PARAMETRIC DATA ASSOCIATION PRIOR FOR MULTI-TARGET TRACKING BASED ON RAO-BLACKWELLIZED MONTE CARLO DATA ASSOCIATION

Oliver Greß and Stefan Posch

Institute of Computer Science, Martin Luther University Halle-Wittenberg,
Von-Seckendorff-Platz 1, Halle (Saale), Germany

Keywords: Probabilistic Tracking, Multiple Targets, Data Association, Monte Carlo.

Abstract: Association of observations to underlying targets is a crucial task in probabilistic tracking of multiple targets. The Rao-Blackwellized Monte Carlo Data Association (RBMCDA) framework circumvents the combinatorial explosion by approximating the joint distribution of targets and association variables by Monte Carlo samples in the space of association variables. We present a parametric data association prior distribution required by RBMCDA, which models the formation of observations. To sample from this distribution an efficient algorithm is developed. The Interacting Multiple Models (IMM) filter is integrated into the RBMCDA framework to model the changing dynamics of targets aiming at tracking small particles in microscopy images. The proposed method is evaluated in a proof of concept and evaluated using synthetic data.

1   INTRODUCTION

Probabilistic tracking techniques are beneficial to track multiple targets with fast motion and similar appearance. Such methods employ a dynamic model to incorporate prior information. In detail, the state distribution of targets is predicted in each time step using the dynamic model and previous estimates. Subsequently this estimate is corrected using an observation of the target, if available. The major issue in multi-target tracking is the combinatorial problem to associate observations to underlying targets.

In (Särkkä et al., 2007) the Rao-Blackwellized Monte Carlo Data Association (RBMCDA) framework was presented for tracking targets with linear dynamics. RBMCDA circumvents the combinatorial problem of data association by sampling in the space of association variables, while distributions of targets' states are kept in analytical form. It also allows to integrate the Interacting Multiple Models (IMM) filter (Bar-Shalom and Blair, 2000) to account for changing dynamics of targets, which is required by the application we are aiming at, namely tracking of small particles in fluorescence microscopy images.

We propose a parametric *data association prior* required for sampling of associations, which models the formation of observations, because of the lack of a generic definition in (Särkkä et al., 2007).

The main contribution of this work is (i) a rigorous definition of the parametric data association prior and (ii) the development of an efficient algorithm to carry out the sampling procedure. We give a proof of concept of the proposed method evaluating performance of tracking on synthetic data.

The remainder of this work is structured as follows: Sec. 2 gives an overview of related work concerned with the problem of tracking multiple targets. An introduction to probabilistic multi-target tracking and the RBMCDA framework is given in Sec. 3. In Sec. 4 the parametric data association prior is developed as well as an efficient algorithm for sampling. Employed models, generation of data, methods for evaluation, and the complexity are described in Sec. 5. The performance is discussed in Sec. 6 and some concluding remarks are given in Sec. 7.

2   RELATED WORK

Various approaches exist for tracking multiple targets. Multi-Hypothesis Tracking (MHT) (Reid, 1979) provides an optimal solution to multi-target tracking for linear systems with Gaussian noise in theory. It further allows for the generation of new targets, extinction of targets as well as classifying observations as

387

clutter. In detail, MHT maintains a tree of hypotheses for all possible associations of observations, which results in an exponentially growing number of hypotheses. However, for practical application this number must be reduced rigorously by heuristic strategies like tree pruning or gating of observations.

The Joint Probabilistic Data Association (JPDA) filter (Bar-Shalom and Fortmann, 1987; Bar-Shalom and Blair, 2000) updates target state distributions by observations in a gate around the target. The updated state distribution results in a Gaussian mixture which is subsequently approximated by a single Gaussian to prevent exponential growth. JPDA was developed to provide robust estimates in environments with many clutter observations (Bar-Shalom and Fortmann, 1987). However, JPDA is limited to a fixed and known number of targets. JPDA was combined with IMM filtering (Bar-Shalom and Blair, 2000) to account for changing dynamics of targets and extended to overcome the limitation to a fixed number of targets (Musicki and Evans, 2002).

Sequential Monte Carlo (SMC) methods approximate state distributions not in closed-form as the above methods but by a finite set of samples (Doucet et al., 2001). SMC methods are capable to deal with nonlinear dynamics and non-Gaussian noise. SMC method are also available for multi-target tracking, but a large number of samples is necessary to avoid overly sparse representation of the high dimensional joint distribution of targets' states. Numerous extensions exist, e.g. to achieve IMM- and JPDA-like behavior (Blom and Bloem, 2003). Monte Carlo methods avoid the combinatorial problem of data association by sampling of association variables. In (Hue et al., 2002) the joint distribution of state and association variables is completely approximated by samples. To reduce the number of samples (Särkkä et al., 2007) introduce Rao-Blackwellization of state variables, i.e. their distributions are kept in analytical form while associations are still sampled.

Targets and observations often are assumed to be statistically independent for efficiency reasons. It thus is reasonable to model dependencies of targets only locally. In (Khan et al., 2005) the joint distribution of otherwise independent targets is penalized using Markov random fields, if targets approach. In (Smal et al., 2008) state samples are re-clustered to reconsider their affiliation to targets in microscopy images.

3 MULTI-TARGET TRACKING AND THE RBMCDA FRAMEWORK

The scenario considered for tracking of multiple targets is as follows: At each time step multiple targets might be present and their unknown number may vary with time because targets may (dis-)appear in the field of view. Individual targets are identified by unique target IDs which are preserved across time steps. The set of target IDs of existing targets at time t is denoted with $\mathfrak{N}^t$. The state of the target with ID n is denoted by $\mathbf{x}_n^t$. Let $\mathbf{X}^t := \{\mathbf{x}_n^t\}_{n \in \mathfrak{N}^t}$ describe the set of the states of the $N^t = |\mathbf{X}^t| = |\mathfrak{N}^t|$ existing targets at time t. Targets may "die" if they disappear and thus their ID and state have to be removed from the set of states and IDs. On the other hand new targets may be "born" in which case their state and ID must be inserted into the corresponding sets. With $\mathbf{X}_-^t$, N_-^t and $\mathfrak{N}_-^t$ we refer to the above sets and variables before target birth and death at time t and with $\mathbf{X}_+^t$, N_+^t and $\mathfrak{N}_+^t$ thereafter. Further on, multiple observations are detected at each time step, comprising observations from existing and newborn targets, but also false alarms. False alarms are referred to as *clutter* observations and do not originate from targets but are due to noise. Furthermore, a target may not be detected in all time steps. Let $\tilde{\mathbf{Z}}^t := (\tilde{\mathbf{z}}_1^t, \ldots, \tilde{\mathbf{z}}_{\tilde{M}^t}^t)$ be the vector of $\tilde{M}^t$ observations at time t, with an arbitrary but fixed order. Note that N^t and $\tilde{M}^t$ are implicitly defined by the number of elements in $\mathbf{X}^t$ and $\tilde{\mathbf{Z}}^t$ respectively. Let $\mathbf{C}^t := (c_1^t, \ldots, c_{\tilde{M}^t}^t)$ denote a vector of hidden association variables, where c_m^t describes association of the m-th observation $\tilde{\mathbf{z}}_m^t$ to a target n ($c_m^t = n$, $n \in \mathfrak{N}_+^t$) or clutter ($c_m^t = 0$). Note that realizations of random variables are marked with $\tilde{\ }$.

This multi target tracking problem can be solved by the iterative estimation of the joint distribution of targets' states analogous to the *predict&update-* scheme well known from single-target tracking (see e.g. (Doucet et al., 2001)). In fact, for the multi-target case we seek for an estimate of the joint distribution of targets' states and association variables given all observations up to the current time t as stated in Eq. (1). The factorization in Eq. (1) demonstrates the combinatorial problem inherent to multi-target tracking: The first term corresponds to the estimate of targets' state distribution for fixed association variables, i.e. we know which observation to use for updating which target. However, estimation of the complete distribution is prohibitive for practical application because the number of valid associations $\mathbf{C}^{1:t}$ to be considered explodes with time and increasing number of targets

and observations.

$$p(\mathbf{X}^t, \mathbf{C}^{1:t} | \tilde{\mathbf{Z}}^{1:t}) = p(\mathbf{X}^t | \tilde{\mathbf{Z}}^{1:t}, \mathbf{C}^{1:t}) \cdot p(\mathbf{C}^{1:t} | \tilde{\mathbf{Z}}^{1:t}) \quad (1)$$

For this reason (Särkkä et al., 2007) propose to approximate Eq. (1) by a finite set of samples $\{^{(s)}\mathbf{C}^{1:t}\}_{s=1,\dots,S}$ in the space of association variables whereas the distributions of targets' states are maintained in analytical form, e.g. by Gaussian distributions.

$$p(\mathbf{X}^t, \mathbf{C}^{1:t} | \tilde{\mathbf{Z}}^{1:t}) \approx$$
$$\frac{1}{S} \sum_{s=1}^{S} p(\mathbf{X}^t | \tilde{\mathbf{Z}}^{1:t}, {}^{(s)}\tilde{\mathbf{C}}^{1:t}) \cdot \delta(\mathbf{C}^t - {}^{(s)}\tilde{\mathbf{C}}^t), \quad (2)$$

with ${}^{(s)}\tilde{\mathbf{C}}^t$ sampled from $p(\mathbf{C}^t | \tilde{\mathbf{Z}}^{1:t}, {}^{(s)}\tilde{\mathbf{C}}^{1:t-1})$ (3)

For each sample ${}^{(s)}\tilde{\mathbf{C}}^{1:t}$ a separate joint distribution of its target states $p(\mathbf{X}^t | \tilde{\mathbf{Z}}^{1:t}, {}^{(s)}\tilde{\mathbf{C}}^{1:t})$ is maintained. This combination of sampled associations and corresponding state distributions will be referred to as a RBMCDA sample. Estimating the distribution of targets' states follows the *predict&update*-scheme mentioned above with an intermediate sampling step to yield realizations ${}^{(s)}\tilde{\mathbf{C}}^t$ before update.

According to (Särkkä et al., 2007) the individual association variables in ${}^{(s)}\tilde{\mathbf{C}}^t$ are sampled sequentially by decomposition of Eq. (3) using Bayes' rule resulting in Eq. (4). The term for the m-th association variable in Eq. (4) is further decomposed (Eq. (5)) into the predicted likelihood of the corresponding observation for the given association and the data association prior, which does not include information about the actual observation.

$$p(\mathbf{C}^t | \tilde{\mathbf{Z}}^{1:t}, \mathbf{C}^{1:t-1}) = p(c_1^t | \tilde{\mathbf{Z}}^{1:t}, \mathbf{C}^{1:t-1})$$
$$\times \prod_{m=2}^{\tilde{M}^t} p(c_m^t | c_{1:m-1}^t, \tilde{\mathbf{Z}}^{1:t}, \mathbf{C}^{1:t-1}) \quad (4)$$

$$p(c_m^t | c_{1:m-1}^t, \tilde{\mathbf{Z}}^{1:t}, \mathbf{C}^{1:t-1})$$
$$\propto p(\tilde{\mathbf{z}}_m^t | c_m^t, \tilde{\mathbf{Z}}^{1:t-1}, \mathbf{C}^{1:t-1}) p(c_m^t | c_{1:m-1}^t, \tilde{M}^t, \mathfrak{N}_-^t) \quad (5)$$

The formulation in Eq. (5) requires the assumption that c_m^t is independent from observations that are not yet associated, i.e. $\tilde{\mathbf{z}}_{m+1:\tilde{M}^t}^t$. This assumption is allowed by the scenario considered in (Särkkä et al., 2007) and exploited to achieve efficient sampling of association variables in Sec. 4, but available information clearly remains unused in our case.

In (Särkkä et al., 2007) only an informal description is provided of how to compute data association prior probabilities. Therefore, we present a rigorous formulation of a data association prior in 4.1 and an efficient sampling algorithm using this prior in 4.2.

4 PARAMETRIC DATA ASSOCIATION PRIOR

4.1 Model of the Data Association Prior

In the following we derive the data association prior from the joint distribution of $\mathbf{C}^t$ when only the number of existing targets N_-^t and their IDs $\mathfrak{N}_-^t$ is known. This joint distribution models the formation of a set of M^t observations composed of k^t observations from the N_-^t existing targets, b^t observations from newborn targets and $u^t = M^t - k^t - b^t$ clutter observations. Existing targets are assumed to be detected with probability P_D. The number of newborn targets, which are all presume to be detected, is assumed to be distributed according to $\nu(b^t)$ and the number of clutter observations according to the distribution $\mu(u^t)$. The probability of a valid set of association variables $\mathbf{C}^t$ is composed of the probability to receive $M^t = k^t + b^t + u^t$ observations and the probability of the specific configuration of association variables, given that the number of existing targets is known. The *joint* data association prior (6) is proportional to this distribution for a fixed number of observations M^t. It is defined to be 0 if the cardinality of $\mathbf{C}^t$ is different from M^t. Otherwise it is proportional to the described distribution:

$$p(\mathbf{C}^t | M^t, \mathfrak{N}_-^t) \propto p(\mathbf{C}^t | \mathfrak{N}_-^t) = \frac{1}{\binom{M^t}{k^t} \binom{N_-^t}{k^t} k^t! \binom{M^t - k^t}{b^t}}$$
$$\times \binom{N_-^t}{k^t} P_D^{k^t} (1 - P_D)^{N_-^t - k^t} \nu(b^t) \mu(u^t) \quad (6)$$

The fraction in Eq. (6) accounts for the number of possible configurations of $\mathbf{C}^t$ for $M^t = k^t + b^t + u^t$ observations. k^t out of M^t observations originate from existing targets with $\binom{N^t}{k^t} k^t!$ different alternatives to assign target IDs. A total of b^t of the remaining $(M^t - k^t)$ observations stem from newborn targets with $\binom{M^t - k^t}{b^t}$ alternatives to choose b^t observations. $u^t = M^t - k^t - b^t$ observations are due to clutter.

The data association prior of an association c_m^t required in Eq. (5) is proportional to Eq. (6) marginalized over all valid configurations $c_{m+1:M^t}^t$ given $\tilde{c}_{1:m-1}^t$ denoted by $\{c_{m+1:M^t}^t | \tilde{c}_{1:m-1}^t\}$:

$$p(c_m^t | \tilde{c}_{1:m-1}^t, M^t, \mathfrak{N}_-^t) = \frac{p(\tilde{c}_{1:m-1}^t, c_m^t | M^t, \mathfrak{N}_-^t)}{p(\tilde{c}_{1:m-1}^t | M^t, \mathfrak{N}_-^t)}$$
$$\propto \sum_{\{c_{m+1:M^t}^t | \tilde{c}_{1:m-1}^t\}} p(\tilde{c}_{1:m-1}^t, c_{m:M^t}^t | \mathfrak{N}_-^t) \quad (7)$$
$$=: q(c_m^t | \tilde{c}_{1:m-1}^t, M^t, \mathfrak{N}_-^t) \quad (8)$$

4.2 Sampling Algorithm

4.2.1 Computation of Data Association Prior

Three different cases have to be considered. The first case is the association of an observation to clutter, i.e. $c_m^t = 0$. The second case is the association to one of the existing targets that is $c_m^t = n_{\text{exist}}$. Association to a newborn target as the third case is given by $c_m^t = n_{\text{new}}$. The marginalization Eq. (7) can be expressed via a sum over possible numbers of associations to existing targets k in $\mathbf{C}^t$ given $\tilde{c}_{1:m-1}^t$ are already assigned (see e.g. Eq (13)). Conditioned on k, an inner sum runs over the possible number of newborn target associations b in $\mathbf{C}^t$ given $\tilde{c}_{1:m-1}^t$. All configurations of $\mathbf{C}^t$ with identical k and b have the same probability, c.f. Eq. (6).

For brevity of notation we use additional definitions in the following. $M := M^t$ denotes the number of current observations and $N := N_-^t = N_+^{t-1}$ is the number of targets that already exist before time step t.

Case 1: Association to Clutter ($c_m = 0$). We consider all feasible numbers k of associations to existing targets in $\mathbf{C}^t$ given $\tilde{c}_{1:m-1}^t$, if the m-th observation is associated to clutter. The minimum number is k_{m-1}, which is defined as the number of associations to existing targets present in $\tilde{c}_{1:m-1}^t$. The maximum number equals $k_m^{\max} := \min(N, M - m + k_{m-1})$ since a maximum of $M - m$ additional associations are possible in $\tilde{c}_{m+1:M}^t$ and we have a total of N existing target available. Likewise we consider the number of associations to newborn targets in $\mathbf{C}^t$ given $\tilde{c}_{1:m-1}^t$. The minimum number corresponds to the number of previous associations to newborn targets in $\tilde{c}_{1:m-1}^t$ denoted by b_{m-1}. The maximum number is $b_m^{\max} := M - m - k + k_{m-1} + b_{m-1}$ as a maximum of $M - m - k$ observations may additionally be associated to clutter. The sum over valid configurations of associations $c_{m+1:M}^t$ given $k_{m-1}, b_{m-1}, \tilde{c}_{1:m-1}^t$ and $c_m^t = 0$ required in Eq. (7) can be expressed by the number of these configurations and is given in Eq. (13). There are $\binom{M-m}{k-k_{m-1}}$ alternatives to select $(k - k_{m-1})$ out of $(M - m)$ observations for additional association with existing targets. $\binom{N-k_{m-1}}{k-k_{m-1}}(k - k_{m-1})!$ alternatives to associate $(k - k_{m-1})$ of the $(N - k_{m-1})$ non-associated existing targets exclusively and $\binom{M-m-k+k_{m-1}}{b-b_{m-1}}$ alternatives to associate $(b - b_{m-1})$ out of the remaining $(M - m - k + k_{m-1})$ observations to newborn targets.

To highlight relations between the different association cases the following definitions will be used:

$$d(m, k_{m-1}; M, N) := \frac{(M-m)! \, (N-k_{m-1})!}{M! \, N!} \quad (9)$$

$$f(k, k_{m-1}; N) := \frac{k! \cdot \binom{N}{k} P_D^k (1 - P_D)^{N-k}}{(k - k_{m-1})!} \quad (10)$$

$$g(b, b_{m-1}) := \frac{b! \cdot \nu(b)}{(b - b_{m-1})!} \quad (11)$$

$$h(m, r, r_{m-1}; M) := \frac{(M-r)! \cdot \mu(M-r)}{(M-r-m+r_{m-1})!} \quad (12)$$

Here we denote $r = k + b$ and $r_{m-1} = k_{m-1} + b_{m-1}$. Their properties are examined in detail in 4.2.2 and are utilized for efficient computation.

Case 2: Association to Existing Target ($c_m = n_{\text{exist}}$). The marginal in Eq. (7) for association to an existing target differs from the clutter case in the sum limits for k and is given in Eq. (14). The limits of the clutter case have to be increased by one, because the number of associations to existing targets in $c_{1:m}$ is $(k_{m-1} + 1)$ in this case. This aspect has also to be considered in the number of possible associations in $c_{m+1:M}$.

Case 3: Association to Newborn Target ($c_m = n_{\text{new}}$). For association to a newborn target the sum limits for b have to be adjusted with regard to the clutter case (see Eq. (15)). Similar to the adjustment of the k-limits in case 1, the limits of b have to be increased by one, because the number of associations to newborn targets in $c_{1:m-1}$ and c_m now equals to $(b_{m-1} + 1)$.

4.2.2 Iterative Properties of Association Prior

The functions d, f, g and h introduced in equations (9) to (12) can be computed iteratively and help to understand the iterative computation of terms used to sequentially sample the data association prior.

If m is incremented $d(m, k_{m-1}; M, N)$ can be easily updated regardless whether $k_m = k_{m-1} + 1$ (Case 2) or $k_m = k_{m-1}$ (Case 1 and Case 3):

$$d(0, 0; M, N) = 1 \quad (16)$$

$$d(m+1, k_{m-1}; M, N) = \frac{1}{M-m} \, d(m, k_{m-1}; M, N)$$

$$d(m, k_{m-1} + 1; M, N) = \frac{1}{N - k_{m-1}} \, d(m, k_{m-1}; M, N)$$

The values of $f(k, k_{m-1}; N)$ only change if an association to an existing target is sampled (Case 2):

$$f(k, 0; N) = \binom{N}{k} P_D^k (1 - P_D)^{N-k} \quad (17)$$

$$f(k, k_{m-1} + 1; N) = (k - k_{m-1}) \, f(k, k_{m-1}; N)$$

Likewise the function $g(b, b_{m-1})$ changes if association to a newborn target occurs (Case 3).

$$g(b, 0) = \nu(b) \quad (18)$$

$$g(b, b_{m-1} + 1) = (b - b_{m-1}) \, g(b, b_{m-1})$$

$$q(c_m^t = 0 | \tilde{c}_{1:m-1}^t, M, \mathfrak{N}_-) =$$

$$= \sum_{k=k_{m-1}}^{k_m^{\max}} \sum_{b=b_{m-1}}^{b_m^{\max}} \binom{M-m}{k-k_{m-1}} \binom{N-k_{m-1}}{k-k_{m-1}} (k-k_{m-1})! \binom{M-m-k+k_{m-1}}{b-b_{m-1}} p(\mathbf{C}^t|N)$$

$$= \sum_{k=k_{m-1}}^{k_m^{\max}} \sum_{b=b_{m-1}}^{b_m^{\max}} \frac{\binom{M-m}{k-k_{m-1}} \binom{N-k_{m-1}}{k-k_{m-1}} (k-k_{m-1})! \binom{M-m-k+k_{m-1}}{b-b_{m-1}}}{\binom{M}{k} \binom{N}{k} k! \binom{M-k}{b}} \binom{N}{k} P_D^k (1-P_D)^{N-k} \nu(b) \mu(M-k-b)$$

$$= \frac{(M-m)! \, (N-k_{m-1})!}{M! \, N!} \left[\sum_{k=k_{m-1}}^{k_m^{\max}} \frac{k!}{(k-k_{m-1})!} \binom{N}{k} P_D^k (1-P_D)^{N-k} \right.$$

$$\left. \times \left[\sum_{b=b_{m-1}}^{b_m^{\max}} \frac{b!}{(b-b_{m-1})!} \nu(b) \frac{(M-k-b)!}{(M-k-b-m+k_{m-1}+b_{m-1})!} \mu(M-k-b) \right] \right]$$

$$= d(m, k_{m-1}; M, N) \sum_{k=k_{m-1}}^{k_m^{\max}} f(k, k_{m-1}; N) \sum_{b=b_{m-1}}^{b_m^{\max}} g(b, b_{m-1}) \, h(m, k+b, k_{m-1}+b_{m-1}; M) \tag{13}$$

$$q(c_m^t = n_{\text{exist}} | \tilde{c}_{1:m-1}^t, M, \mathfrak{N}_-) =$$

$$= \sum_{k=k_{m-1}+1}^{k_m^{\max}+1} \sum_{b=b_{m-1}}^{b_m^{\max}} \frac{\binom{M-m}{k-k_{m-1}-1} \binom{N-k_{m-1}-1}{k-k_{m-1}-1} (k-k_{m-1}-1)! \binom{M-m-k+k_{m-1}+1}{b-b_{m-1}}}{\binom{M}{k} \binom{N}{k} k! \binom{M-k}{b}} \binom{N}{k} P_D^k (1-P_D)^{N-k} \nu(b) \mu(M-k-b)$$

$$= d(m, k_{m-1}+1; M, N) \sum_{k=k_{m-1}+1}^{k_m^{\max}+1} f(k, k_{m-1}+1; N) \sum_{b=b_{m-1}}^{b_m^{\max}} g(b, b_{m-1}) h(m, k+b, k_{m-1}+1+b_{m-1}; M) \tag{14}$$

$$q(c_m^t = n_{\text{new}} | \tilde{c}_{1:m-1}^t, M, \mathfrak{N}_-) =$$

$$= \sum_{k=k_{m-1}}^{k_m^{\max}} \sum_{b=b_{m-1}+1}^{b_m^{\max}+1} \frac{\binom{M-m}{k-k_{m-1}} \binom{N-k_{m-1}}{k-k_{m-1}} (k-k_{m-1})! \binom{M-m-k+k_{m-1}}{b-b_{m-1}-1}}{\binom{M}{k} \binom{N}{k} k! \binom{M-k}{b}} \binom{N}{k} P_D^k (1-P_D)^{N-k} \nu(b) \cdot \mu(M-k-b)$$

$$= d(m, k_{m-1}; M, N) \sum_{k=k_{m-1}}^{k_m^{\max}} f(k, k_{m-1}; N) \sum_{b=b_{m-1}+1}^{b_m^{\max}+1} g(b, b_{m-1}+1) h(m, k+b, k_{m-1}+b_{m-1}+1; M) \tag{15}$$

If both m and r_{m-1} are incremented simultaneously (Cases 2 and 3) h remains unchanged. If m is incremented but the total number target associations r_{m-1} remains unchanged (Case 1) the last recursion is used:

$$h(1, r, 0; M) = (M-r) \, \mu(M-r) \tag{19}$$
$$h(1, r, 1; M) = \mu(M-r)$$
$$h(m+1, r, r_{m-1}+1; M) = h(m, r, r_{m-1}; M)$$
$$h(m+1, r, r_{m-1}; M) = (M-r-m+r_{m-1})$$
$$\times h(m, r, r_{m-1}; M)$$

4.2.3 Sampling Algorithm

To sample a complete vector of associations $^{(s)}\tilde{\mathbf{C}}^t$ for time step t given $^{(s)}\tilde{\mathbf{C}}^{1:t-1}$, the individual realization $^{(s)}\tilde{c}_m^t$ are sampled sequentially for $m = 1$ to $m = M$. An iteration consists of two parts. First the probabilities for sampling (see Eq. (5)) are computed. Subsequently these probabilities are used to sample

the m-th association $^{(s)}\tilde{c}_m^t$. In the second part variables and arrays representing functions d, f, g and h are updated if necessary using the recursions derived in the last subsection. This sampling is performed independently for each RBMCDA sample $^{(s)}\tilde{\mathbf{C}}^t$ in turn. For clarity the RBMCDA sample index $^{(s)}$ is omitted in the following description of the algorithm.

INITIALIZE
$d(1, 0; M, N)$, $f(k, 0; M, N)$, $f(k, 1; M, N)$
$g(b, 0)$, $g(b, 1)$, $h(1, r, 0; M)$, $h(1, r, 1; M)$
$k_{m-1} \leftarrow 0$, $b_{m-1} \leftarrow 0$
$p_{\tilde{C}} \leftarrow 1$ // $p(\mathbf{C}^t | \mathbf{Z}^{1:t}, \tilde{\mathbf{C}}^{1:t-1})$
$\tilde{\mathbf{C}} \leftarrow \{\}$ // $\tilde{\mathbf{C}}^t$

// SEQUENTIAL SAMPLING of $\mathbf{C}^t$
for $m = 1 \rightarrow M$ **do**
    // Part I: Compute probabilities and sample $\tilde{c}_m^t$

    // compute data association prior probabilities
    $q_{\text{clutter}} \leftarrow q(c_m^t = 0 | \tilde{c}_{1:m-1}, M, N)$ // Eq. (13)

(a) $N_{\text{init}} = 50$.   (b) $N_{\text{init}} = 100$.   (c) $N_{\text{init}} = 200$.
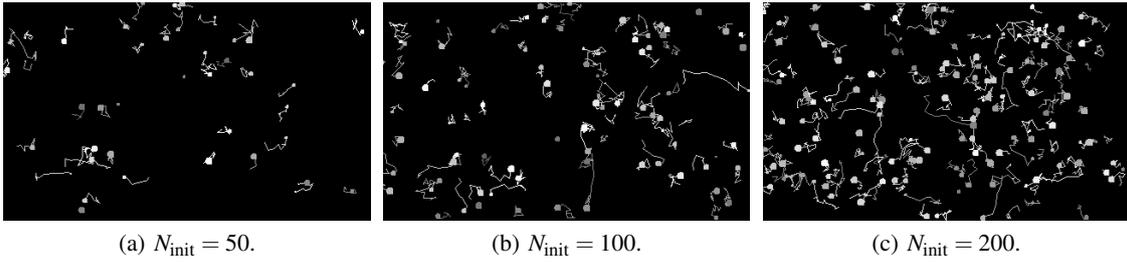
Figure 1: Images of the generated sequences at $t = 10$: Observations are drawn as spots together with their current trajectories.

$q_{\text{exist}} \leftarrow q(c_m^t = n_{\text{exist}} | \tilde{c}_{1:m-1}, M, N)$ // Eq. (14)
$q_{\text{newborn}} \leftarrow q(c_m^t = n_{\text{new}} | \tilde{c}_{1:m-1}, M, N)$ // Eq. (15)

// compute probabilities for association to clutter
$\hat{p}(c_m^t = 0) \leftarrow q_{\text{clutter}} \cdot p_{\text{clutter}}(\mathbf{z}_m^t)$

// ... to each target not associated in $\tilde{c}_{1:m-1}$
for all $n \in \mathfrak{N}_-^t \backslash \tilde{c}_{1:m-1}$ **do**
     $\hat{p}(c_m^t = n) \leftarrow q_{\text{exist}}$
           $\times p(\mathbf{z}_m^t | c_m^t = n, c_{1:m-1}^t, \mathbf{Z}^{1:t} \backslash \mathbf{z}_m^t, \mathbf{C}^{1:t-1})$
end for

// ... to newborn target
$\hat{p}(c_m^t = n_{\text{new}}) \leftarrow q_{\text{newborn}} \cdot p_{\text{newborn}}(\mathbf{z}_m^t)$

normalize $\hat{p}(c_m^t) \rightarrow p(c_m^t)$

// sample association $\tilde{c}_m^t$
$\tilde{c}_m^t \sim p(c_m^t)$

$\tilde{\mathbf{C}} \leftarrow (\tilde{\mathbf{C}}, \tilde{c}_m^t)$
$p_{\tilde{\mathbf{C}}} \leftarrow p_{\tilde{\mathbf{C}}} \cdot p(\tilde{c}_m^t)$

// Part II: Update variables and arrays

// update d for $m \leftarrow m+1$
$d(m+1, k_{m-1}; M, N) \leftarrow$ Eq. (16)
if $\tilde{c}_m^t = 0$ **then**
     // update h for $m \leftarrow m+1$
     $h(m+1, r, k_{m-1} + b_{m-1}; M) \leftarrow$ Eq. (19)
     $h(m+1, r, k_{m-1} + b_{m-1} + 1; M) \leftarrow$ Eq. (19)
else if $\tilde{c}_m^t \in \mathfrak{N}_-^t$ **then**
     // update f and d for $k_{m-1} \leftarrow k_{m-1} + 1$
     // $f(k, k_{m-1} + 1; M, N)$ remains unaffected
     $f(k, k_{m-1} + 2; M, N)$ according to (17)
     $d(m+1, k_{m-1} + 1; M, N) \leftarrow$ according to (16)
     $k_{m-1} \leftarrow k_{m-1} + 1$
else
     // update g for $b_{m-1} \leftarrow b_{m-1} + 1$
     // $g(b, b_{m-1} + 1)$ remains unaffected
     $g(b, b_{m-1} + 2)$ according to (18)
     $b_{m-1} \leftarrow b_{m-1} + 1$
end if
end for
return $(\tilde{\mathbf{C}}, p_{\tilde{\mathbf{C}}})$

5 RESULTS

Model Specification. The state of a target comprises its current position (x^t, y^t) and size, defined as the squareroot of the occupied area, as well as

its position in the previous time step, i.e. $\mathbf{x}^t := (x^t, y^t, x^{t-1}, y^{t-1}, \sqrt{\text{area}^t})^T$. Observations comprise position and size only, i.e. $\mathbf{z}^t := (x^t, y^t, \sqrt{\text{area}^t})^T$. The observation of a target is modeled to obtain the current location and size of the target perturbed with additive Gaussian noise. The measurement noise covariance matrix R is defined as a diagonal matrix with variances for coordinates and size. The dynamic behavior of targets is modeled by switching linear models with Gaussian noise. The corresponding noise covariance matrices Q again are defined as diagonal matrix with the variances for current and previous coordinates as well as for size. Two dynamic models are employed: A random walk model (*rw*) accounts for stopping and change of direction by no deterministic change in position and size of a target, but Gaussian noise only. A second model for directional motion (*fle*) extrapolates the next position from the current and last position with additional Gaussian noise. See (Genovesio and Olivo-Marin, 2008) for more details. Switching of dynamic models between time steps is described by the probabilities $P_{rw \rightarrow rw}$ to stay in the *rw*-model and $P_{rw \rightarrow fle}$, $P_{fle \rightarrow rw}$ and $P_{fle \rightarrow fle}$ defined analogously. We integrate the IMM filter (Bar-Shalom and Blair, 2000) into the RBMCDA framework, which estimates the state distribution of a target with changing dynamics by a Gaussian mixture.

Poisson distributions $\nu(b; \lambda_B)$ and $\mu(u; \lambda_c)$ account for the number of observations from newborn targets and clutter observations respectively, controlled by parameters λ_B and λ_c. The death of targets follows an exponential distribution of the time to last association with parameter λ_d. The proposed model is generative and thus may be used to generate synthetic data.

Synthetic Data Generation. Three time sequences of 50 time steps were generated for evaluation. Each sequence was generated for a different number of initial targets in an image of fixed size to investigate the behavior of the algorithm regarding different numbers of targets and different spatial density of observations. The sequences are referred to by the number of initial

Table 1: Parameters used for data generation and tracking.

	$N_{\text{init}} = 50$	$N_{\text{init}} = 100$	$N_{\text{init}} = 200$
λ_B	0.6	1.4	2.25
Q		$\begin{pmatrix} 5 & 0 & 0 & 0 & 0 \\ 0 & 5 & 0 & 0 & 0 \\ 0 & 0 & 1.67 & 0 & 0 \\ 0 & 0 & 0 & 1.67 & 0 \\ 0 & 0 & 0 & 0 & 1.1 \end{pmatrix}$	
λ_c	λ_d	P_D	R
5	0.5	0.97	$\begin{pmatrix} 5 & 0 & 0 \\ 0 & 5 & 0 \\ 0 & 0 & 1.1 \end{pmatrix}$
$P_{rw \to rw}$	$P_{rw \to fle}$	$P_{fle \to rw}$	$P_{fle \to fle}$
0.7	0.3	0.5	0.5

targets $N_{\text{init}} \in \{50, 100, 200\}$. Fig. 1 shows images of the different sequences at $t = 10$ to give an impression of how densely observations are situated. All parameters are fixed for generation of the three different sequences except for λ_B, which was adjusted for each sequence to achieve a nearly balanced number of deaths and births of targets. The parameters for data generation are presented in Tab. 1. True associations of all observations in this synthetic data are known and used for evaluation of tracking results. Tracking was performed with a number of RBMCDA samples $S \in \{10, 100, 1000\}$ for each sequence using the corresponding parameters from data generation.

Performance Evaluation. Evaluation of RBMCDA tracking performance is not straightforward as target IDs of different RBMCDA samples do not relate to each other nor to IDs of ground-truth data. Thus trajectories cannot directly be related to each other using target IDs or compared to ground-truth. For the same reason state distributions estimated within the RBMCDA samples cannot easily be combined. We therefore consider *track graphs* to evaluate tracking results. The nodes of a track graph are constituted by all observations of a time sequence. Edges are defined as follows. For each target we consider all observations associated to this target. These observations are ordered with respect to time and neighboring observations in this sequence are connected by an edge. As a consequence all observations associated to clutter are not incident to any edge. In the same way the ground-truth track graph is constructed using true associations available from synthetic data. The graph of each RBMCDA sample is compared to this ground-truth graph as follows. Edges that are common to both graphs are considered as *true positives* (TP), edges that appear only in the ground-truth graph are counted as *false negatives* (FN) and edges only in the graph from the current sample are regarded as *false positives* (FP). An example is given in Fig. 2, where in the ground-truth track graph observations are horizontally connected,
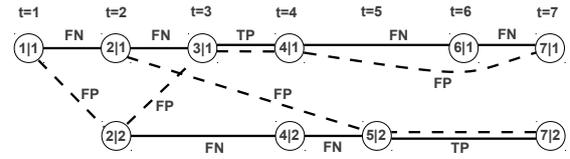
Figure 2: Evaluation of track graphs: nodes represent observations, solid edges ground-truth trajectories, dashed edges tracking results. TP are edges present in both graphs, FN only in ground-truth and FP only in tracking results.

(a) Precision $\mathcal{P}$.   (b) Recall $\mathcal{R}$.

Figure 3: Tracking performance for synthetic data N_{init} tracked with different number of RBMCDA samples S.

but observations 2|1 and 2|2 are interchanged in the track graph of the RBMCDA sample. Furthermore observations 4|2 and 6|1 are wrongly associated to clutter. To evaluate the performance of tracking we use *precision* $\mathcal{P} := \frac{\text{TP}}{\text{TP}+\text{FP}}$ and *recall* $\mathcal{R} := \frac{\text{TP}}{\text{TP}+\text{FN}}$. Note that the track graph is only assessed locally with respect to time, which may underestimate the performance. E.g., in Fig. 2 observation 6|1 is associated to clutter by the tracker resulting in two FN and one FP, while the association of 4|1 to 7|1 is obviously correct. Still we prefer this measure for its clear definition.

The precision $\mathcal{P}$ of tracking results in the different experiments is presented in Fig. 3(a) and recall $\mathcal{R}$ in Fig. 3(b). Both measures increase for any of the sequences, when the number of RBMCDA samples is increased. On the other hand $\mathcal{P}$ and $\mathcal{R}$ decrease, when more targets are presents in the domain.

Computation Times. Tracking was conducted on AMD Opteron 848 processors with clock speed 2.2 GHz using a Java implementation on Sun's 64-Bit virtual machine. Probabilities are represented by their logarithms to avoid numerical problems. Their summation thus is expensive and required extensively by Eqs. (13), (14) and (15). Computation times are measured in UNIX' *user CPU time* and presented in Tab. 2 for the different sequences and sample sizes.

393

Table 2: Computation time (seconds) of tracking sequences N_{init} with different number of RBMCDA samples S.

	$S = 10$	$S = 100$	$S = 1000$
$N_{\text{init}} = 50$	21.1	167.3	2,083.1
$N_{\text{init}} = 100$	89.8	886.3	11,021.9
$N_{\text{init}} = 200$	384.8	3,714.1	38,478.0

6 DISCUSSION

The proposed method is evaluated for tracking performance in order to prove its concept. Sequence $N_{\text{init}} = 50$ was generated to pose a tracking problem of relatively low complexity (see Fig. 1(a)). Targets are not located densely and noise in position and size is reasonable. The performance is given in Fig. 2 for varying sample size S. With regard to the pessimistic characteristics of the employed measures, the resulting performance indicates expectedly good performance for tracking with $S = 10$ RBMCDA samples. In general $\mathcal{R}$ is smaller than $\mathcal{P}$ because, e.g., wrong associations to clutter have no impact on $\mathcal{P}$, however on $\mathcal{R}$. If the number of RBMCDA samples is increased performance improves with maximum precision of 0.961 and recall of 0.886 for $S = 1000$.

The tracking problem gets more difficult as the number of targets increases in sequence $N_{\text{init}} = 100$ (see Fig. 1(b)). Wrong associations become more likely when targets and their observations approach and thus both $\mathcal{P}$ and $\mathcal{R}$ decrease in general, but are still very satisfactory (see Figs. 3(a) and 3(b)). Again, the performance improves with sample size S. Maximum precision is 0.93 with a recall of 0.899.

Even more targets are generated in sequence $N_{\text{init}} = 200$, posing a very challenging tracking problem (see Fig. 1(c)). Here, a maximum precision of 0.841 is achieved where the maximum recall is 0.807. This is well acceptable taking the complexity of the data into account. It could be speculated that the performance may be further increased using a larger set of RBMCDA samples.

In summary, tracking performance shows that the proposed method is able to successfully track large numbers of targets in demanding data.

For practical application we recommend to represent probabilities by their logarithms to track larger numbers of targets, because sampling would fail as the small probabilities in Eq. (5) cannot be represented by double precision any more. Calculation of the data association prior becomes more expensive due to the summation of probabilities, but the numerical stability allows to successfully track at least 3600 targets. This is in contrast to about 70 targets for a conventional representation of probabilities.

7 CONCLUSIONS

In this work we propose a parametric data association prior for use with RBMCDA presented in (Särkkä et al., 2007) to track a varying number of targets. This prior models the formation of observations from existing and newborn targets as well as clutter observations. We developed an efficient algorithm to sample associations using this prior. In a proof of concept we show that this sampling procedure allows to successfully sample associations in the presence of hundreds of targets. Computation times are moderate using a sample size of 1000 and as much as 200 targets on average. For synthetic data of demanding complexity the performance of sampled associations is well acceptable. We integrated the IMM filter to enable tracking of targets with changing dynamics for an application to microscopy image analysis in mind.

REFERENCES

Bar-Shalom, Y. and Blair, W. (2000). *Multitarget-multisensor tracking: Applications and advances. Vol III*. Artech House.

Bar-Shalom, Y. and Fortmann, T. (1987). *Tracking and data association*. Academic Press Professional.

Blom, H. and Bloem, E. (2003). Joint IMMPDA particle filter. In *Information Fusion, Proc 6th International Conference*.

Doucet, A., de Freitas, N., and Gordon, N. (2001). An introduction to sequential monte carlo methods. *Sequential Monte Carlo methods in practice*, pages 3–14.

Genovesio, A. and Olivo-Marin, J. (2008). Particle tracking in 3D+t biological imaging. In Rittscher, J., editor, *Microscopic Image Analysis for Life Science Applications*, pages 223–282. Artech House.

Hue, C., Le Cadre, J., and Pérez, P. (2002). Sequential monte carlo methods for multiple target tracking and data fusion. *Signal Processing, IEEE Trans*, 50(2):309–325.

Khan, Z., Balch, T., and Dellaert, F. (2005). MCMC-based particle filtering for tracking a variable number of interacting targets. *Pattern Analysis and Machine Intelligence, IEEE Trans*, 27:2005.

Musicki, D. and Evans, R. (2002). Joint integrated probabilistic data association-JIPDA. In *Information Fusion, 2002. Proc 5th International Conference*, volume 2, pages 1120–1125. IEEE.

Reid, D. (1979). An algorithm for tracking multiple targets. *Automatic Control, IEEE Trans*, 24(6):843–854.

Särkkä, S., Vehtari, A., and Lampinen, J. (2007). Rao-blackwellized particle filter for multiple target tracking. *Information Fusion*, 8(1):2–15.

Smal, I. et al. (2008). Multiple object tracking in molecular bioimaging by rao-blackwellized marginal particle filtering. *Medical image analysis*, 12(6):764–777.