135–1148, 1998.

[7] N. Li and X.-R. Li, "Target perceivability and its applications," *IEEE Transactions on Signal Processing*, vol. 49, no. 11, pp. 2588–2604, 2001.

[8] X. R. Li, N. Li, and V. P. Jilkov, "SPRT-based track confirmation and rejection," in *Proceedings of the Fifth International Conference on Information Fusion*, vol. 2. IEEE, 2002, pp. 951–958.

[9] X.-R. Li and V. P. Jilkov, "Survey of maneuvering target tracking. Part I. Dynamic models," *IEEE Transactions on aerospace and electronic systems*, vol. 39, no. 4, pp. 1333–1364, 2003.

[10] R. P. Mahler, *Statistical multisource-multitarget information fusion*. Artech House, Inc., 2007.

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

[12] D. Reid, "An algorithm for tracking multiple targets," *IEEE transactions on Automatic Control*, vol. 24, no. 6, pp. 843–854, 1979.

[13] A. Wald, "Sequential tests of statistical hypotheses," *The annals of mathematical statistics*, vol. 16, no. 2, pp. 117–186, 1945.

[14] M. Wieneke and W. Koch, "On sequential track extraction within the PMHT framework," *EURASIP J. Adv. Signal Process*, vol. 2008, pp. 90:1–90:13, Jan. 2008.

[15] J. L. Williams, "Marginal multi-Bernoulli filters: RFS derivation of MHT, JIPDA, and association-based member," *IEEE Transactions on Aerospace and Electronic Systems*, vol. 51, no. 3, pp. 1664–1687, 2015.

[16] E. F. Wilthil, A. L. Flåten, and E. F. Brekke, "A target tracking system for ASV collision avoidance based on the PDAF," in *Sensing and Control for Autonomous Vehicles: Applications to Land, Water and Air Vehicles*, T. I. Fossen, K. Y. Pettersen, and H. Nijmeijer, Eds. Springer International Publishing, 2017, pp. 269–288.

Bibliography

[1] O. Asplin. Performance evaluation of track initiation methods. Norwegian University of Science and Technology, Specialization Project, 2017.

[2] Y. Bar-Shalom and X. Li. *Multitarget-Multisensor-Tracking: Principles And Techniques*. YBS, 1995.

[3] E. Brekke, O. Hallingstad, and J. Glattetre. The signal-to-noise ratio of human divers. In *Proceedings of OCEANS'10*, Sydney, Australia, May 2010.

[4] R. Brown and P. Hwang. *Introduction to Random Signals and Applied Kalman Filtering*. John Wiley & Sons, Inc, 2012.

[5] P. Chanzy, L. Devroye, and C. Zamora-Cura. Analysis of range search for random k-d trees. *Acta Informatica*, 37(4):355–383, Jan 2001.

[6] T. Cover and P. Hart. Nearest neighbor pattern classification. *IEEE Transactions on Information Theory*, 13(1):21–27, January 1967.

[7] Y. Cui, J. Yang, Y. Yamaguchi, G. Singh, S. Park, and H. Kobayashi. On semiparametric clutter estimation for ship detection in synthetic aperture radar images. *IEEE Transactions on Geoscience and Remote Sensing*, 51(5):3170–3180, May 2013.

[8] International Maritime Orginization. AIS transponders - Regulations for carriage of AIS.

[9] G. Keuk. Sequential track extraction. *IEEE Transactions on Aerospace and Electronic Systems*, 34(4):1135–1148, 1998.

[10] T. Kirubarajan and Y. Bar-Shalom. Probabilistic data association techniques for target tracking in clutter. *Proceedings of the IEEE*, 92(3):536–557, Mar 2004.

[11] N. Levanon. *Radar signals*. Wiley, Hoboken, N.J, 2004.

[12] N. Li and X. Li. Target perceivability and its applications. *IEEE Transactions on Signal Processing*, 49(11):2588–2604, 2001.

[13] X. Li and V. Jilkov. Survey of maneuvering target tracking. part I. dynamic models. *IEEE Transactions on aerospace and electronic systems*, 39(4):1333–1364, 2003.

[14] D. Musicki and R. Evans. Joint integrated probabilistic data association: Jipda. *IEEE Transactions on Aerospace and Electronic Systems*, 40(3):1093–1099, July 2004.

[15] D. Musicki, R. Evans, and S. Stankovic. Integrated probabilistic data association. *IEEE Transactions on Automatic Control*, 39(6):1237–1241, 1994.

[16] D. Musicki and M. Morelande. Non parametric target tracking in non uniform clutter. In *2005 7th International Conference on Information Fusion*, July 2005.

[17] D. Musicki, S. Suvorova, M. Morelande, and B. Moran. Clutter map and target tracking. In *2005 7th International Conference on Information Fusion*, July 2005.

[18] D. Reid. An algorithm for tracking multiple targets. *IEEE Transactions on Automatic Control*, 24(6):843–854, December 1979.

[19] E. Wilthil, E. Brekke, and O. Asplin. Track initiation for maritime radar tracking with and without prior information. In *2018 21st International Conference on Information Fusion*, Accepted for publication at Fusion 2018, Cambridge, UK.

[20] E. Wilthil, A. Flåten, and E. Brekke. A target tracking system for asv collision avoidance based on the PDAF. *Sensing and Control for Autonomous Vehicles*, pages 269–288, 2017.