TRAJECTORY RETRIEVAL FROM MONTE CARLO DATA ASSOCIATION SAMPLES FOR TRACKING IN FLUORESCENCE MICROSCOPY IMAGES

Oliver Greß, Stefan Posch

Institute of Computer Science, Martin Luther University Halle-Wittenberg
Von-Seckendorff-Platz 1, 06120 Halle (Saale), Germany

ABSTRACT

The dynamic behavior of sub-cellular structures is of great interest to bio-medical research. We propose to use probabilistic multi-target tracking to analyze the dynamics of stress granules (SGs) to alleviate the effort of manual analysis. Inherent to multi-target tracking approaches is the combinatorial problem to associate observations to underlying targets. Rao-Blackwellized Monte Carlo Data Association circumvents this problem by sampling in the space of associations. As each sample provides its own hypothesis of SG trajectories, we employed a graph partitioning algorithm to extract one single set of trajectories. This is shown to outperform the sample with maximum probability on both synthetic data and fluorescence microscopy images.

Index Terms— Probabilistic Tracking, Multiple Targets, Monte Carlo, Data Association, Fluorescence Microscopy

1. INTRODUCTION

The understanding of cellular processes is a central question in systems biology where the behavior of cellular structures under varying conditions is one important issue. We consider the aggregation of stress granules (SGs), which are observed in the cytoplasm of eukaryotic cells in response to environmental stress [1]. Their dynamic characteristics like accumulation, motion and interaction offer valuable information.

SGs appear as spot-like structures of different and changing size in fluorescence microscopy images (see Fig. 2). Tracking SGs is difficult, as their appearance contains few clues to identify individual SGs. Furthermore, SGs suddenly start or stop moving and change their direction. Thus modeling dynamics may be beneficial for tracking SGs, as is the basis of probabilistic tracking methods. Here the state distribution of targets, SGs in our case, is predicted in each time step using a dynamic model and previous estimates. This estimate is corrected by an observation of the target. The major issue in multi-target tracking are the combinatorics of associating observations in each image to underlying targets.

Rao-Blackwellized Monte Carlo Data Association (RBMCDA) proposed in [2] avoids this combinatorial problem by sampling in the space of associations while keeping state distributions in closed form. We propose to use RBMCDA for tracking SGs due to its data association approach and the ability to integrate Interacting Multiple Models (IMM) filtering [3] required by the changing dynamics of SGs.

The main contributions of this work are the successful application of RBMCDA for tracking SGs and the proposal of a method to extract a single set of trajectories from all RBMCDA samples. This set outperforms the obvious choice of a single RBMCDA sample with maximum probability on synthetic as well as on real microscopy data.

2. RELATED WORK

Multi-Hypothesis Tracking (MHT) [4] provides an optimal solution to multi-target tracking for linear systems with Gaussian noise in theory. It accounts for the birth and death of targets as well as observations from noise (*clutter*). MHT maintains a tree of hypotheses for all possible associations, which results in an exponentially growing number of hypotheses that must be reduced heuristically for practical applications.

The Joint Probabilistic Data Association (JPDA) filter [3] updates the state distribution of each target using nearby observations assuming exclusive association to this target. The resulting Gaussian mixture is approximated by a single Gaussian to prevent exponential growth. JPDA provides robust estimates in case of clutter. Extensions combine it with IMM [3] or overcome its constraint on a fixed number of targets [5].

Sequential Monte Carlo (SMC) methods approximate state distributions by a set of samples [6]. SMC methods can avoid the combinatorial problem in multi-target tracking by sampling association variables. In [7] the joint distribution of both target states and association variables is approximated by samples requiring a large number due to high dimensionality. For this reason [2] introduce Rao-Blackwellization of state variables, i.e. their distributions are modeled analytically.

Tracking techniques were successfully applied to microscopy images. In [8] SMC state samples are re-clustered to reconsider their affiliation to targets. Here the state distribution of targets is maintained in analytical form only for their intensity variable, in contrast to [2] where the whole state distribution is maintained in closed form. In [9] a suboptimal matching strategy decides for associations.

3. METHODS

3.1. Rao-Blackwellized Monte Carlo Data Association

At each time step, probabilistic multi-target tracking updates the state distribution of each target with the corresponding observation. For this, each observation needs to be associated to one target or to clutter. To consider all possible combinations of associations is not feasible due to combinatorics. The other extreme, to decide on one single of these associations will provoke errors which accumulate with time. RBMCDA presented in [2] provides a compromise. Here a set of RBMCDA samples (R-samples) is maintained where each R-sample comprises for each observation one association to either target or to clutter. This association is sampled according to an underlying distribution independently in each R-sample. Furthermore, each R-sample contains – in analytical form – a state distributions for each target which depends on the sampled association of the R-sample and all observations. Thus, the set of R-samples can be understood as approximating the joint distribution of target states and association variables given all observations. Furthermore, the associations of one R-sample convey the trajectories of the detected spots by connecting spots of the same target in the time sequence.

3.2. Statistical Models

For RBMCDA based tracking, statistical models of the dynamic process and the observation of targets are required as for any filtering approach. In addition, a model for data association is needed to account for birth, death and detection of targets, and for clutter observations. We use the model given in [10]: A binomial distribution models the number of observations from existing targets with probability P_D to detect a target. The numbers of observations from newborn targets and clutter are both Poisson distributed with parameter λ_b and λ_c, respectively. Targets die according to an exponential distribution of the time to last association with parameter λ_d.

Dynamic modeling needs to account for the changing dynamics of SGs. Thus we integrate the IMM filter [3, 9], which allows us to switch between two dynamic models: A random walk (*rw*) model and a model for directional motion (*fle*), which extrapolates the next position of a SG linearly from its last and current position. The hidden state of a target is denoted by $\mathbf{x}^t := (x^t, y^t, x^{t-1}, y^{t-1}, a^t)^T$. (x^t, y^t) and (x^{t-1}, y^{t-1}) are the current resp. previous position of the SG's centroid, and a^t the square-root of the SG's size at time t. An observation comprises the centroid of the SG and the square-root of its size. It is denoted by $\mathbf{z} := (x^t, y^t, a^t)^T$. The change of an SG's state from time $t-1$ to t is a linear system with $\mathbf{F}_{rw}$ and $\mathbf{F}_{fle}$ for the *rw* and *fle* model respectively:

$$\mathbf{x}^t = \mathbf{F}\mathbf{x}^{t-1} + \mathbf{q}^t, \text{ where} \tag{1}$$

$$\mathbf{F}_{rw} := \begin{pmatrix} 1&0&0&0&0\\ 0&1&0&0&0\\ 0&0&1&0&0\\ 0&0&0&1&0\\ 0&0&0&0&1 \end{pmatrix}, \mathbf{F}_{fle} := \begin{pmatrix} 2&0&-1&0&0\\ 0&2&0&-1&0\\ 1&0&0&0&0\\ 0&1&0&0&0\\ 0&0&0&0&1 \end{pmatrix} \tag{2}$$

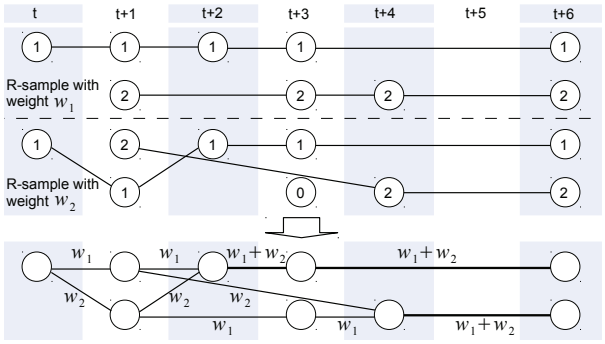


Fig. 1. *Top:* Trajectory hypotheses of two R-samples presented as graphs. Nodes correspond to observations, enclosed numbers to associated target ID (> 0) or clutter ($= 0$). *Bottom:* J-graph resulting from the above R-samples.

and $\mathbf{q}^t \sim \mathcal{N}(\mathbf{0}, \mathbf{Q})$ models additive Gaussian process noise.

The observation model assumes the observation to be the target's state with mean-free Gaussian noise $\mathbf{r}^t \sim \mathcal{N}(\mathbf{0}, \mathbf{R})$:

$$\mathbf{z}^t = \mathbf{H}\mathbf{x}^t + \mathbf{r}^t, \text{ where } \mathbf{H} := \begin{pmatrix} 1&0&0&0&0\\ 0&1&0&0&0\\ 0&0&0&0&1 \end{pmatrix} \tag{3}$$

See [9] for details on these dynamic and observation models.

3.3. Retrieval of Trajectories from R-samples

As stated in Subsec. 3.1, each R-sample contains one hypothesis for the trajectories of all spots detected. However, in many applications one single interpretation of the data is of interest. An obvious option is to choose the R-sample of maximum probability, but it does not represent the complete distribution in contrast to the ensemble of R-samples. Merging of all R-samples cannot easily be done as target IDs are not shared between the associations of different R-samples. Only the identities of the observations, i.e. spots, are shared between R-samples. This is exploited to construct a joint graph of trajectories (J-graph) from tracking results of all R-samples that is used to extract one joint set of trajectories.

The J-graph contains all observations as nodes and is multipartite with respect to time points. The trajectories of each R-sample connect nodes of this graph from different partitions (Fig. 1 top). Edges hypothesized by one or more R-samples are merged in the J-graph and weighted by the sum of weights of R-samples involved (Fig. 1 bottom).

Extracting trajectories from the J-graph may now be formulated as an assignment problem, where a set of nodes from different partitions, i.e. observations from different time points, are to be assigned to stem from the same target or to clutter. Nodes of the same target induce a subgraph, and as each node can be assigned exactly once, these subgraphs are disjoint. We aim at extracting subgraphs such that their sum of edge weights is maximal. This is a generalization of the assignment problem for bipartite graphs solved by the

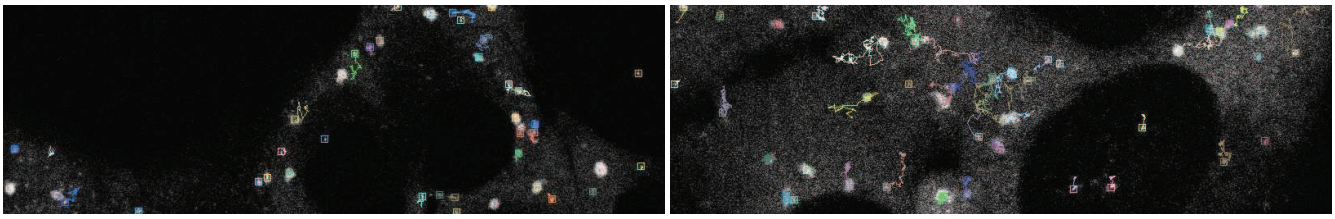


Fig. 2. Trajectories of SGs extracted with *greedyGourmet*. SGs in sequence 1 (left) move less than in sequence 2 (right).

Hungarian algorithm [11] but no efficient solution for the multipartite case is known.

Therefore we employ the *greedyGourmet* algorithm proposed in [12] in the context of mass spectrometry to group mass peaks. This problem is cast as an assignment problem for multipartite graphs. The idea of the algorithm is to consider each node v_t of a partition t and find its best match $v_{\tilde{t}}$ in subsequent partitions $\tilde{t} > t$. The best match is the node of partition $\tilde{t}$ connected to v_t by an edge with maximal weight. If in turn v_t is also the best match of $v_{\tilde{t}}$, the corresponding edge and node are added to the subgraph with all edges from the new node to other nodes of the subgraph. If the node to be added is already member of another subgraph or the current subgraph contains another node of the new node's partition, the new node is only added if the total sum of weights increases by this operation. The algorithm assigns a unique target ID to all observations of each subgraph constructed.

greedyGourmet connects nodes by weak edges rather than leaving them unconnected. Thus, e.g., an observation which was associated to a target by only one R-sample would still be assigned to this target, however with little support. Therefore, we prune an edge before applying *greedyGourmet* if its weight is below a threshold θ. This threshold is the sum of the $(N - \lfloor N_{\text{eff}} \rfloor)$ lowest R-sample weights where N is the number of R-samples and N_{eff} is the effective sample size computed from R-sample probabilities according to [9].

4. RESULTS

The proposed method is evaluated on two image sequences acquired by confocal laser scanning microscopy. Both sequences were reduced from 3D-volumes to 2D by averaging along the z-dimension to alleviate manual labeling. Both data contain a similar number of SGs ($\approx$50), which move faster in sequence 2 (see Fig. 2). SGs were detected using [13] and manually labeled with target IDs. In addition 56 synthetic sequences were created using the generative tracking model for a variety of parameters and tracked with fixed parameters to examine the proposed approach when the model deviates from the data tracked. The tracker's behavior is further discussed in [10]. In all experiments 100 R-samples were used.

Evaluation is based on the trajectories where two observations of - not necessarily directly - consecutive time points are connected by an edge if observations originate from the same target. Comparing to ground-truth edges common to both graphs are considered as *true positives* (TP), edges only in the ground-truth graph as *false negatives* (FN) and edges only in the tracking graph as *false positives* (FP). Precision $\mathcal{P} := \frac{\text{TP}}{\text{TP+FP}}$ and recall $\mathcal{R} := \frac{\text{TP}}{\text{TP+FN}}$ are computed to characterize the tracking performance. See [10] for more details.

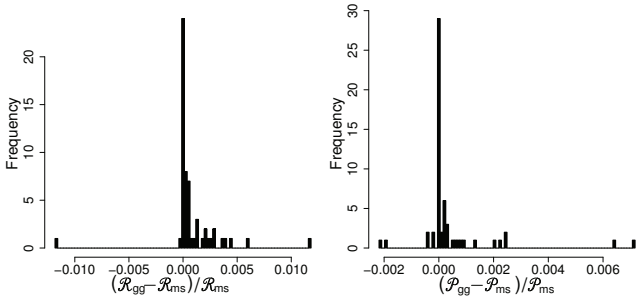


Fig. 3. Histograms of the relative difference of quality measures of *greedyGourmet* and maximum-weight R-sample in synthetic data experiments. Left: $\frac{\mathcal{R}_{gg}\text{-}\mathcal{R}_{ms}}{\mathcal{R}_{ms}}$. Right: $\frac{\mathcal{P}_{gg}\text{-}\mathcal{P}_{ms}}{\mathcal{P}_{ms}}$.

We compare $\mathcal{R}$ and $\mathcal{P}$ of *greedyGourmet* trajectories (denoted by $\mathcal{R}_{gg}$ and $\mathcal{P}_{gg}$) with the quality of trajectories from R-samples, especially from the maximum weight R-sample (called $\mathcal{R}_{ms}$ and $\mathcal{P}_{ms}$), the obvious choice (see Subsec. 3.3).

In 32 of the 56 experiments with synthetic data $\mathcal{R}_{gg} > \mathcal{R}_{ms}$, in 24 experiments $\mathcal{R}_{gg} = \mathcal{R}_{ms}$ and only in 2 experiments $\mathcal{R}_{gg} < \mathcal{R}_{ms}$. Identical $\mathcal{R}$ is observed if a large subset of R-samples shares the maximum weight and their associations represent the same trajectories. This happens if the set of R-samples is resampled to maintain an acceptable effective sample size N_{eff} during tracking [9]. Here *greedyGourmet* extracts the same trajectories as the R-samples with maximum weight. The findings for $\mathcal{P}$ are similar. In 30 of the 56 experiments $\mathcal{P}_{gg} > \mathcal{P}_{ms}$, in 22 experiments their values equal and in 6 experiments $\mathcal{P}_{gg} < \mathcal{P}_{ms}$. Fig. 3 illustrates these findings by histograms of the relative difference of $\mathcal{R}$ and $\mathcal{P}$ of *greedyGourmet* and the maximum-weight sample, i.e. $\frac{\mathcal{R}_{gg}\text{-}\mathcal{R}_{ms}}{\mathcal{R}_{ms}}$ and $\frac{\mathcal{P}_{gg}\text{-}\mathcal{P}_{ms}}{\mathcal{P}_{ms}}$, for all these experiments. Note that some of the positive differences coincide with the 0-bin. In summary these experiments advise to rather use trajectories from *greedyGourmet* than from the R-sample with highest weight.

For real data the weight (abscissa) and $\mathcal{R}/\mathcal{P}$ (ordinate) of

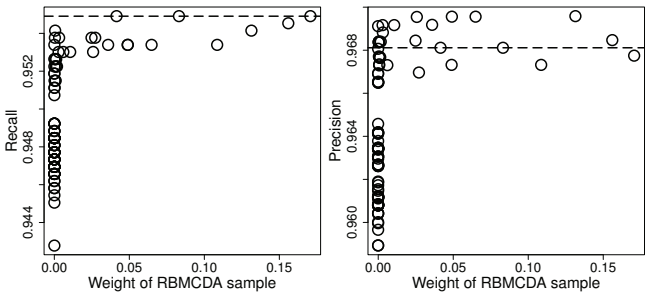


Fig. 4. Tracking performance of *greedyGourmet* (dashed line) vs. the individual weighted R-samples (○) for sequence 1.

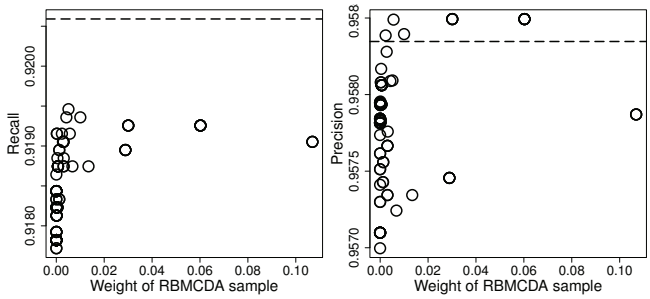


Fig. 5. Tracking performance of *greedyGourmet* (dashed line) vs. the individual weighted R-samples (○) for sequence 2.

each R-sample is plotted in Figs. 4 and 5. The dashed line shows the performance of *greedyGourmet*. Note that weights of the leftmost samples are very small, yet not zero. The plots show that the maximum weight R-sample, i.e. the rightmost circle, does not give the best performance among all R-samples. This again suggests not to choose this maximum weight R-sample as the only result of tracking. In the following values are rounded to 4 decimal places. For sequence 1 $\mathcal{R}_{\text{gg}}$=$\mathcal{R}_{\text{ms}}$=0.9549 is equal to the maximum $\mathcal{R}$ achieved by all R-samples. $\mathcal{P}_{\text{gg}}$=0.9681 is below the maximum among all R-samples, however exceeds $\mathcal{P}_{\text{ms}}$=0.9677. Tracking becomes harder in sequence 2. Here *greedyGourmet* with $\mathcal{R}_{\text{gg}}$=0.9206 exceeds not only $\mathcal{R}_{\text{ms}}$=0.9191 but also $\mathcal{R}$ of all R-samples. When $\mathcal{P}$ is regarded, $\mathcal{P}_{\text{gg}}$=0.9584 is exceeded by few R-samples but is clearly superior to $\mathcal{P}_{\text{ms}}$=0.9579.

The employed method achieves successful tracking of SGs in both sequences. Furthermore, the findings on synthetic data are confirmed by the real data, i.e. extracting one single set of trajectories using *greedyGourmet* gives superior performance compared to the maximum weight choice.

5. CONCLUSION

In this work we demonstrate the application of RBMCDA for tracking SGs in image data from fluorescence microscopy. We propose to use *greedyGourmet* to extract a single set of trajectories from trajectory hypotheses represented by all R-samples. The quality of trajectories from *greedyGourmet* outperforms trajectories from the highest weighted R-sample for synthetic data as well as for real microscopy images.

6. REFERENCES

[1] N. Kedersha et al., "Stress granules and processing bodies are dynamically linked sites of mRNP remodeling," *Cell Biology*, vol. 169, no. 6, pp. 871–884, 2005.

[2] S. Särkkä et al., "Rao-Blackwellized particle filter for multiple target tracking," *Information Fusion*, vol. 8, no. 1, pp. 2–15, 2007.

[3] Y. Bar-Shalom and W.D. Blair, *Multitarget-multisensor tracking: Applications and advances. Vol III*, Artech House, 2000.

[4] D. Reid, "An algorithm for tracking multiple targets," *IEEE Trans. Autom. Control*, vol. 24, no. 6, pp. 843–854, 1979.

[5] D. Musicki and R. Evans, "Joint integrated probabilistic data association-JIPDA," in *5th Int. Conf. on Information Fusion*, 2002, vol. 2, pp. 1120–1125.

[6] A. Doucet et al., *Sequential Monte Carlo methods in practice*, Springer, 2001.

[7] C. Hue et al., "Sequential Monte Carlo methods for multiple target tracking and data fusion," *IEEE Trans. Signal Process.*, vol. 50, no. 2, pp. 309–325, 2002.

[8] I. Smal et al., "Multiple object tracking in molecular bioimaging by Rao-Blackwellized marginal particle filtering," *Medical Image Analysis*, vol. 12, no. 6, pp. 764–777, 2008.

[9] A. Genovesio and J.C. Olivo-Marin, "Particle tracking in 3D+t biological imaging," in *Microscopic Image Analysis for Life Science Applications*, J. Rittscher et al., Eds., pp. 223–282. Artech House, 2008.

[10] O. Greß et al., "Model dependency of RBMCDA for tracking multiple targets in fluorescence microscopy," in *Bildverarbeitung für die Medizin*, 2012, accepted.

[11] H.W. Kuhn, "The hungarian method for the assignment problem," *Naval research logistics quarterly*, vol. 2, no. 1-2, pp. 83–97, 1955.

[12] J. Kutzera, "Gruppierung von LC/MS-Pseudospektren aus multiplen Messungen," Diploma thesis, Martin Luther University Halle-Wittenberg, 2009.

[13] O. Greß et al., "Scale-adaptive wavelet-based particle detection in microscopy images," in *Bildverarbeitung für die Medizin*, H.-P. Meinzer et al., Eds. 2010, pp. 266–270, Springer.