An Uncertainty-Aware Performance Measure for Multi-Object Tracking

Downloaded from: <https://research.chalmers.se>, 2023-04-26 08:50 UTC

Citation for the original published paper (version of record):

Pinto, J., Xia, Y., Svensson, L. et al (2021). An Uncertainty-Aware Performance Measure for Multi-Object Tracking. IEEE Signal Processing Letters, 28: 1689-1693.
<http://dx.doi.org/10.1109/LSP.2021.3103488>

N.B. When citing this work, cite the original published paper.

©2021 IEEE. Personal use of this material is permitted.

However, permission to reprint/republish this material for advertising or promotional purposes

An Uncertainty-Aware Performance Measure for Multi-Object Tracking

Juliano Pinto, Yuxuan Xia, *Student Member, IEEE*, Lennart Svensson, *Senior Member, IEEE*, and Henk Wymeersch, *Senior Member, IEEE*

Abstract—Evaluating the performance of multi-object tracking (MOT) methods is not straightforward, and existing performance measures fail to consider all the available uncertainty information in the MOT context. This can lead practitioners to select models which produce uncertainty estimates of lower quality, negatively impacting any downstream systems that rely on them. Additionally, most MOT performance measures have hyperparameters, which makes comparisons of different trackers less straightforward. We propose the use of the negative log-likelihood (NLL) of the multi-object posterior given the set of ground-truth objects as a performance measure. This measure takes into account all available uncertainty information in a sound mathematical manner without hyperparameters. We provide efficient algorithms for approximating the computation of the NLL for several common MOT algorithms, show that in some cases it decomposes and approximates the widely-used GOSPA metric, and provide several illustrative examples highlighting the advantages of the NLL in comparison to other MOT performance measures.

Index Terms—Multitarget tracking, Multi-object tracking, Performance measure, Uncertainty evaluation.

I. INTRODUCTION

Multi-object tracking (MOT) is the task of tracking an unknown number of objects through time using noisy measurements, with important applications in various areas [1]–[5]. The main challenge for MOT is the unknown correspondence between objects and measurements, which makes it necessary for the algorithms to infer such information [6]. Additionally, in many applications it is important that tracking systems provide accurate uncertainty estimates of their outputs, so that decision-making systems can take robust actions [7]. Evaluating the quality of MOT methods is also challenging, due to the lack of knowledge about the correct correspondence between the ground-truth set of object states and the tracker’s state estimates and the associated uncertainties [8].

Several performance measures have been proposed for MOT, including the Hausdorff metric [9], the Wasserstein metric [9], multi-object tracking accuracy and precision (MOTA and MOTP) [10], higher-order tracking accuracy (HOTA) [11], optimal subpattern assignment (OSPA) [12], and generalized optimal subpattern assignment (GOSPA) [13]. MOTA and MOTP are prevalent in computer vision tasks, such as multiple people tracking [14], while OSPA and GOSPA are the most widely applied for general MOT problems. These performance measures cope with the unknown associations between estimates and objects by relying on minimum-cost associations

based on a user-defined distance measure. None of these methods can assess the quality of the uncertainty estimates, and existing measures such as normalized estimation error squared [15] are not easily applicable to MOT due to the unknown data associations.

To address this, there have been efforts to extend the MOT performance measures to incorporate uncertainty information, but they only evaluate some of the available uncertainty [16], [17], and/or rely on proxies for ground-truth state uncertainties which are not applicable to general MOT applications [17]. To the best of our knowledge, no general, mathematically sound incorporation of all the uncertainties in MOT has been proposed in a performance measure. Additionally, most MOT performance measures have hyperparameters (e.g., distance metric and thresholds). Such choices are often non-trivial, difficult to generalize to new contexts, and make comparisons between different trackers less straightforward. Moreover, MOT methods themselves may have hyperparameters, optimized for a certain metric, leading to a chicken-and-egg-problem.

In this letter, we introduce a new MOT performance measure that incorporates all uncertainties in a sound mathematical manner, while at the same time having zero hyperparameters: the negative log-likelihood (NLL) of the model (i.e., the MOT method posterior), given the ground-truth objects. NLL has been widely applied in statistics [18] and optimization [19], but not yet in MOT. We provide the following contributions. First, we propose the use of the NLL as a new MOT performance measure, and provide efficient algorithms for computing/approximating it for a number of common MOT algorithms based on random finite sets (RFSs): (C)PHD [20], [21], PMBM [22], PMB [23], MBM [24], and MBM₀₁ [25]. Second, we show that for the PMB MOT family, the NLL decomposes into separate terms that provide additional transparency into the performance, and that under certain assumptions, NLL is closely related to GOSPA. Lastly, we provide illustrative examples to highlight the advantages of NLL.

Notations: Scalars and vectors are denoted by lowercase or uppercase letters with no special typesetting x , matrices by uppercase boldface letters $\mathbf{X}$, and sets by uppercase blackboard bold letters $\mathbb{X}$. In addition, we define $\mathbb{N}_a = \{i \in \mathbb{N} \mid i \leq a\}$, $a \in \mathbb{N}$.

II. NLL AS A PERFORMANCE MEASURE

In this section, we provide the definition of NLL, along with examples on how to compute it efficiently for different families of MOT densities. Given a ground-truth set of object states $\mathbb{Y} = \{y_1, \dots, y_{|\mathbb{Y}|}\}$ and a posterior density of the tracked objects $f_M(\cdot)$ from method M (i.e., the multi-object posterior,

This work was supported, in part, by a grant from the Chalmers AI Research Centre Consortium (CHAIR).

The authors are with the Department of Electrical Engineering, Chalmers University of Technology, Sweden. (e-mail: juliano@chalmers.se).

which describes the distribution of the set of object states), the NLL is defined as

$$\text{NLL}(\mathbb{Y}, f_M) = -\log f_M(\mathbb{Y}). \quad (1)$$

We can then rank different algorithms (say M_1 and M_2) by comparing $\text{NLL}(\mathbb{Y}, f_{M_1})$ and $\text{NLL}(\mathbb{Y}, f_{M_2})$ for the same set $\mathbb{Y}$, or by computing an expectation with respect to different trials. This type of comparison naturally incorporates the uncertainty information estimates by the trackers, since to score well algorithms must have most of the mass of their posterior in regions where it is likely that the objects in $\mathbb{Y}$ will be, without being overly confident (see Section IV). We now provide examples of several important MOT filters and show how the NLL can be computed.

1) *(C)PHD Filters*: (C)PHD filters [20], [21] have been widely used due to their low complexity, simple implementation, and relatively good performance. The multi-object posterior for a CPHD filter $f_{\text{CPHD}}(\mathbb{X})$ takes the form

$$f_{\text{CPHD}}(\mathbb{X}) = |\mathbb{X}|! p(|\mathbb{X}|) \prod_{x \in \mathbb{X}} s(x), \quad (2)$$

where $p(|\mathbb{X}|)$ is the cardinality distribution (a Poisson distribution for the PHD filter) of $\mathbb{X}$, and $s(x)$ is the single-object state density. Computing $\text{NLL}(\mathbb{Y}, f_{\text{CPHD}})$ according to (1) therefore yields the expression

$$\text{NLL}(\mathbb{Y}, f_{\text{CPHD}}) = -\log(|\mathbb{Y}|!) - \log p(|\mathbb{Y}|) - \sum_{y \in \mathbb{Y}} \log s(y),$$

which can be computed with complexity $\mathcal{O}(|\mathbb{Y}|)$.

2) *PMBM Filters*: Poisson multi-Bernoulli mixture (PMBM) filters [22], [23] are the optimal solution to MOT with standard multi-object dynamic and measurement models with Poisson birth [5, Chap. 13]. The multi-object posterior for a PMBM filter is defined as:

$$f_{\text{PMBM}}(\mathbb{X}) = \sum_{\mathbb{X}^U \cup \mathbb{X}^D = \mathbb{X}} f_{\text{PPP}}(\mathbb{X}^U) f_{\text{MBM}}(\mathbb{X}^D) \quad (3)$$

$$f_{\text{PPP}}(\mathbb{X}) = \exp\left(-\int \lambda(x') dx'\right) \prod_{x \in \mathbb{X}} \lambda(x) \quad (4)$$

$$f_{\text{MBM}}(\mathbb{X}) = \sum_{h=1}^H w_h \sum_{\substack{\mathbb{X}^U = \mathbb{X} \\ \mathbb{X}^D = \mathbb{X}}} \prod_{k=1}^m f_k^h(\mathbb{X}_k) \quad (5)$$

where $\lambda(\cdot)$ is the intensity function of the Poisson point process (PPP), w_h are the weights of each of the H MB components of the MBM ($\sum_h w_h = 1$), m is the number of Bernoulli components in each of the MB components of the MBM (set as identical for each mixture component without loss of generality), and $f_k^h(\mathbb{X}_k)$ is the k -th Bernoulli density of the h -th hypothesis, with

$$f_k^h(\mathbb{X}_k) = \begin{cases} 1 - r_k^h, & \text{if } \mathbb{X}_k = \emptyset \\ r_k^h p_k^h(x), & \text{if } \mathbb{X}_k = \{x\} \\ 0 & \text{otherwise.} \end{cases} \quad (6)$$

where r_k^i is the existence probability, and $p_k^i(\cdot)$ is the single-object density.

When evaluating $f_{\text{PMBM}}(\mathbb{Y})$, (3) can be interpreted as summing the likelihoods of all the possible assignments between the elements of $\mathbb{Y}$ and either the PPP component or one of the Bernoulli components of the PMB density, for each hypothesis h . Since the number of such possible assignments grows super-exponentially in $\mathbb{X}$, computing the NLL has a complexity that

also grows at the same rate in $\mathbb{Y}$. However, among all these assignments, generally only a few contribute significantly to the overall sum, and the likelihood can be approximated by neglecting all other terms (e.g., when the ground-truth objects in $\mathbb{Y}$ are reasonably well-separated and so are the p_k^h for each h).

To find such terms for each hypothesis h , we solve an optimal assignment problem for $\mathbf{A}^h \in \{0, 1\}^{(m+|\mathbb{Y}|) \times |\mathbb{Y}|}$ [6, Chap. 7]:

$$\min_{\mathbf{A}^h} \sum_k \sum_l C_{k,l}^h A_{k,l}^h \quad (7a)$$

$$\text{s.t.} \quad \sum_{k=1}^{m+|\mathbb{Y}|} A_{k,l}^h = 1, \quad \sum_{l=1}^{|\mathbb{Y}|} A_{k,l}^h \leq 1, \quad (7b)$$

where C^h is a cost matrix defined as

$$C_{k,l}^h = \begin{cases} -\log\left(\frac{p_k^h(y_l)}{1-r_k^h} r_k^h\right), & \text{if } k \leq m \\ -\log \lambda(y_l), & \text{if } k = l + m \\ \infty, & \text{otherwise,} \end{cases} \quad (8)$$

and $\mathbf{A}^h$ is the assignment matrix between ground-truth objects and the components of the PMBM. If $[\mathbf{A}^h]_{i,j} = 1$, then y_j is assigned to the i -th component of the PMBM, where $1 \leq i \leq m$ corresponds to the m Bernoulli components in hypothesis h , and all $i > m$ to the PPP component.

Murty's algorithm [26] allows for efficient computation of the Q -lowest cost associations $\mathbf{A}_1^{h,*}, \dots, \mathbf{A}_Q^{h,*}$ to this assignment problem. We find that

$$\text{NLL}(\mathbb{Y}, f_{\text{PMBM}}) \approx \int \lambda(y') dy' \quad (9)$$

$$- \log \left(\sum_{h=1}^H \sum_{q=1}^Q w_h \prod_{y \in \mathbb{Y}^U(\mathbf{A}_q^{h,*})} \lambda(y) \prod_{k=1}^m f_k^h(\mathbb{Y}_k(\mathbf{A}_q^{h,*})) \right)$$

where $\mathbb{Y}_k(\mathbf{A}_q^{h,*}) = \{y_j \in \mathbb{Y} \mid [\mathbf{A}_q^{h,*}]_{k,j} = 1\}$, $\mathbb{Y}^U(\mathbf{A}_q^{h,*}) = \mathbb{Y} \setminus \cup_{i=1}^m \mathbb{Y}_i(\mathbf{A}_q^{h,*})$. The worst-case time complexity of this approximation scales as $\mathcal{O}(HQ(m+|\mathbb{Y}|)^3)$ [26].

3) *Other Common MOT Filters*: The same assignment problem defined in (7) can be used to efficiently compute other special cases of the PMBM density. For instance, the PMB density is a special case of the PMBM with a single MB component ($H = 1$) [5], and the MBM density is a PMBM where $\lambda(x) = 0, \forall x$ [5]. Additionally, the MBM₀₁ density is a special case of the MBM density when existence probabilities p_k^i of all the Bernoulli components are set to either 0 or 1 [22]. For all these densities, the multi-object posterior has the same form as (3), and can therefore be efficiently approximated using (9). Finally, any MOT filter that produces a set of predictions with state uncertainties and/or existence probabilities (such as some deep-learning-based methods, e.g., [27]) can be seen as having a PMBM density posterior with $H = 1$ and $\lambda(x) = 0 \forall x$, and therefore its NLL can also be approximated by (9).

III. DECOMPOSITION OF THE NLL FOR PMB DENSITIES

An attractive property for a performance measure in MOT is that it decomposes into meaningful terms, which can then be individually analyzed for providing additional insights into the types of errors made by the algorithms. In this section we show that the NLL for PMB densities decomposes into three separate terms that depend on the ability of the posterior to

explain matched objects, missed objects, and false detections, while taking into account all the uncertainties in the posterior. In addition, we provide a connection between the decomposed form of the NLL and the GOSPA metric, showing the under certain conditions the latter is a special case of the former, up to additive offsets.

A. Decomposing the NLL

As a special case of (9) with $H = 1$ and $Q = 1$, the negative log-likelihood (NLL) of a PMB density can be approximated as:

$$\text{NLL}(\mathbb{Y}, f_{\text{PMB}}) \approx \quad (10)$$

$$\min_{\mathbf{A} \in \mathbb{A}} - \sum_{k=1}^m \log f_k(\mathbb{Y}_k(\mathbf{A})) + \int \lambda(y') dy' - \sum_{y \in \mathbb{Y}^U(\mathbf{A})} \log \lambda(y),$$

where we drop the dependency on h and q , and where $\mathbb{A}$ is the set of all matrices in $\{0, 1\}^{(m+|\mathbb{Y}|) \times |\mathbb{Y}|}$ that satisfy (7b). We can also express the optimization in (10) in terms of assignment sets γ , i.e., the set of matched indices (i, j) , where $\gamma(\mathbf{A}) = \{(i, j) \in \mathbb{N}_m \times \mathbb{N}_{|\mathbb{Y}|} \mid [\mathbf{A}]_{i,j} = 1\}$. This change, together with (6), yields

$$\begin{aligned} \text{NLL}(\mathbb{Y}, f_{\text{PMB}}) \approx & \min_{\gamma \in \Gamma} - \underbrace{\sum_{(i,j) \in \gamma} \log(r_i p_i(y_j))}_{\text{Localization}} \\ & - \underbrace{\sum_{i \in \mathbb{F}(\gamma)} \log(1 - r_i)}_{\text{False detections}} + \underbrace{\int \lambda(y') dy' - \sum_{j \in \mathbb{M}(\gamma)} \log \lambda(y_j)}_{\text{Missed objects}}, \end{aligned} \quad (11)$$

where Γ is the set of all possible assignment sets, $\mathbb{F}(\gamma) = \{i \in \mathbb{N}_m \mid \nexists j : (i, j) \in \gamma\}$ is the set of indices of the Bernoullis not matched to any ground-truth, and $\mathbb{M}(\gamma) = \{j \in \mathbb{N}_{|\mathbb{Y}|} \mid \nexists i : (i, j) \in \gamma\}$ is the set of indices of ground-truths not matched to any Bernoulli component. This decomposition separates the NLL score of a PMB density into three parts: the *localization* part, which measures how well the Bernoulli components explain the matched ground-truth, the *false detections* part, which measures the existence probabilities for unmatched Bernoullis, and the *missed objects* part, which measures how well the PMB's PPP component explains the objects not matched to any Bernoulli component. Such decomposition allows practitioners to identify which types of mistakes the algorithms being evaluated are committing.

B. Connection to GOSPA

We proceed to show that (11) can be related to the GOSPA metric [13]. Recall that the GOSPA metric ($p = 1$, $\alpha = 2$) between two sets $\mathbb{X}$ and $\mathbb{Y}$ is

$$\text{GOSPA}_c(\mathbb{X}, \mathbb{Y}) = \min_{\gamma \in \Gamma} \sum_{(i,j) \in \gamma} D(x_i, y_j) + \frac{c}{2} (|\mathbb{X}| + |\mathbb{Y}| - 2|\gamma|),$$

where $D(\cdot, \cdot)$ is a metric on the state-space and c is a cut-off distance. Observe that similarly to NLL, GOSPA decomposes into individual terms for the quality of the localized objects and penalties for false detections and missed objects, but does not account for existence probabilities or object densities. This similarity can be made explicit in the following theorem.

Theorem 1. *If (i) $\lambda(x) = \bar{\lambda}/V \leq 1$, inside a field-of-view with volume V (0 elsewhere), (ii) $r_i = \rho$, $\forall i$, (iii)*

$p_i(x) = \exp(-D(x_i, x))/k$, $\forall i$, where D is a translation-invariant metric,¹ k is the normalization constant, and (iv) $\log(1 - \rho) = \log(\bar{\lambda}/V) \doteq -c/2$, then

$$\text{NLL}(\mathbb{Y}, f_{\text{PMB}}) = V(1 - \rho) \quad (12)$$

$$+ \min_{\gamma} \sum_{(i,j) \in \gamma} \left(D(x_i, y_j) + \log \frac{k}{\rho} \right) + \frac{c}{2} (|\mathbb{X}| + |\mathbb{Y}| - 2|\gamma|)$$

where $\mathbb{X} = \{x_1, \dots, x_m\}$.

Proof. See Appendix A. $\square$

The necessary conditions can be interpreted as (i) the Poisson intensity function takes a constant value over the field-of-view with volume V ; (ii) all the Bernoulli existence probabilities are identical; (iii) the state distribution for all the Bernoulli components of the PMB are set to the same functional form; (iv) there is a relation between the existence probabilities and the Poisson intensity. This result shows that when such conditions hold, GOSPA is a special case of the NLL performance measure up to certain additive constants.

IV. ILLUSTRATIVE EXAMPLES

In this section, we study the properties of the NLL performance measure in two illustrative examples and compare it to three popular MOT performance measures: MOTA [10], MOTP [10], and GOSPA [13]. We use 2.0 as the cutoff distance for establishing matches, Euclidean distance to compute errors, and $\alpha = 2$, $p = 1$ for the GOSPA hyperparameters.

Example 1: Consider two trackers, M_1 and M_2 in a MOT problem where the state is 2-dimensional position, and $\mathbb{Y} = \{(2, 5), (6, 3)\}$. Both trackers have a posterior in the form of a two-component MB with Gaussian state densities (which can be seen as PMBM filters), illustrated in Fig. 1 along with the corresponding existence probability for each component. M_1 's predictions are better than M_2 's: its Bernoulli component closer to A is more certain of the object's existence, and the state density for the component closer to B explains the ground-truth position significantly better than the component from M_2 . In order to compute MOTA, MOTP, and GOSPA scores for each tracker, one would first need to extract "hard" estimates from these posteriors. A common practice is to use the means of the Bernoulli state densities [22], resulting in prediction sets $\mathbb{X}_{M_1} = \{(3, 5), (7, 4)\}$ and $\mathbb{X}_{M_2} = \{(1, 5), (5, 2)\}$. Since the ground-truth positions are equidistant to these estimates, all such performance measures rank M_1 and M_2 equally: $\text{MOTA} = 1$, $\text{MOTP} = \frac{1+\sqrt{2}}{2}$, and $\text{GOSPA} = 1 + \sqrt{2}$ for both M_1 and M_2 . In contrast, the NLL also uses the uncertainty information available in the posteriors of the trackers to compare M_1 and M_2 . Specifically, since the state densities are Gaussian, we have that

$$f_{M_k}(\mathbb{Y}) \approx \prod_{i=1}^2 r_{i,k} \frac{e^{-\frac{1}{2}(y_i - \mu_{i,k})^\top \Sigma_{i,k}^{-1}(y_i - \mu_{i,k})}}{2\pi \sqrt{|\Sigma_{i,k}|}} \quad (13)$$

because the other possible matches between Bernoullis and predictions have negligible contribution. Hence, applying the definition of the NLL from (1), we obtain

¹Translational invariance is needed to allow k to be independent of x_i .

Fig. 1. *Example 1*: existence probabilities and 2-sigma regions for the Bernoulli components of M_1 and M_2 , along with ground-truth object states as black stars.

Fig. 2. *Example 2*: now M_1 has an additional PPP component, with 2-sigma region of the intensity function $\lambda(\cdot)$ illustrated as a dashed ellipse. Same legend as Fig. 1.

$$\begin{aligned} \text{NLL}(\mathbb{Y}, f_{M_k}) &\approx \sum_{i=1}^2 -\log r_{i,k} + \frac{1}{2} \log |\Sigma_{i,k}| \\ &+ \frac{1}{2} (y_i - \mu_{i,k})^\top \Sigma_{i,k}^{-1} (y_i - \mu_{i,k}) + \log(2\pi), \end{aligned} \quad (14)$$

where $\mu_{i,k}$ and $\Sigma_{i,k}$ are the mean and covariance matrix of the i -th Bernoulli component of M_k , and y_i is the ground-truth matched to that component according to the optimal match. Note how these three terms incorporate all the uncertainty information available in the posterior: $(y_i - \mu_{i,k})^\top \Sigma_{i,k}^{-1} (y_i - \mu_{i,k})$ penalizes overconfident state densities, $-\log r_{i,k}$ accounts for the quality of the existence probabilities, while $\log |\Sigma_{i,k}|$ penalizes underconfident state densities. Evaluating (14) for M_1 and M_2 results in $\text{NLL}(\mathbb{Y}, f_{M_1}) \approx 4.6$ and $\text{NLL}(\mathbb{Y}, f_{M_2}) \approx 16.3$, in line with the intuition that M_1 's predictions are considerably better than M_2 's (lower NLL scores are better, indicating higher likelihood).

Example 2: In this example the state is also the 2-dimensional position of the objects, this time with $\mathbb{Y} = \{(2,5), (7,6)\}$. M_1 is now a PMB filter (PMB density as its posterior), while M_2 is a Bernoulli mixture with a single component, illustrated in Fig. 2. The PPP component of M_1 has an intensity function $\lambda(\cdot)$ which is Gaussian, and its 2-sigma region is illustrated as a dashed ellipse. From the figure, we see that the posterior from M_1 is superior in this case too, being vastly better at explaining the ground-truth than M_2 . Although both trackers localize object A equally well, M_2 's posterior is unable to explain any $\mathbb{Y}$ such that $|\mathbb{Y}| > 1$ (likelihood 0), whereas the PPP component of M_1 is able to explain any number of missed objects. Extracting estimates from these

posteriors as in Example 1 results in $\mathbb{X}_{M_1} = \{(2,6)\}$ and $\mathbb{X}_{M_2} = \{(2,4)\}$. Again, MOTA, MOTP, and GOSPA all rank M_1 and M_2 equally, since both estimates are equidistant to ground-truth object A, and both trackers miss the ground-truth object B: MOTA = 0.5, MOTP = 1, and GOSPA = 2 for both trackers. Therefore, all these performance measures completely miss the fact that according to f_{M_2} the realization $\mathbb{Y}$ is impossible. In contrast, the NLL performance measure takes into account the available uncertainty information, revealing the correct ranking of these trackers. In this case, $\text{NLL}(\mathbb{Y}, f_{M_2}) = \infty$, since $(f_{M_2}(\mathbb{Y}) = 0)$, whereas

$\text{NLL}(\mathbb{Y}, f_{M_1}) \approx -\log(r_1 p_1(y_A)) + 1 - \log \lambda(y_B) \approx 8.2$, therefore showing that the NLL performance measure correctly ranks M_1 to be the best tracker in this context.

V. CONCLUSION

We proposed the use of the negative log-likelihood of a model as a MOT performance measure and showed that it incorporates all available uncertainties in the evaluation in a sound mathematical manner, without the need of any hyperparameters. We also provided efficient algorithms for approximating the NLL of common MOT methods, along with a special case for the PMB density where the NLL decomposes and is an approximate generalization of the GOSPA metric. Examples indicate that NLL is better at capturing the expected performance than conventional performance measures. The adoption of NLL can hopefully provide a fairer way of comparing MOT methods and guide novel designs, especially for deep learning methods, which often fail to provide uncertainty estimates and thus have poor NLL.

APPENDIX A PROOF OF THEOREM 1

Starting from (11):

$$\begin{aligned} \text{NLL}(\mathbb{Y}, f_{\text{PMB}}) &\approx \min_{\gamma \in \Gamma} - \sum_{(i,j) \in \gamma} \log r_i p_i(y_j) \\ &- \sum_{i \in \mathbb{F}(\gamma)} \log(1 - r_i) + \int \lambda(y') dy' - \sum_{j \in \mathbb{M}(\gamma)} \log \lambda(y_j) \end{aligned}$$

and invoking assumptions (i)–(iv), we obtain

$$\begin{aligned} &\stackrel{(i)}{=} \min_{\gamma \in \Gamma} - \sum_{(i,j) \in \gamma} \log r_i p_i(y_j) \\ &- \sum_{i \in \mathbb{F}(\gamma)} \log(1 - r_i) + \bar{\lambda} - \left(\log \frac{\bar{\lambda}}{V} \right) (|\mathbb{Y}| - |\gamma|) \\ &\stackrel{(ii)}{=} \min_{\gamma \in \Gamma} - \sum_{(i,j) \in \gamma} \left(\log p_i(y_j) + \log \rho \right) \\ &- \log(1 - \rho) (|\mathbb{X}| - |\gamma|) + \bar{\lambda} - \left(\log \frac{\bar{\lambda}}{V} \right) (|\mathbb{Y}| - |\gamma|) \\ &\stackrel{(iii)}{=} \min_{\gamma \in \Gamma} \sum_{(i,j) \in \gamma} \left(D(x_i, y_j) + \log \frac{k}{\rho} \right) \\ &- \log(1 - \rho) (|\mathbb{X}| - |\gamma|) + \bar{\lambda} - \left(\log \frac{\bar{\lambda}}{V} \right) (|\mathbb{Y}| - |\gamma|) \\ &\stackrel{(iv)}{=} V(1 - \rho) + \min_{\gamma \in \Gamma} \sum_{(i,j) \in \gamma} \left(D(x_i, y_j) + \log \frac{k}{\rho} \right) \\ &+ \frac{c}{2} (|\mathbb{X}| + |\mathbb{Y}| - 2|\gamma|). \end{aligned}$$

REFERENCES

- [1] Y. Bar-Shalom, X. R. Li, and T. Kirubarajan, *Estimation with applications to tracking and navigation: theory algorithms and software*. John Wiley & Sons, 2004.
- [2] S. Challa, M. R. Morelande, D. Mušicki, and R. J. Evans, *Fundamentals of object tracking*. Cambridge University Press, 2011.
- [3] F. Meyer, T. Kropfreiter, J. L. Williams, R. Lau, F. Hlawatsch, P. Braca, and M. Z. Win, "Message passing algorithms for scalable multitarget tracking," *Proceedings of the IEEE*, vol. 106, no. 2, pp. 221–259, 2018.
- [4] R. L. Streit, R. B. Angle, and M. Efe, *Analytic combinatorics in multiple object tracking*. Springer, 2021.
- [5] R. P. Mahler, *Statistical Multisource-Multitarget Information Fusion*. Artech House Norwood, MA, 2007.
- [6] Y. Bar-Shalom, *Multitarget-multisensor tracking: advanced applications*. Norwood, 1990.
- [7] S. Shafaei, S. Kugele, M. H. Osman, and A. Knoll, "Uncertainty in machine learning: A safety perspective on autonomous driving," in *International Conference on Computer Safety, Reliability, and Security*. Springer, 2018, pp. 458–464.
- [8] O. E. Drummond and B. E. Fridling, "Ambiguities in evaluating performance of multiple target tracking algorithms," in *Signal and Data Processing of Small Targets 1992*, O. E. Drummond, Ed., vol. 1698, International Society for Optics and Photonics. SPIE, 1992, pp. 326 – 337. [Online]. Available: <https://doi.org/10.1117/12.139399>
- [9] J. R. Hoffman and R. P. Mahler, "Multitarget miss distance via optimal assignment," *IEEE Transactions on Systems, Man, and Cybernetics-Part A: Systems and Humans*, vol. 34, no. 3, pp. 327–336, 2004.
- [10] K. Bernardin and R. Stiefelhagen, "Evaluating multiple object tracking performance: the clear mot metrics," *EURASIP Journal on Image and Video Processing*, vol. 2008, pp. 1–10, 2008.
- [11] J. Luiten, A. Osep, P. Dendorfer, P. Torr, A. Geiger, L. Leal-Taixé, and B. Leibe, "Hota: A higher order metric for evaluating multi-object tracking," *International journal of computer vision*, vol. 129, no. 2, pp. 548–578, 2021.
- [12] D. Schuhmacher, B.-T. Vo, and B.-N. Vo, "A consistent metric for performance evaluation of multi-object filters," *IEEE transactions on signal processing*, vol. 56, no. 8, pp. 3447–3457, 2008.
- [13] A. S. Rahmathullah, Á. F. García-Fernández, and L. Svensson, "Generalized optimal sub-pattern assignment metric," in *20th International Conference on Information Fusion (Fusion)*. IEEE, 2017, pp. 1–8.
- [14] P. Dendorfer, H. Rezatofighi, A. Milan, J. Shi, D. Cremers, I. Reid, S. Roth, K. Schindler, and L. Leal-Taixé, "Mot20: A benchmark for multi object tracking in crowded scenes," *arXiv:2003.09003[cs]*, Mar. 2020, arXiv: 2003.09003. [Online]. Available: <http://arxiv.org/abs/1906.04567>
- [15] X. R. Li, Z. Zhao, and V. P. Jilkov, "Estimator's credibility and its measures," in *Proc. IFAC 15th World Congress*, 2002.
- [16] S. Nagappa, D. E. Clark, and R. Mahler, "Incorporating track uncertainty into the OSPA metric," in *14th International Conference on Information Fusion*. IEEE, 2011, pp. 1–8.
- [17] X. He, R. Tharmarasa, T. Kirubarajan, and T. Thayaparan, "A track quality based metric for evaluating performance of multitarget filters," *IEEE Transactions on Aerospace and Electronic Systems*, vol. 49, no. 1, pp. 610–616, 2013.
- [18] T. A. Severini, *Likelihood methods in statistics*. Oxford University Press, 2000.
- [19] I. J. Myung, "Tutorial on maximum likelihood estimation," *Journal of mathematical Psychology*, vol. 47, no. 1, pp. 90–100, 2003.
- [20] R. P. Mahler, "Multitarget Bayes filtering via first-order multitarget moments," *IEEE Transactions on Aerospace and Electronic systems*, vol. 39, no. 4, pp. 1152–1178, 2003.
- [21] R. Mahler, "PHD filters of higher order in target number," *IEEE Transactions on Aerospace and Electronic systems*, vol. 43, no. 4, pp. 1523–1543, 2007.
- [22] Á. F. García-Fernández, J. L. Williams, K. Granström, and L. Svensson, "Poisson multi-Bernoulli mixture filter: direct derivation and implementation," *IEEE Transactions on Aerospace and Electronic Systems*, vol. 54, no. 4, pp. 1883–1901, 2018.
- [23] J. L. Williams, "Marginal multi-Bernoulli filters: RFS derivation of MHT, JIPDA, and association-based MeMBer," *IEEE Transactions on Aerospace and Electronic Systems*, vol. 51, no. 3, pp. 1664–1687, 2015.
- [24] Á. F. García-Fernández, Y. Xia, K. Granström, L. Svensson, and J. L. Williams, "Gaussian implementation of the multi-Bernoulli mixture filter," in *22th International Conference on Information Fusion (FUSION)*. IEEE, 2019, pp. 1–8.
- [25] B.-T. Vo and B.-N. Vo, "Labeled random finite sets and multi-object conjugate priors," *IEEE Transactions on Signal Processing*, vol. 61, no. 13, pp. 3460–3475, 2013.
- [26] D. F. Crouse, "On implementing 2D rectangular assignment algorithms," *IEEE Transactions on Aerospace and Electronic Systems*, vol. 52, no. 4, pp. 1679–1696, 2016.
- [27] A. Harakeh, M. Smart, and S. L. Waslander, "Bayesod: A Bayesian approach for uncertainty estimation in deep object detectors," in *ICRA*. IEEE, 2020, pp. 87–93.