Probability Hypothesis Density Filter Implementation and Application

Zachary Prihoda, Arec Jamgochian, Ben Moore, Bernard Lange

June 12th 2019

Abstract

In this project, we explore multiple-target tracking and filtering in noise-cluttered sensing environments. Specifically, we formulate the Gaussian Mixture Probability Hypothesis Density (GM-PHD) filter, the analytical solution to multi-target Bayes filtering under linear Gaussian assumptions and time-varying numbers of targets. We implement the GM-PHD filter as a Julia package, showcase simulations of its effective use, and benchmark against examples in literature.

I Introduction

Making decisions in a robotic environment requires accurate state information of targets in the environment. The objective of multi-target tracking is to use sensor data to jointly estimate, at each time step, the number of targets in an environment and their states [1]. Typically, this requires relying on a sequence of noisy and cluttered observation sets and using probabilistic methods to form posterior estimates on state. This becomes particularly challenging in a multi-target environment, where not only do the states of the targets vary with time, but the number of targets also change as targets appear and disappear [1]. An active area of research revolves around correlating cluttered sensor data with objects evolving dynamically in the scene.

Many solution implementations involve gating noisy observations (e.g. radar data) above some threshold to estimate the presence of objects in a scene, then track-ing those objects individually. However, a large body of literature exists around modeling joint distributions over potential objects and propagating forward those distributions. Some implementations such as Probabilstic Data Association (PDA) [2], Joint Probabilistic Data Association (JPDA) [2, 3], and Multi-Hypothesis Tracking (MHT)[4], require explicit data association to propagate filters. Later research uses finite-set statistics (FISST) - namely the mathematical concept of Random Finite Sets (RFS) - to formulate multi-hypothesis Bayes filtering without requiring explicit data association [5]. These methods overcome the significant computational bottlenecks that accompany explicit data association when a there are a large number of targets. They also allow us to handle uncertainty in the number of targets at any point in time.

Using the RFS formulation, one can formulate the multi-target Bayes filter by propagating a Probability Hypothesis Density, or a first-order statistic of the RFS of states in time [5]. Under assumptions which we highlight in Section III, one can obtain solutions analogous to the Kalman filter for single-target Bayes. If we assume dynamics and measurements are linear Gaussian, and object birthing and spawning intensities follow Gaussian mixtures, then the Gaussian Mixture Probability Hypothesis Density (GM-PHD) filter is the resulting analytical-form filter [1].

Because of its inclusion of target birth, spawn, and death models, the GM-PHD filter is a useful solution to a wide range of multiple-target tracking problems in cluttered sensing environments. Some examples include:

- Modeling motion of unknown number of neighboring cars in a driving environment

- Tracking cars in an intersection as they come in and out of scene

- Monitoring airspace over an airport as airplanes take off and land

- Tracking motion of an aircraft carrier as it births aircraft

- Tracking multiple warheads of an incoming missile as they separate

In this project we discuss the formulation of the GM-PHD filter and implement it in a Julia package[1]. The rest of the paper is organized as follows. Section II describes background on multiple target tracking, including the differences and advantages of different filters. Section III goes into depth on the GM-PHD filter and its formulation. Section IV discusses testing of our implementation, as well as benchmarking with an example available in literature. Section V discusses another example implemented using our package. We conclude and discuss future work in section VI.

II Background

A Multi-target Tracking

With an established number of targets and number of measurements, one can filter target states using the multiple-hypothesis Kalman filter (MHKF), an extension of the Kalman filter which propagates a Gaussian Mixture Model (GMM) over potential states. However in practice, the number of sensor measurements is not fixed, but rather is cluttered. The notion of multiple target tracking (MTT) was introduced in [6] which used a Kalman Filter to estimate the target states from cluttered measurements. To account for the computational requirements, the set of targets are divided into independent clusters.

One extremely common approach to target tracking in noise-cluttered sensing environments is through the Probabilistic Data

Association (PDA) filter [2]. The PDA filter assumes that there only one target in that noise, that the track for the target has been initialized, that past information can be summarized by a Gaussian posterior sufficient statistic, if detected, there is only one true target measurement and many independent and identically distributed (i.i.d.) uniform in space, Poisson in number false alarms, and target detection occurs with known probability. Under these assumptions, the PDA filter works by selecting valid measurements in the validation region of the target at the current time, computing an association probability for each of those measurements, computing a weighted combined innovation based on these association probabilities, and using the combined innovation to update state.

The extension of this filter to multiple targets in the environment, where the number of targets is known, is the Joint Probabilistic Data Association (JPDA) filter [3, 2]. In the JPDA filter, the states are Gaussian with approximated conditional means and covariances propagated as a sufficient statistic. Measurements from one target can fall in the validation region of a neighboring target, and dynamics and measurement models for different targets do not have to be identical. In the JPDA filter, the measurement-to-target association probabilities are computed jointly, and the association probabilities are computed only for the latest set of measurements. The state estimation is then done separately as in a PDA filter in either a coupled or a decoupled fashion.

Alternatively, another method for MTT with explicit data association which many modern multiple target tracking systems implement is multiple hypothesis tracking (MHT) [4]. MHT typically includes gating observations, forming new tracks and hypotheses combinatorially from those gated observations and previously predicted tracks, evaluating those tracks and hypotheses, outputting most likely hypothesis to the user, and predicting tracks from the surviving tracks and hypothesis. [7] highlighted the major challenges of the assignment of the observations to the multiple targets and introduced an extension to the particle filter (i.e. non-parametric implementation of the Bayes filter) to do so.

A problem with all of these aforementioned

---

methods is their dependence on explicit data association, which can be a major bottleneck when the number of targets is high. Another issue is the assumption that the number of targets in clutter is known, which forces reliance on non-probabilistic methods for handling when targets are introduced to and leave from the scene. Because of these drawbacks, it is has been practical to formulate methods of filtering state estimation in high signal-to-noise-ratio (SNR) environments in ways that do not require explicit data association.

B Probability Hypothesis Surface, Density, and Filter

Towards modeling objects in a noise-cluttered scene evolving under a joint distribution without explicit data association, a formulation is made using Random Finite Sets (RFS), a concept from finite-set statistics (FISST) [8]. RFS-based approaches encode the number of states present at any given time as a random variable, and the value of each of those states a random vector. The RFS-based filtering methods have the benefits of not requiring explicit data association between measurements and targets, and not requiring the number of targets to be known explicitly. PHD filters allow therefore for Bayesian recursion considering random target births, spawns, and deaths.

The Probability Hypothesis Surface (PHS), which encodes the first moment of a RFS, was first introduced by Stein, Winter and Tenney in a series of unpublished manuscripts from 1993 to 1995 [9]. PHS has the property that when the surface is integrated over a region, it gives the expected number of agents [9]. Furthermore, the location of those agents can be inferred from the peaks of the PHS. This laid the foundation for the Probability Hypothesis Density (PHD) filter, whose propagating sufficient statistic - the Probability Hypothesis Density - is defined by these properties.

[9] outlines the first full account of the material discussed by Stein and Winter. Specifically, he discusses the background for multi-target moments, identifies the PHD as the first-order statistical moment, derives the recursive Bayes filtering equations, the prediction and the update equations for the PHD

filter. [5] attempted a Multi-target Bayes Filtering method using First-Order Multi-target Moments derived in [9], but ran into several issues. In particular, Mahler found that even for single-target problems, the filter was too computationally expensive. He goes so far as to state "Multi-target Bayes will never be of practical interest without the development of drastic but principled approximation strategies [5]. Despite his lack of initial success, he did lay some groundwork for later work.

The PHD Filter was later solved analytically for the case where targets follow a linear Gaussian dynamical model and the sensors follows a linear Gaussian measurement model. [1] introduces a tractable, exact algorithm for this case, with further assumptions that survival and detection probabilities are state independent, and intensities of birth and spawn Random Finite sets follow Gaussian mixtures. This algorithm, denoted the Gaussian Mixture Probability Hypothesis Density (GM-PHD) filter, provided significant improvement in computational efficiency over previous Monte-Carlo-based methods used to approximate solutions to the PHD filter recursion. The GM-PHD filter was shown in [1] to always have better tracking performance compared to the JPDA filter, and especially when the clutter rate increases or detection probability decreases. Uniform convergence in the GM-PHD Filter algorithm is proven and convergence results of the extended Kalman PHD Filter are given in [10]. The computational complexity of the PHD filter recursion is formalized in [11].

Previously, Vo and Singh had decided to approach the problem from an alternative direction. Rather than using optimal Bayesian multi-target filtering, linearizing the models (EK-PHD) or performing stochastic linearization (UK-PHD), they approximated the PHD by a set of weighted random samples which are propagated through time using importance sampling and re-sampling strategies. In order to do that, the probabilistic interpretation of the PHD was presented with the clever definition of the Prediction Operator and the Update Operator which enabled the application of the Sequential Monte Carlo methods (SMC). This became known as the Sequential Monte Carlo PHD (SMC-PHD) filter [12] and allowed a significantly less "restrictive" environment than a GM-PHD fil-

ter.

However, the major drawback is the approximation of the number of targets which is heavily unstable when using the SMC-PHD filter and the methods for the particles clustering (e.g. k-means and EM) which either assume the Gaussian density (EM) [13] or implicitly assume the spherical shape of the clusters (k-means) which becomes problematic when the sensors outputs do not exhibit radial symmetry (L2-Norm) or equivalently in L1-Norm interpretation. These issues were addressed by Ristic, Clark and Vo in a later paper [13] by intelligent thresholding of the weights assigned to each particle set created by each measurement.

[1] also outlines the extensions to the non-linear Gaussian models and puts forward PHD Filter analogs to the Extended Kalman Filter (EKF) and Unscented Kalman Filter (UKF). Extented Kalman and Unsceneded Kalman PHD filters (EK-PHD and UK-PHD) were shown to provide a good approximation while being computationally more efficient than the Particle rendition of the PHD filter [1].

Clark and Bell explored alternative suboptimal methods for multi-target tracking. Specifically, they tried two particle filter methods (k-mean clustering and mixture modelling) within the PHD framework which carry respective assumptions explained above. They found that both of their algorithms were able to tracking multi-target sets successfully when the assumptions hold [14].

One disadvantage of the GM-PHD and SMC-PHD filters is the lack of identity information of individual targets and hence lack of the temporal association for agents over time. In [15], they propose a multi-agent tracker extension to the GM-PHD filter which addresses this issues. However, further work is required to effectively track the crossing targets.

C PHD Filter Extensions

Since the application of Random Finite Sets to state estimation, significant research has been done extending the PHD Filter under differing assumptions.

The Cardinalized PHD (CPHD) filter propagates the entire probability distribution on number of targets in order to present a better-performance filter that is first-order in states of individual targets but higher-order in target number [16]. Mahler also proposed a recursion for propagating multi-target posterior density [11], namely the parameters of a multi-Bernoulli RFS that approximate the posterior multi-target RFS. An analytical Gaussian Mixture solution as well as a Sequential Monte Carlo solution which solve a cardinality bias in the original recursion formulation are presented in [17]. The resulting filter, which has gained recent popularity, is known as the Multi-Bernoulli filter.

Even more recently, the idea of labeled RFS was introduced to address target trajectories and their uniqueness by using conjugate priors [18]. Using labeled RFS theory, the labeled multi-Bernoulli (LMB) [19, 20] and generalized labeled multi-Bernoulli (GLMB) [21] filters have been formulated, and have performance advantages in track estimation with low signal to noise ratios.

While implementing these extensions qualifies as great future work, the main goal of this project is to explore and implement the GM-PHD filter.

III GM-PHD Filter [1]

For a more detail on this section, the reader is highly encouraged to read [1] on which this derivation highly based upon.

A Random Finite Set

The Random Finite Set (RFS) formulation, defines the collection of individual targets as a set-valued state, and the collection of individual observations as a set-valued observation. Given M targets with states x_k, N measurements z_k, at time-step k the respective collections can be represented as finite sets [1].

$$X_k = \{x_{k,1}, \ldots, x_{k,M(k)}\} \qquad (1)$$

$$Z_k = \{z_{k,1}, \ldots, z_{k,N(k)}\} \qquad (2)$$

The order in which the states are listed has no significance, the origins of the measurements are unknown and are indistinguishable from misdetections. Using the RFS formulation to model set-valued states and set-valued allows

the problem of dynamically estimating multiple targets in the presence of clutter and association uncertainty to be cast in a Bayesian filtering framework [1]. The RFS formulation enables the death and survival of targets at the next time step, surviving targets evolve to their new states and new targets can appear, resulting in new states.

B PHD Filter

The Probability Hypothesis Density (PHD) filter is a multiple-target filter for recursively estimating the number and the state of a set of targets given a set of observations [22].

Given a state x_{k-1} at time $k-1$, the probability density of a transition to the state x_k at time k is,

$$f_{k|k-1}(x_k|x_{k-1}) \qquad (3)$$

Given a state x_k at time k, the probability density of receiving the observation z_k is,

$$g_k(z_k|x_k) \qquad (4)$$

The probability density of the state x_k at time k given all observations up to time k is,

$$p_k(x_k|z_{1:k}) \qquad (5)$$

Given the computational complexity of the single-target Bayes nonlinear filtering equations, which is exacerbated in a multi-target case [1]. The multi-target Bayes filter propagates the multiple target posterior in time via the recursion,

$$p_{k|k-1}(X_k|Z_{1:k-1}) =$$
$$\int f_{k|k-1}(X_k|X)p_{k-1}(X|Z_{1:k-1})dX$$
$$p_k(X_k|Z_{1:k}) = \qquad (6)$$
$$\frac{g_k(Z_k|X_k)p_{k|k-1}(X_k|Z_{1:k-1})}{\int g_k(Z_k|X)p_{k|k-1}(X|Z_{1:k-1})dX}$$

Due to the computational complexity associated with the multi-target Bayes filter, an approximation strategy is required. The PHD filter propagates the intensity of the targets Random Finite Set (RFS) in time instead of the full multi-target posterior density [22].

The propagation of the intensity of the targets RFS is achieved using the following definitions. Let x be a fixed target state and Z is a data stream, $Z^k = [z_1, ..., z_k]$. Then for any $i >= 1$ the marginal-density value is given by:

$$\frac{1}{i!} \int f_{k|k}([x, x_1, ..., xi]|Z^{(k)})dx_1...dx_i \qquad (7)$$

Equation 7 is the total posterior likelihood that the multi-target system has $i + 1$ targets and that one of these targets has state x. Consequently, for each x the marginal-density value is given by;

$$v_{k|k}(x|Z^{(k)}) =$$
$$\sum_{i=1}^{\infty} \frac{1}{i!} \int f_{k|k}([x, x_1, ..., xi]|Z^{(k)})dx_1...dx_i \qquad (8)$$

Equation 8 is the total posterior likelihood that the multi-target system contains a target that has state x. As a result, $v_{k|k}(x|Z^{(k)})$ will tend to have maxima approximately at the locations of the targets [9]. Note that, the state vector x in a PHD filter represents the collection of multi-target states as the targets in the observation region which is constantly changing due to targets exiting or entering the observation region as well as misdetections, spontaneous target birthing and spawning. This state definition is in contrast to the conventional fixed single-target state.

A unique property of the PHD of density, $v_{k|k}(x|Z^{(k)})$ is that when it is integrated over a region in the target state space, $\int_S v_{k|k}(x|Z^{(k)})dx$ is equal to the expected number of targets contained in S [9].

As proven in [9], the PHD is equal to the expectation density or first factorial-moment density of the RFS at a time-step k. Given this property, the PHD, $v_{k|k}(x|Z^{(k)})$ is a form of least-squares best-fit approximation of the multi-target posterior $f_{k|k}(X|Z^{(k)})$. By propagating the PHD instead of the full multi-target posterior, the PHD is a multi-target analog of the single-target first-order moment [9].

It should be noted that the PHD is only a sufficient approximation of the full multi-target posterior in a multi-target sensing situation where the signal-to-clutter ratio (SCR)

as well as signal to noise ratio (SNR) is large. If this is not the case the filter target state estimates will diverge from the true target states.

The clutter model, target existence and birth and spawn dynamics for the filter are represented below where ζ is the previous state of a target [1]:

$\gamma_k(\cdot)$ intensity of the birth RFS Γ_k at time k;

$\beta_{k|k-1}(\cdot|\zeta)$ intensity of the RFS $\beta_{k|k-1}(\zeta)$ spawned at time k by a target with previous state ζ;

$p_{S,k}(\zeta)$ probability that a target still exists at time k given that its previous state is ζ;

$p_{D,k}(x)$ probability of detection given a state x at time k;

$\kappa_k(\cdot)$ intensity of clutter RFS K_k at time k;

The following assumptions are made to the clutter, RFS generation, target evolution and detection for the model [1].

1. Each target evolves and generates observations independently of one another.

2. Clutter is Poisson and independent of target-originated measurements.

3. The predicted multiple-target RFS governed by $p_{k|k-1}$ is Poisson.

Assumptions 1 and 2 are standard assumptions used in most tracking applications [1]. Assumption 3 is a reasonable approximation in applications where interactions between targets are negligible.

The posterior intensity is propagated using the PHD recursion given by:

$$v_{k|k-1}(x) = \int p_{S,k}(\zeta)f_{k|k-1}(x|\zeta)v_{k-1}(\zeta)d\zeta +$$

$$\int \beta_{k|k-1}(x|\zeta)v_{k-1}(\zeta)d\zeta + \gamma_k(x)$$

$$v_k(x) = [1 - p_{D,k}(x)]\,v_{k|k-1}(x) +$$

$$\sum_{z \in Z_k} \frac{p_{D,k}(x)g_k(z|x)v_{k|k-1}(x)}{\kappa_k(z)\int p_{D,k}(\xi)g_k(z|\xi)v_{k|k-1}(\xi)}$$

$$\tag{9}$$

C GM-PHD

C.1 Linear Gaussian Multi-target Model

One issue associated with the PHD recursion is that in general a closed-form solution cannot be found. Therefore, a linear Gaussian multi-target model is employed to find a closed-form solution. The linear Gaussian multi-target model requires the following assumptions [1].

Each target and sensor follows a linear Gaussian dynamical and measurement model [1]:

$$f_{k|k-1}(x|\zeta) = \mathcal{N}(x; F_{k-1}\zeta, Q_{k-1}) \tag{10}$$

$$g_k(z|x) = \mathcal{N}(z; H_k x, R_k) \tag{11}$$

where $N(\,\cdot\,; m, P)$ denotes a Gaussian density with mean m and covariance P, F_{k-1} is the state transition matrix, Q_{k-1} is the process noise covariance, H_k is the observation matrix, and R_k is the observation noise covariance.

The survival and detection probabilities of the targets are assumed to be independent of the state,

$$p_{S,k}(x) = p_{S,k}$$
$$p_{D,k}(x) = p_{D,k} \tag{12}$$

The intensities of the birth and spawn RFSs are modelled as Gaussian mixtures of the form;

$$\gamma_k(x) = \sum_{i=1}^{J_{\gamma,k}} w_{\gamma,k}^{(i)}\mathcal{N}(x; m_{\gamma,k}^{(i)}, P_{\gamma,k}^{(i)}) \tag{13}$$

where $J_{\gamma,k}, w_{\gamma,k}^{(i)}, m_{\gamma,k}^{(i)}, P_{\gamma,k}^{(i)}, i = 1, \ldots, J_{\gamma,k}$, are model parameters which determine the shape of the birth intensity of a new target. $m_{\gamma,k}^{(i)}, i = 1, \ldots, J_{\gamma,k}$ are the peaks of the spontaneous birth intensity in, which represent the highest local concentrations of expected number of spontaneous births. These peaks are locations where new targets are most likely to appear, for example a junction entrance or aircraft carrier. The covariance matrix $P_{\gamma,k}^{(i)}$ determines the spread of the birth intensity around the particular peak and the

weight $w_{\gamma,k}^{(i)}$ gives the expected number of new targets originating from the peak [1].

$$\beta_{k|k-1}(x|\zeta) =$$

$$\sum_{j=1}^{J_{\beta,k}} w_{\beta,k}^{(j)} \mathcal{N}(x; F_{\beta,k-1}^{(j)}\zeta + d_{\beta,k-1}^{(j)}, Q_{\beta,k-1}^{(j)}) \quad (14)$$

where $J_{\beta,k}$, $w_{\beta,k}^{(j)}$, $F_{\beta,k-1}^{(j)}$, $d_{\beta,k-1}^{(j)}$, and $Q_{\beta,k-1}^{(j)}, j = 1,\ldots, J_{\beta,k}$, are model parameters which determine the shape of the spawning intensity of a target with previous state ζ. The peak of the spawning intensity of a target is an affine function of the previous state ζ, $F_{\beta,k-1}^{(j)}\zeta + d_{\beta,k-1}^{(j)}$ is an affine function of , as a new spawned target is modelled to be spawned in the proximity of its parent. To extend the previous example mentioned where ζ is state of the aircraft carrier at time $k-1$, while $F_{\beta,k-1}^{(j)}\zeta + d_{\beta,k-1}^{(j)}$ is the expected state of airplanes spawned at time k.

C.2 Gaussian Mixture PHD Recursion

A closed form solution to the PHD recursion is presented by applying the following two propositions detailed in [1].

Proposition 1: Suppose that the assumptions made in the definition of the linear Gaussian Multi-target model hold and that the posterior intensity at time $k-1$ is a Gaussian mixture,

$$v_{k-1}(x) = \sum_{i=1}^{J_{k-1}} w_{k-1}^{(i)} \mathcal{N}(x; m_{k-1}^{(i)}, P_{k-1}^{(i)}) \quad (15)$$

Then, the predicted intensity for time k is a Gaussian mixture, given by,

$$v_{k|k-1}(x) = v_{S,k|k-1}(x) + v_{\beta,k|k-1}(x) + \gamma_k(x) \quad (16)$$

where,

$$v_{S,k|k-1}(x) = p_{S,k} \sum_{j=1}^{J_{k-1}} w_{k-1}^{(j)} \mathcal{N}(x; m_{S,k|k-1}^{(j)}, P_{S,k|k-1}^{(j)}) \quad (17)$$

---

given $\{w_{k-1}^{(i)}, m_{k-1}^{(i)}, P_{k-1}^{(i)}\}_{i=1}^{J_{k-1}}$, and the measurement set Z_k.

step 1. (prediction for birth targets)

$i = 0$.

for $j = 1, \ldots, J_{\gamma,k}$

$\quad i := i + 1$.

$\quad w_{k|k-1}^{(i)} = w_{\gamma,k}^{(j)}, \quad m_{k|k-1}^{(i)} = m_{\gamma,k}^{(j)}, \quad P_{k|k-1}^{(i)} = P_{\gamma,k}^{(j)}$.

end

for $j = 1, \ldots, J_{\beta,k}$

$\quad$ for $\ell = 1, \ldots, J_{k-1}$

$\quad\quad i := i + 1$.

$\quad\quad w_{k|k-1}^{(i)} = w_{k-1}^{(\ell)} w_{\beta,k}^{(j)}$,

$\quad\quad m_{k|k-1}^{(i)} = d_{\beta,k-1}^{(j)} + F_{\beta,k-1}^{(j)} m_{k-1}^{(\ell)}$,

$\quad\quad P_{k|k-1}^{(i)} = Q_{\beta,k-1}^{(j)} + F_{\beta,k-1}^{(j)} P_{k-1}^{(\ell)} (F_{\beta,k-1}^{(j)})^T$.

$\quad$ end

end

step 2. (prediction for existing targets)

for $j = 1, \ldots, J_{k-1}$

$\quad i := i + 1$.

$\quad w_{k|k-1}^{(i)} = p_{S,k} w_{k-1}^{(j)}$,

$\quad m_{k|k-1}^{(i)} = F_{k-1} m_{k-1}^{(j)}, \quad P_{k|k-1}^{(i)} = Q_{k-1} + F_{k-1} P_{k-1}^{(j)} F_{k-1}^T$,

end

$J_{k|k-1} = i$.

step 3. (construction of PHD update components)

for $j = 1, \ldots, J_{k|k-1}$

$\quad \eta_{k|k-1}^{(j)} = H_k m_{k|k-1}^{(j)}, \quad S_k^{(j)} = R_k + H_k P_{k|k-1}^{(j)} H_k^T$,

$\quad K_k^{(j)} = P_{k|k-1}^{(j)} H_k^T [S_k^{(j)}]^{-1}, \quad P_{k|k}^{(j)} = [I - K_k^{(j)} H_k] P_{k|k-1}^{(j)}$.

end

step 4. (update)

for $j = 1, \ldots, J_{k|k-1}$

$\quad w_k^{(j)} = (1 - p_{D,k}) w_{k|k-1}^{(j)}$,

$\quad m_k^{(j)} = m_{k|k-1}^{(j)}, \quad P_k^{(j)} = P_{k|k-1}^{(j)}$.

end

$\ell := 0$.

for each $z \in Z_k$

$\quad \ell := \ell + 1$.

$\quad$ for $j = 1, \ldots, J_{k|k-1}$

$\quad\quad w_k^{(\ell J_{k|k-1}+j)} = p_{D,k} w_{k|k-1}^{(j)} \mathcal{N}(z; \eta_{k|k-1}^{(j)}, S_k^{(j)})$.

$\quad\quad m_k^{(\ell J_{k|k-1}+j)} = m_{k|k-1}^{(j)} + K_k^{(j)}(z - \eta_{k|k-1}^{(j)})$.

$\quad\quad P_k^{(\ell J_{k|k-1}+j)} = P_{k|k}^{(j)}$.

$\quad$ end

$\quad w_k^{(\ell J_{k|k-1}+j)} := \dfrac{w_k^{(\ell J_{k|k-1}+j)}}{\kappa_k(z) + \sum_{i=1}^{J_{k|k-1}} w_k^{(\ell J_{k|k-1}+i)}}$, for $j = 1, \ldots, J_{k|k-1}$.

end

$J_k = \ell J_{k|k-1} + J_{k|k-1}$.

output $\{w_k^{(i)}, m_k^{(i)}, P_k^{(i)}\}_{i=1}^{J_k}$.

Algorithm 1: GM-PHD Filter Algorithm [1]

$$m_{S,k|k-1}^{(j)} F_{k-1} m_{k-1}^{(j)} \quad (18)$$

$$P_{S,k|k-1}^{(j)} Q_{k-1} + F_{k-1} P_{k-1}^{(j)} F_{k-1}^T \quad (19)$$

7

$$v_{\beta,k|k-1}(x) =$$

$$\sum_{j=1}^{J_{k-1}} \sum_{\ell=1}^{J_{\beta,k}} w_{k-1}^{(j)} w_{\beta,k}^{(\ell)} \mathcal{N}(x; m_{\beta,k|k-1}^{(j,\ell)}, P_{\beta,k|k-1}^{(j,\ell)}) \tag{20}$$

$$m_{\beta,k|k-1}^{(j,\ell)} = F_{\beta,k-1}^{(\ell)} m_{k-1}^{(j)} + d_{\beta,k-1}^{(\ell)} \tag{21}$$

$$P_{\beta,k|k-1}^{(j,\ell)} = Q_{\beta,k-1}^{(\ell)} + F_{\beta,k-1}^{(\ell)} P_{\beta,k-1}^{(j)} (F_{\beta,k-1}^{(\ell)})^T \tag{22}$$

Proposition 2: Again suppose that the the assumptions made in the definition of the linear Gaussian Multi-target model hold and that the posterior intensity at time k is a Gaussian mixture given by,

$$v_{k|k-1}(x) = \sum_{i=1}^{J_{k|k-1}} w_{k|k-1}^{(i)} \mathcal{N}(x; m_{k|k-1}^{(i)}, P_{k|k-1}^{(i)}) \tag{23}$$

Then, the posterior intensity at time k is also a Gaussian mixture and is given by,

$$v_k(x) = (1 - p_{D,k}) v_{k|k-1}(x) + \sum_{z \in Z_k} v_{D,k}(x; z) \tag{24}$$

where,

$$v_{D,k}(x; z) = \sum_{j=1}^{J_{k|k-1}} w_k^{(j)}(z) \mathcal{N}(x; m_{k|k}^{(j)}(z), P_{k|k}^{(j)}) \tag{25}$$

$$w_k^{(j)}(z) = \frac{p_{D,k} w_{k|k-1}^{(j)} q_k^{(j)}(z)}{\kappa_k(z) + p_{D,k} \sum_{\ell=1}^{J_{k|k-1}} w_{k|k-1}^{(\ell)} q_k^{(\ell)}(z)} \tag{26}$$

$$m_{k|k}^{(j)}(z) = m_{k|k-1}^{(j)} + K_k^{(j)}(z - H_k m_{k|k-1}^{(j)}) \tag{27}$$

$$P_{k|k}^{(j)} = [I - K_k^{(j)} H_k] P_{k|k-1}^{(j)} \tag{28}$$

$$K_k^{(j)} = P_{k|k-1}^{(j)} H_k^T (H_k P_{k|k-1}^{(j)} H_k^T + R_k)^{-1} \tag{29}$$

Notice that Propositions 1 and 2 are the prediction and update steps of the PHD recursion for a linear Gaussian multiple-target model, respectively. Proposition 2 provides closed-form expressions to calculate the means, covariances, and weights of v_k from previous the PHD of the previous step, v_{k-1} when a set of measurements at time k arrive.

The steps involved in the GM-PHD filter are detailed in Algorithm 1.

---

given $\{w_k^{(i)}, m_k^{(i)}, P_k^{(i)}\}_{i=1}^{J_k}$, a truncation threshold T, a merging threshold U, and a maximum allowable number of Gaussian terms J_{max}.
Set $\ell = 0$, and $I = \{i = 1, \ldots, J_k | w_k^{(i)} > T\}$.
repeat

$\quad \ell := \ell + 1$.
$\quad j := \arg\max_{i \in I} w_k^{(i)}$.
$\quad L := \left\{ i \in I \mid (m_k^{(i)} - m_k^{(j)})^T (P_k^{(i)})^{-1} (m_k^{(i)} - m_k^{(j)}) \leq U \right\}$.
$\quad \tilde{w}_k^{(\ell)} = \sum_{i \in L} w_k^{(i)}$.
$\quad \tilde{m}_k^{(\ell)} = \frac{1}{\tilde{w}_k^{(\ell)}} \sum_{i \in L} w_k^{(i)} x_k^{(i)}$.
$\quad \tilde{P}_k^{(\ell)} = \frac{1}{\tilde{w}_k^{(\ell)}} \sum_{i \in L} w_k^{(i)} (P_k^{(i)} + (\tilde{m}_k^{(\ell)} - m_k^{(i)})(\tilde{m}_k^{(\ell)} - m_k^{(i)})^T)$.
$\quad I := I \backslash L$.

until $I = \emptyset$.
if $\ell > J_{max}$ then replace $\{\tilde{w}_k^{(i)}, \tilde{m}_k^{(i)}, \tilde{P}_k^{(i)}\}_{i=1}^{\ell}$ by those of the J_{max} Gaussians with largest weights.
output $\{\tilde{w}_k^{(i)}, \tilde{m}_k^{(i)}, \tilde{P}_k^{(i)}\}_{i=1}^{\ell}$ as pruned Gaussian components.

Algorithm 2: Pruning Algorithm [1]

---

given $\{w_k^{(i)}, m_k^{(i)}, P_k^{(i)}\}_{i=1}^{J_k}$.
Set $\hat{X}_k = \emptyset$.
for $i = 1 \cdots, J_k$
$\quad$ if $w_k^{(i)} > 0.5$,
$\quad\quad$ for $j = 1, \ldots, \text{round}(w_k^{(i)})$
$\quad\quad\quad$ update $\hat{X}_k := \left[\hat{X}_k, m_k^{(i)} \right]$
$\quad\quad$ end
$\quad$ end
end
output $\hat{X}_k$ as the multi-target state estimate.

Algorithm 3: Extraction Algorithm [1]

As the GM-PHD filter propagates Gaussian mixtures in time, the number of Gaussian components required to represent the PHD, v_k grows un-boundedly over time. A time k the number of Gaussian components required to represent v_k is given by 30,

$$(J_{k-1}(1 + J_{\beta,k}) + J_{\gamma,k})(1 + |Z_k|) = \mathcal{O}(J_{k-1}|Z_k|) \tag{30}$$

where J_{k-1} is the number of components of v_{k-1}. In order to limit both the the number of Gaussian components and thus computation time the Gaussian components whose weights are below a certain threshold are discarded. This process is completed using the pruning algorithm detailed in Algorithm 2.

Once the posterior density v_k has been computed the multiple-target state estimates are then extracted using the extraction algorithm in Algorithm 3. The Gaussian mixture representation allow for easy extraction of the state estimates as the means of the Gaussian components are the local maxima of v_k. This process is far more complex for example, when using a particle-PHD filter, the estimated number of targets N^k is given by the total mass of the particles representing v_k and the estimated states are then obtained by partitioning these particles into N^k clusters.

Figure 4: 1-Dimensional Prune Test

IV Testing

In order to verify our implementation is working correctly, we generated a couple of tests and implemented an example from [1].

First, we wanted to test our implementation of the pruning method. We generated 1000 random Gaussian distributions with random weights and covariances and fed them through our pruning function described in Algorithm 2. This was performed for both a 1-dimensional and 2-dimensional case. The results of the tests are shown in figures 4 and 5. Note that after pruning there are much less high frequency fluctuations associated with low sigma Gaussians. The results are harder to see in figure 5, but under close inspection, it can be seen that the peaks are less noisy while maintaining their general shape.

From these results, we see that the pruning is working as expected. In both cases, we were able to reduce the number of Gaussian models from 1000 to less than 50 without losing too much fidelity in the model. An important property of the pruning algorithm is that is in how the expected number of targets over any region remains roughly the same. This is handled by clever merging of nearby Gaussian models. Pruning is extremely important in enabling the GM-PHD filter to be used in real-time applications without requiring an infeasible amount of computation time, as such the verification of this process is important in designing a usable GM-PHD filter.

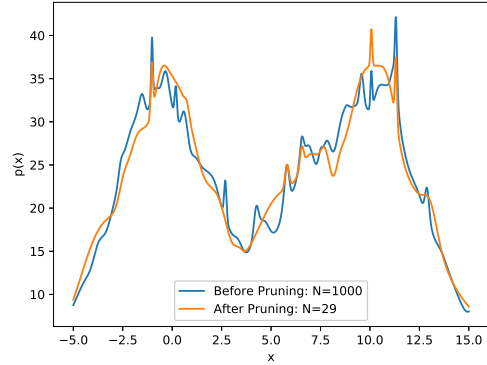Secondly, we wanted to test the extraction method. We generated 6 random Gaussian
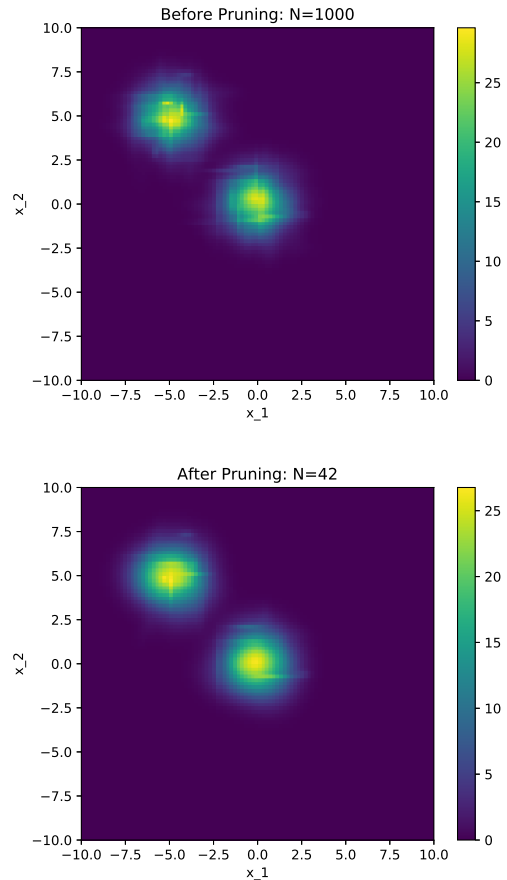

Figure 5: 2-Dimensional Prune

9

Figure 6: Extraction Tests

Figure 7: Surveillance Example
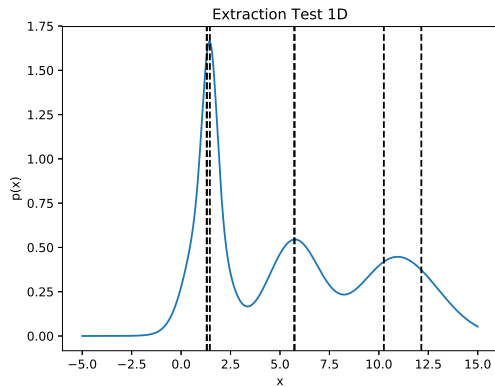
distributions with random weights and extracted the belief states from Gaussian Mixture. Once again, this was performed for a 1D and a 2D case. The results are shown in figure 6. The extracted states are represented by vertical dashed lines in the 1D case and crosses in the 2D case.

Once again, the tests verified that our functions were working properly. The extraction method is purely used to determine the expected locations of targets within our Gaussian Mixture. This is required for many applications that require an actual estimated state.

For the final test, to verify the full functionality of our implementation, we implemented the Surveillance example from [1]. This example runs through the entirety of the PHD filter algorithm stack (filter step, prune, extraction). The example consists of two tracking targets amid clutter and noisy measure-
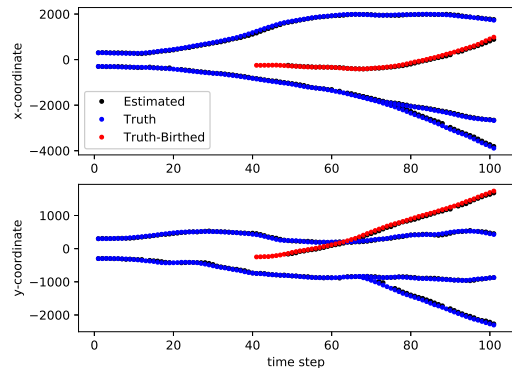
ments with an extra target spawning from one of the targets at an unknown. Additionally, to extend the example, we added a spontaneous birth to see if the filter would track it.

The resulting trajectories are shown in Figure 7. The blue dots represent the truth trajectories for the original and the spawned targets. The red dots represent the trajectory of the spontaneously birthed target. The black dots represent the output of the extracted states from the filtered GMM.

These results match those presented in [1]. We were successfully able to track the existing, spawned and birthed targets with no issues. With all of these tests done, we consider our implementation verified.

V Application

Now that we have built and tested our own implementation of the GM-PHD Filter, we want to display it's capabilities in a more physical and complex example.

Consider tracking a moving aircraft carrier and airplanes in the middle of the ocean. The aircraft carrier has its own dynamics and has the ability to spawn aircraft. Additionally, if we are monitoring a fixed region of the ocean, there is a possibility of aircraft entering the region (birthing) from the edges. Let's assume we are able to obtain a noisy measurement of the x and y positions of all targets within the region of interest.

From the PHD-Filters perspective, let's assume the dynamics can be modeled as a first

order integrator. The discrete dynamics can then be written as:

$$x_{k+1} = F_k x_k + w_k \tag{31}$$

Where:

$$x_k = [p_x, p_y, v_x, v_y]^T \tag{32}$$

$$F_K = \begin{bmatrix} I_2 & \Delta I_2 \\ 0_2 & I_2 \end{bmatrix} \tag{33}$$

$$w_k \sim \mathcal{N}(0, Q) \tag{34}$$

$$Q = \sigma_p^2 \begin{bmatrix} \frac{\Delta^4}{4} I_2 & \frac{\Delta^3}{2} I_2 \\ \frac{\Delta^3}{2} I_2 & \Delta^2 I_2 \end{bmatrix} \tag{35}$$

where Δ is the time step, σ_p is the standard-deviation of process noise and w_k is zero-mean Gaussian white noise.

The measurement model is simply:

$$y_k = C x_k + v_k \tag{36}$$

Where:

$$C = [I_2 \quad 0_2] \tag{37}$$

$$v_t \sim \mathcal{N}(0, R) \tag{38}$$

$$R = \sigma_m^2 I_2 \tag{39}$$

where σ_m is the standard deviation of the measurement noise and v_k is zero-mean Gaussian white noise.

Since we are dealing with a square region of interest, we define a birth region at each of the four edges. Specifically the birth model is:

$$w_\gamma = 0.01 \tag{40}$$

$$m_{\gamma 1} = [-r, 0, v, 0]^T \tag{41}$$

$$m_{\gamma 2} = [r, 0, -v, 0]^T \tag{42}$$

$$m_{\gamma 3} = [0, -r, 0, v]^T \tag{43}$$

$$m_{\gamma 4} = [0, r, 0, -v]^T \tag{44}$$

$$P_1 = P_2 = \text{Diag}([1, 2r, 1, v]) \tag{45}$$

$$P_3 = P_4 = \text{Diag}([2r, 1, v, 1]) \tag{46}$$

where r is the distance from the center of the region of interest to one of the edges, and v is the expected velocity of targets entering the space (we chose $r = 50, v = 5$). Note that we are assuming the expected birthing point to be the center of each edge, and we
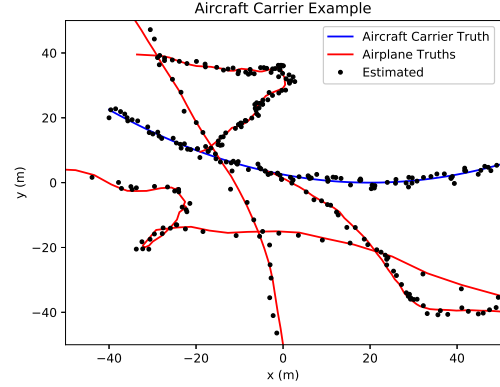

Figure 8: Aircraft Carrier Example: Trajectories

defined the uncertainty to be an extended ellipse along each edge. The birthing weight essentially states the expected probability that a target births from each of these regions per time step.

Since we expect targets to be spawning from the aircraft carrier, we define the spawn model as:

$$w_\beta = 0.01$$

$$Q_\beta = \text{Diag}([1.0, 1.0, 5.0, 5.0])$$

We choose Q such that the positional uncertainty is quite low, but the velocity uncertainty enables spawned targets to exit the aircraft carrier in any direction.

We choose the clutter density to be $\kappa = 12.5 * 10^{-6}$.

With all of this defined, we simulated a system with one aircraft carrier, two aircraft spawning from the aircraft, and two aircraft birthing from the boundaries. The results of running this scenario through our PHD Filter can be seen in Figures 8 and 9.

To explain Figure 8, the aircraft carrier's trajectory is outlined in blue. There are two aircraft that fly through the region of interest. The first enters from the bottom boundary and exits at the top. The second enters from the left, performs a small loop and exits through the right. The two spawned planes can be identified since their trajectories start along the aircraft carriers trajectory.

Figure 9 is simply the maximum of $p(x)$ along the time axis and shows an alternative depiction of the same information
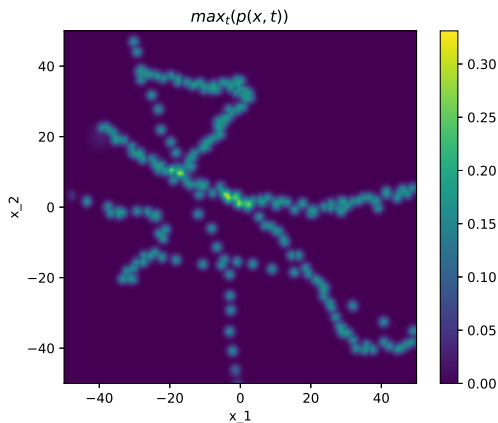
11

$max_t(p(x,t))$

Figure 9: Aircraft Carrier Example: probability map

Note that even though the control inputs are unknown and our assumed discrete dynamics describe a straight line, the uncertainty in the process noise allows us to track more complicated trajectories. Specifically, the filter is able to track an airplane as it performs a small loop maneuver.

Additionally, it is worth noting that it took a few measurement cycles before the our filter was able to track the airplane entering on the left. It simply didn't have enough information to produce an state estimate until it had obtained a few measurements.

VI Conclusion and Future Work

In this paper, we survey different methods for multi-hypothesis target filtering. We outline the formulation of the Gaussian Mixture Probability Hypothesis Density filter, the closed-form analytical solution to the multi-target Bayes filter under linear Gaussian assumptions. We implement the GM-PHD Filter in a Julia package and implement examples using the package. We compare our benchmarks against examples in literature, and highlight other uses for the filter.

Future work involves adding to the package to create a useful, well-maintained Gaussian filtering package Julia. This includes adding parametric single-target Bayes filter-

ing solutions (e.g. Kalman filter, Extended Kalman filter, Unscented Kalman filter), as well as extensions of the PHD filter under differing assumptions (e.g. Extended PHD filter, Unscented PHD filter, SMC-PHD filter, Multi-Bernoulli filter). We would also like to implement the Multi-Bernoulli filter, as well as multi-target trackers (e.g. GM-PHD tracker).

References

[1] B. . Vo and W. . Ma. The gaussian mixture probability hypothesis density filter. *IEEE Transactions on Signal Processing*, 54(11):4091–4104, Nov 2006.

[2] Yaakov Bar-Shalom, Fred Daum, and Jim Huang. The probabilistic data association filter. *IEEE Control Systems Magazine*, 29(6):82–100, 2009.

[3] Thomas Fortmann, Yaakov Bar-Shalom, and Molly Scheffe. Sonar tracking of multiple targets using joint probabilistic data association. *IEEE journal of Oceanic Engineering*, 8(3):173–184, 1983.

[4] Samuel S Blackman. Multiple hypothesis tracking for multiple target tracking. *IEEE Aerospace and Electronic Systems Magazine*, 19(1):5–18, 2004.

[5] Ronald PS Mahler. Multitarget bayes filtering via first-order multitarget moments. *IEEE Transactions on Aerospace and Electronic systems*, 39(4):1152–1178, 2003.

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

[7] C. Hue, J. . Le Cadre, and P. Perez. Tracking multiple objects with particle filtering. *IEEE Transactions on Aerospace and Electronic Systems*, 38(3):791–812, July 2002.

[8] Ronald PS Mahler. " statistics 101" for multisensor, multitarget data fusion. *IEEE Aerospace and Electronic Systems Magazine*, 19(1):53–64, 2004.

[9] Ronald P.S. Mahler. A theoretical foundation for the stein- winter " probability hypothesis density (phd)" multitarget tracking approach. 2000.

[10] D. Clark and B. Vo. Convergence analysis of the gaussian mixture phd filter. *IEEE Transactions on Signal Processing*, 55(4):1204–1212, April 2007.

[11] Ronald PS Mahler. *Statistical multisource-multitarget information fusion*. Artech House, Inc., 2007.

[12] Ba-Ngu Vo, Sumeetpal Singh, Arnaud Doucet, et al. Sequential monte carlo implementation of the phd filter for multi-target tracking. In *Proc. Intl Conf. on Information Fusion*, pages 792–799, 2003.

[13] Branko Ristic, Daniel Clark, and Ba-Ngu Vo. Improved smc implementation of the phd filter. In *2010 13th International Conference on Information Fusion*, pages 1–8. IEEE, 2010.

[14] D. E. Clark and J. Bell. Multi-target state estimation and track continuity for the particle phd filter. *IEEE Transactions on Aerospace and Electronic Systems*, 43(4):1441–1453, October 2007.

[15] K. Panta, D. E. Clark, and B. Vo. Data association and track management for the gaussian mixture probability hypothesis density filter. *IEEE Transactions on Aerospace and Electronic Systems*, 45(3):1003–1016, July 2009.

[16] Ronald Mahler. Phd filters of higher order in target number. *IEEE Transactions on Aerospace and Electronic systems*, 43(4):1523–1543, 2007.

[17] Ba-Tuong Vo, Ba-Ngu Vo, and Antonio Cantoni. The cardinality balanced multi-target multi-bernoulli filter and its implementations. *IEEE Transactions on signal processing*, 57(2):409–423, 2008.

[18] Ba-Tuong Vo and Ba-Ngu Vo. Labeled random finite sets and multi-object conjugate priors. *IEEE Transactions on Signal Processing*, 61(13):3460–3475, 2013.

[19] Stephan Reuter, Ba-Tuong Vo, Ba-Ngu Vo, and Klaus Dietmayer. The labeled multi-bernoulli filter. *IEEE Transactions on Signal Processing*, 62(12):3246–3260, 2014.

[20] Ba-Ngu Vo, Ba-Tuong Vo, and Dinh Phung. Labeled random finite sets and the bayes multi-target tracking filter. *IEEE Transactions on Signal Processing*, 62(24):6554–6567, 2014.

[21] Francesco Papi, Ba-Ngu Vo, Ba-Tuong Vo, Claudio Fantacci, and Michael Beard. Generalized labeled multi-bernoulli approximation of multi-object densities. *IEEE Transactions on Signal Processing*, 63(20):5487–5497, 2015.

[22] Michele Pace. Multi-target tracking with phd filters. *PHD Filters*.