A Confidence Measure for Feature-based Localization using Random Finite Sets

Manuel Stübler, Jürgen Wiest, Stephan Reuter and Klaus Dietmayer*

Abstract: A reliable confidence measure for the localization of road vehicles is a crucial requirement to enable highly automated driving with an inherent self-awareness of its robustness and capabilities. Therefore, the present contribution introduces a novel approach to estimate not only the spatial uncertainty of a feature-based localization algorithm but also an associated confidence measure. The methodology makes use of Random Finite Sets to model and quantitatively describe the (expected) difference between a given map of landmarks and a set of corresponding online measurements. Based on a probabilistic representation of the two sets, the above-mentioned confidence measure is derived. Simulation results based on a Monte-Carlo localization framework are provided.

Keywords: Confidence, Reliability, Localization, Random Finite Sets

1 Introduction

Self-localization is one of the well-studied topics in robotics and autonomous systems. A detailed summary of common methodologies can for example be found in [1]. These include Markov localization and multi-hypothesis tracking with Kalman filters as well as grid and Monte-Carlo localization approaches with histogram and particle filters respectively. Despite the long-term research in this area, there is hardly any contribution found that addresses the problem of online measuring the confidence and reliability of such a localization process.

A basic concept to measure the confidence of Monte-Carlo localization (MCL) algorithms was outlined in [1] where the fundamental idea is to compare the current measurement likelihood to its average, which is for example learned from previously collected data. However, this approach cannot be directly adopted to the multi-object domain due to the incomparability of the multi-object likelihood between different measurement-to-landmark constellations, especially regarding the cardinality of the landmark map. Another way to evaluate the performance of a (single-object) state estimator is to calculate the Normalized Innovation Squared (NIS). The NIS is typically used for the online verification of a Kalman filter (and its extensions) and follows a χ^2 distribution for a consistent state estimator if the uncertainty is assumed to be Gaussian. In the multi-object case, e.g. regarding feature-based self-localization algorithms using multi-object likelihoods, this concept is again not applicable anymore. Therefore, the Multi-object Generalized NIS (MGNIS) was derived in [2] which facilitates an online evaluation of the

*The authors are with driveU / Institute of Measurement, Control and Microtechnology, Ulm University, 89081 Ulm, Germany. E-Mail: {firstname.lastname} at uni-ulm.de

performance for multi-object tracking algorithms. However, the MGNIS has the major drawback that it does not follow a χ^2 distribution anymore and additionally is dominated by clutter. Therefore, its interpretation is quite problematic [3].

A different approach to evaluate the performance of multi-object filtering algorithms with ground truth data is based on the Wasserstein metric. Examples of a Wasserstein metric are, e.g., the Optimal Mass Transform (OMAT) metric [4] and the Optimal Subpattern Assignment (OSPA) metric [5]. Both metrics provide a measure for the distance of two Random Finite Sets (RFSs) that do not necessarily have the same cardinality and are therefore especially well-suited for the evaluation of tracking algorithms based on RFSs. The fact that clutter measurements and misdetections in general increase the OMAT distance in a disproportionate manner led to the introduction of a cut-off parameter in the OSPA metric. Another Wasserstein metric, namely the Cardinalized Optimal Linear Assignment (COLA) metric, that especially facilitates the evaluation of Simultaneous Localization And Mapping (SLAM) algorithms was recently published in [6] and is derived from the OSPA metric. In fact, the two latter metrics try to incorporate both, the cardinality and the spatial difference of RFSs at a time where the OSPA metric delivers a spatial distance and the COLA metric delivers a cardinality distance. Because the Wasserstein metrics provide a distance measure between two RFSs, they can also be used to measure the difference between a map of landmarks and a set of online measurements. But again, they suffer from the effect that the distance between two sets tends to increase with their cardinality. Another metric that facilitates the comparison of RFSs is the Hausdorff metric which is typically dominated by outliers and therefore is hardly applicable when clutter measurements and misdetections are present.

The proposed method circumvents the previously outlined drawbacks and provides a probabilistic measure of the confidence and the spatial uncertainty of feature-based localization algorithms. The derived confidence measure is based on the multi-object likelihood [7] though its corresponding spatial uncertainty estimate is closely related to the OSPA metric. The remaining parts of this contribution are structured as follows: first, in Section 2, there is a brief introduction into the basic concepts and fundamentals that are required for the understanding of the following definition of the confidence measure in Section 3. Evaluation results based on a simulation framework are provided in Section 4 and a final conclusion is given in Section 5.

2 Basics

2.1 Random Finite Sets

A Random Finite Set (RFS) is defined by

$$X = \{x_1, \dots, x_n\}, \tag{1}$$

where the elements $x_i \in X$ are distributed according to a probability density function $f(x_i)$ and the number of elements $n = |X|$ follows a cardinality distribution $\rho(n)$. Regarding the tasks of mapping and localization, RFSs are a suitable way of representing feature maps and sets of measurements as they do not imply an intrinsic order in contrast to random vectors [8]. Furthermore, with the Finite Set Statistics (FISST) [7] a Bayesian framework is available that facilitates the usage of RFSs for feature-based mapping and

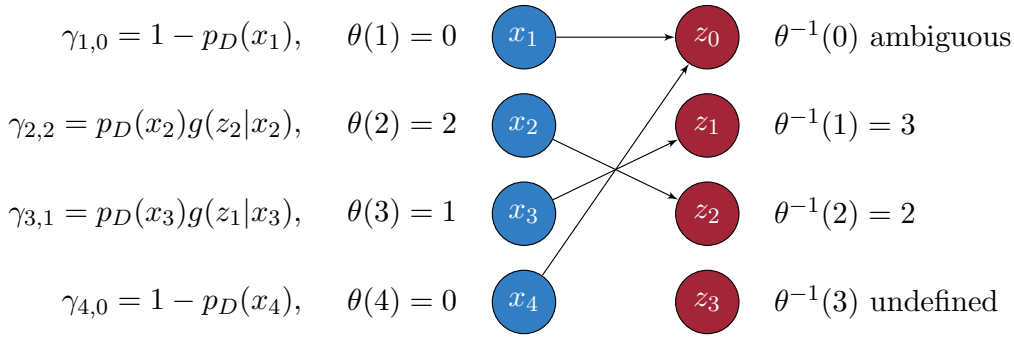

Figure 1: Illustration of an association function θ with $n = 4$ and $m = 3$, where z_3 is a clutter measurement and x_1 as well as x_4 are misdetections that were assigned to the notional misdetection measurement z_0. Each landmark x_i has a corresponding association likelihood $\gamma_{i,\theta(i)}$ based on its assigned measurement $z_{\theta(i)}$.

localization. In general, a map excerpt that lies within the Field of View (FoV) of the respective sensor $X = M \cap \text{FoV}$ is used. The map typically comprises landmark positions in global coordinates whereas the measurements represent the feature detections in a local coordinate system. The two-dimensional vehicle pose $\hat{\boldsymbol{p}} = [\, x \; y \; \psi \,]^{\mathrm{T}}$ can for example be estimated by maximizing the (multi-object) likelihood function $g(Z|\hat{\boldsymbol{p}}, X)$ because the likelihood for a set of measurements accordingly depends on the map of the environment and the pose of the vehicle. An additional process model can be incorporated by utilizing a particle filtering method. In the literature, this approach is known under the term MCL [9, 10]. The fact that the true association between map landmarks and measurements is unknown is addressed by summing over all possible hypotheses within the multi-object likelihood following [7]. Because of the exponentially increasing computational power, the multi-object likelihood is typically approximated by summing only over the k best assignments.

2.2 Association function

Given a set of measurements $Z = \{z_1, \dots, z_m\}$ and a map of landmarks $X = \{x_1, \dots, x_n\}$, the association function $\theta : \{1, \dots, n\} \to \{0, 1, \dots, m\}$ assigns to each landmark x_i with $i \in \{1, \dots, n\}$ either a (true) measurement z_j with $j \in \{1, \dots, m\}$ or the notional misdetection element z_0. An important property of the association function is that a measurement cannot be assigned to more than one landmark, thus $\theta(i_1) = \theta(i_2) > 0 \Rightarrow i_1 = i_2$. However, several undetected landmarks may be assigned to z_0 which is actually not a true measurement but a mathematical construct. Figure 1 illustrates an exemplary association function with 4 landmarks and 3 measurements. An association function θ can also be expressed through an equivalent assignment matrix A and vice versa (c.f. Section 3.1.2).

3 Confidence measure

Based on the concept of the multi-object likelihood derived in [7], in this section the mean association likelihood $\bar{\gamma}_\theta(Z|X)$ for a set of measurements Z, a map of landmarks X and a known association function θ is introduced as a measure for the confidence of feature-based localization estimates. A related approach is outlined in [1] (p. 257ff.) by comparing

the current (single-object) measurement likelihood $g(z|x)$ to the average likelihood that is for example learned from the recorded dataset. However, the multi-object likelihood degenerates in the number of measurements as well as in the number of map landmarks and consequently would favor smaller cardinalities over larger ones. Thus, the basic idea of this contribution is to average the (single-object) association likelihood

$$\gamma_{i,\theta(i)} = \begin{cases} p_D(x_i) \cdot g(z_{\theta(i)}) & \text{if } \theta(i) > 0 \\ 1 - p_D(x_i) & \text{if } \theta(i) = 0 \end{cases} \tag{2}$$

of a landmark x_i and an associated measurement $z_{\theta(i)}$.

3.1 Mean association likelihood

By disregarding clutter measurements for the moment, the mean association likelihood of a set of measurements $Z = \{z_1, \dots, z_m\}$ and a map of landmarks $X = \{x_1, \dots, x_n\}$ can be written as

$$\bar{\gamma}_\theta(Z|X) = \left(\prod_{i=1}^{n} (1 - p_D(x_i)) \prod_{i \in \mathcal{D}_\theta} \frac{p_D(x_i) \cdot g(z_{\theta(i)}|x_i)}{(1 - p_D(x_i))} \right)^{\frac{1}{n}}, \tag{3}$$

where $g(\cdot|\cdot) \in [0, 1]$ denotes a normalized single-object likelihood function, θ is an association function according to Section 2.2 and $p_D(\cdot)$ represents the detection probability. The set of indices of all detected map landmarks is given by $\mathcal{D}_\theta := \{i : \theta(i) > 0\}$ (c.f. Figure 1).

3.1.1 Single-object likelihood

One possibility of defining a (normalized) single-object likelihood is based on the Mahalanobis distance:

$$g(z_j|x_i) = \exp\left(-\frac{d(x_i, z_j)^2}{2\sigma_j^2}\right). \tag{4}$$

In this case, $d(\cdot, \cdot)$ is the (Euclidean) distance and the covariance matrix $\Sigma_j = \sigma_j^2 I_N$ that corresponds to a measurement z_j is assumed to be a diagonal matrix of size $N = \dim(z_j)$ where the diagonal elements are σ_j^2. Furthermore, the objects x_i are assumed to be point landmarks with an infinitesimal extent. However, this is just an exemplary definition and might be interchanged by another one that fits better to the particular (sensor) setup.

3.1.2 Association function estimate

In general, the true association function is unknown. Therefore, multi-object tracking algorithms consider all possible hypotheses resulting in an exponential growth of the required computational power. A common approximation is to maintain in each step only the most likely hypotheses. In the present contribution, the mean association likelihood between a map X and a set of measurements Z is assumed to be the one corresponding to the association that maximizes the confidence measure following

$$\bar{\gamma}^*(Z|X) = \max_\theta \bar{\gamma}_\theta(Z|X). \tag{5}$$

C	z_1	$\dots$	z_m	$z_0^{(1)}$	$\dots$	$z_0^{(n)}$
x_1	η_{11}	$\dots$	η_{1m}	η_{10}	$\dots$	η_{10}
$\vdots$	$\vdots$	$\ddots$	$\vdots$	$\vdots$	$\ddots$	$\vdots$
x_n	η_{n1}	$\dots$	η_{nm}	η_{n0}	$\dots$	η_{n0}

(a) The general cost matrix $C \in \mathbb{R}^{n\times(m+n)}$ that solves the assignment problem. The columns $z_0^{(i)}$ represent a misdetection of the respective landmark.

A	z_1	z_2	z_3	$z_0^{(1)}$	$z_0^{(2)}$	$z_0^{(3)}$	$z_0^{(4)}$	$\sum$
x_1	0	0	0	1	0	0	0	1
x_2	0	1	0	0	0	0	0	1
x_3	1	0	0	0	0	0	0	1
x_4	0	0	0	0	0	0	1	1

(b) An assignment matrix A that corresponds to the exemplary association function from Figure 1.

Table 1: Illustration of the cost matrix C and the assignment matrix A.

This assumption is particularly reasonable for distinct associations between measurements and map landmarks where the multi-object likelihood is dominated by the association with the highest weight (likelihood). Furthermore, this ensures a confidence measure within the interval $[0, 1]$ if the single-object likelihood itself is already normalized. The maximization of the mean association likelihood can be realized with the Munkres assignment algorithm [12]. The cost matrix $C \in \mathbb{R}^{n\times(m+n)}$ that solves the respective assignment problem is illustrated in Table 1a. The individual costs for assigning a measurement z_j to a map landmark x_i with $i \in \{1, \dots, n\}$ are given by

$$\eta_{ij} = \begin{cases} -\log(p_D(x_i)g(z_j|x_i)) & \text{if } j \in \{1, \dots, m\} \\ -\log(1 - p_D(x_i)) & \text{if } j = 0 \end{cases}, \tag{6}$$

where $j = 0$ represents a misdetection, $p_D(\cdot)$ is the detection probability and $g(\cdot|\cdot)$ states the single-object likelihood. The result of the Munkres assignment algorithm is an assignment matrix $A \in \{0, 1\}^{n\times(m+n)}$ as illustrated in Table 1b.

3.1.3 Cut-off distance

With the assumption that all measurements z_j are well separable and originate from a distinct map landmark x_i, the contribution of such a particular map landmark to the mean association likelihood is given by

$$\gamma_i = \max\left\{1 - p_D(x_i), p_D(x_i) \max_j g(z_j|x_i)\right\}. \tag{7}$$

Consequently, the misdetection probability is kind of a cut-off parameter that takes effect whenever the distance between a measurement z_j and its corresponding map landmark x_i is greater than

$$\check{d}_{ij} = \sqrt{-2\sigma_j^2 \log\left(\frac{1 - p_D(x_i)}{p_D(x_i)}\right)} \tag{8}$$

with $p_D(x_i) > 0.5$. Equation (8) represents the state where detection and misdetection of a map landmark are equally likely, again disregarding the influence of the clutter source for the moment. If $p_D(x_i) \leq 0.5$, then $\check{d}_{ij}$ is not a real valued (positive) number and a misdetection is always preferred, no matter how close a measurement z_j is located to x_i.

The cut-off is somehow comparable to the one introduced in the OSPA metric, except that it depends on the detection probability $p_D(x_i)$ of a landmark x_i and the standard deviation σ_j of a measurement z_j. This also means that a false association will be made if z_j is originating from x_i and has a higher distance than the cut-off $\check{d}_{ij}$. An illustration of equation (7) is shown in Figure 3b for the cut-off distances $d_{e,1} = \sigma$, $d_{e,2} = 2\sigma$ and $d_{e,3} = 3\sigma$ that correspond to the respective detection probabilities $p_{D,1} \approx 0.62$, $p_{D,2} \approx 0.88$ and $p_{D,3} \approx 0.99$.

3.2 Localization error estimate

The localization error of order p between X and Y that corresponds to an association function θ is defined by:

$$\bar{\epsilon}_{\theta,p}(X,Z) = \left(\frac{1}{|\mathcal{D}_\theta|} \sum_{i \in \mathcal{D}_\theta} d(x_i, z_{\theta(j)})^p \right)^{\frac{1}{p}}. \tag{9}$$

This estimate takes only those landmarks into account that were actually assigned to a measurement. This is reasonable because misdetections and clutter measurements should have an impact on the confidence measure and not on the localization error estimate. The choice of p directly influences the weighting of outliers. For $p = 1$ the localization error estimate is equivalent to the average distance between map landmarks and their assigned measurements. For $p = 2$ the error estimate corresponds to the root mean square distance and weights outliers higher than the arithmetic mean does. The minimum localization error estimate $\bar{\epsilon}_{\theta,p}$ over all hypotheses θ is equal to the localization error $e^{(c)}_{p,loc}$ [5] of the OSPA metric by disregarding clutter measurements and misdetections.

3.3 Clutter-extended mean association likelihood

By modeling the clutter as being Poisson distributed in the cardinality, uniformly distributed in the FoV and originating from a single clutter source that is treated like a permanent map element of X, the mean association likelihood can be extended to incorporate clutter measurements as follows:

$$\bar{\gamma}_{\theta,\lambda}(Z|X) := \left(p_\lambda(m - |\mathcal{D}_\theta|) \prod_{i=1}^{n} (1 - p_D(x_i)) \prod_{i \in \mathcal{D}_\theta} \frac{p_D(x_i) \cdot g(z_{\theta(i)}|x_i)}{(1 - p_D(x_i))} \right)^{\frac{1}{n+1}}, \tag{10}$$

where $m - |\mathcal{D}_\theta|$ is the number of unassigned measurements and $p_\lambda(k) = \frac{\lambda^k}{k!} e^{-\lambda}$ is the Poisson distributed probability of having k clutter measurements in case of an expectation value $\lambda > 0$. This can be thought of as an extension of X by a notional clutter source x_0 to which all the clutter measurements are assigned. Therefore, the clutter source is treated just like any other landmark $x_i \in X$ with the only difference that it can produce several scattered measurements within the FoV. Analagous to Section 3.1.2, the association function θ can again be estimated via maximization:

$$\bar{\gamma}^*_\lambda(Z|X) = \max_\theta \bar{\gamma}_{\theta,\lambda}(Z|X). \tag{11}$$

If the number of expected clutter measurements λ is unknown, it can for example be additionally estimated alongside with the association function θ.

3.4 Implementation

A major drawback of solving equation (11) lies within the dependence regarding the assignment of measurements to landmarks with respect to the clutter source. To be precise, the association likelihood of the clutter source $p_\lambda(m-|\mathcal{D}_\theta|)$ depends on the overall association function θ itself. In contrast, solving equation (5) is much simpler because in this case the cost η_{ij} for assigning a measurement z_j to a map landmark x_i is completely independent from the assignment of the other measurements. Therefore, the cost matrix C is generated according to equation (6) (c.f. Table 1) which corresponds to finding the best association for $\bar{\gamma}_\theta(Z|X)$ instead of $\bar{\gamma}_{\theta,\lambda}(Z|X)$. The optimal assignment problem is then solved (e.g. by using the Munkres algorithm) for the mean association likelihood by neglecting clutter measurements. The result of this step is an assignment matrix A which is an equivalent representation of the association function $\theta^* = \arg\max_\theta \bar{\gamma}_\theta(Z|X)$. In particular, this means that all measurements that were not assigned to a (regular) map landmark following equation (5) will be assigned to the clutter source x_0 afterwards. The confidence measure is thus approximated by

$$\bar{\gamma}^*_\lambda(Z|X) \approx \bar{\gamma}_{\theta^*,\lambda}(Z|X) = \left(p_\lambda(m-|\mathcal{D}_A|)\cdot e^{-\eta_{A,C}}\right)^{\frac{1}{n+1}} \tag{12}$$

where $|\mathcal{D}_A| = \sum_{i=1}^{n}\sum_{j=1}^{m} a_{ij}$ is the number of detected landmarks (those with an associated measurement) and $\eta_{A,C} = \sum_{i=1}^{n}\sum_{j=1}^{m+n} a_{ij}c_{ij}$ is the overall cost for the respective assignment. Certainly, the approximated confidence measure might be underestimated but this, indeed, is less noticeable for a large number of map landmarks due to the reduced influence of the clutter association likelihood in this case.

4 Evaluation

The evaluation was done with a simulation framework that implements a feature-based Monte-Carlo localization method. The localization algorithm is similar to the one running in the autonomous driving vehicle at Ulm University [13].

4.1 Simulation framework

The simulation scenario consists of $N_L = 41$ landmarks that are randomly scattered alongside a pre-defined trajectory in each run whereas the trajectory itself is invariant. The two-dimensional pose of the vehicle is defined by $\boldsymbol{p} = [x\ y\ \psi]^{\mathrm{T}}$ and is estimated using a particle filter with a total of $N = 100$ particles. The modeled sensor has a FoV of 360° and is able to detect landmarks within a range of $R = (1\text{m}, 20\text{m})$. Landmarks are detected by the sensor with a constant probability of $p_D = 0.88$ and have an unbiased Gaussian spatial uncertainty with a standard deviation of $\sigma = 0.1\text{m}$. The innovation step is carried out based on the multi-object likelihood $g(Z|X)$ having a set of measurements Z for a given set of map landmarks X [7]. The number of clutter measurements per frame follows a Poisson distribution with expectation value $\lambda = 1$. The confidence measure $\bar{\gamma}^*_\lambda(Z|X)$ is approximated using equation (12) and the corresponding localization error estimate follows equation (9). Both measures were calculated only for the pose estimate $\hat{\boldsymbol{p}}$ which equals the weighted mean of the particle set. An extract of an exemplary Monte-Carlo run is shown in Figure 2a.

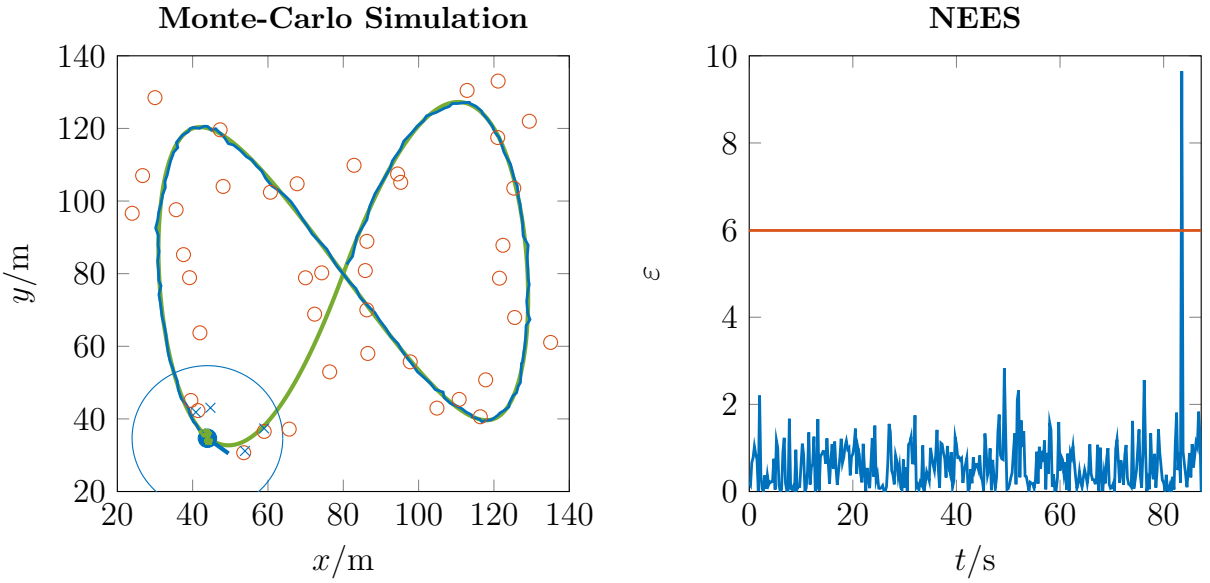

(a) Extract of an exemplary Monte-Carlo run: The green and blue line represent the true and estimated trajectory. Red circles are map landmarks, blue crosses are measurements and green dots represent the particles.

(b) The corresponding NEES (blue line) and inverse χ^2 distribution for two degrees of freedom with a one-sided confidence region of 95% (red line).

Figure 2: An exemplary Monte-Carlo run using the simulation framework.

4.2 Consistency check

The Normalized Estimation Error Squared (NEES) for a given timestep k is defined by $\varepsilon(k) = \boldsymbol{x}^{\mathrm{T}}(k)P^{-1}(k)\boldsymbol{x}(k)$. It is illustrated in Figure 2b for the related excerpt shown in Figure 2a. The corresponding one-sided confidence interval for two degrees of freedom and a confidence region of 95% is represented by the red line. A consistent state estimator is assumed to violate this threshold in 5% of the cases. Consistency thereby is defined such that for an unbiased estimator the true and estimated uncertainty coincide.

Another consistency check using N independent samples can be carried out by considering the mean NEES which is given by $\bar{\varepsilon} = \frac{1}{N}\sum_{k=1}^{N}\varepsilon(k)$. For a consistent estimator, $N \cdot \bar{\varepsilon}$ should follow a χ^2 distribution with $N \cdot \dim(\boldsymbol{x})$ degrees of freedom [14]. This implies that the mean NEES becomes arbitrarily close to $\dim(\boldsymbol{x})$ for $N \to \infty$. The mean confidence measure for the exemplary run shown in Figure 2 is approximately 0.51 where the mean NEES of this run is approximately 0.54. A mean NEES below $\dim(\boldsymbol{x})$ indicates that the uncertainty is overestimated. One reason for this is that an orientation error is not covered by equation (9) and therefore accounts for an accordingly higher uncertainty regarding the position error of the vehicle.

In Figure 3a, the mean NEES of 100 Monte-Carlo runs is plotted against different confidence threshold values by taking only samples into account that satisfy $\bar{\gamma}(k) > \bar{\gamma}_{th}$. There is a quite sharp rise for $\bar{\gamma}_{th} \lesssim 0.39$ which can be explained as follows: If there are no landmarks within the FoV, the confidence measure only depends on the clutter association likelihood that follows a Poisson distribution. With the assumption that $\lambda = 1$, the highest possible confidence measure is $\bar{\gamma}_1^* = e^{-1} \approx 0.37$ which dominates even if there are a few coincidental matches between the two sets X and Z.

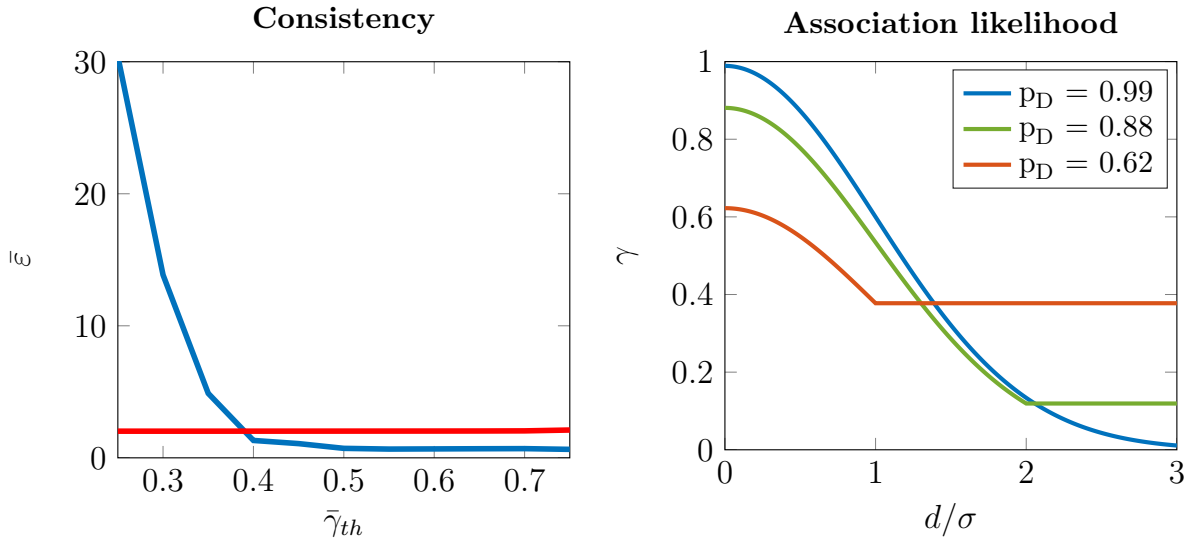

(a) Simulation results based on 100 Monte-Carlo runs: Illustrated is the mean NEES for samples with a confidence measure higher than $\bar{\gamma}_{th}$ (blue line) and the inverse χ^2 distribution with $\alpha = 0.05$ (red line).

(b) The (single-object) association likelihood for different detection probabilities p_D. The cut-off distances correspond to σ (red line), 2σ (green line) and 3σ (blue line).

Figure 3: Evaluation results of the consistency measure based on the mean NEES (left) and illustration of the single-object association likelihood γ_i (right).

5 Conclusion

In this contribution, a novel approach to estimate the localization error alongside with a corresponding confidence measure was proposed. The concept was derived for feature-based localization algorithms using a probabilistic model based on Random Finite Sets. The implementation makes use of the Munkres assignment algorithm and has the complexity $\mathcal{O}(n^3)$. The proposed methodology was evaluated using a simulation framework that implements a feature-based Monte-Carlo localization approach. It showed promising results in terms of a reliable and meaningful measure for the confidence of a multi-object state and uncertainty estimate. An additional evaluation with real-world data is one of the next steps that should be carried out in order to provide further revealing conclusions. Moreover, an appropriate filtering method to smoothen the confidence measure would be beneficial.

References

[1] S. Thrun, W. Burgard, and D. Fox, *Probabilistic Robotics (Intelligent Robotics and Autonomous Agents)*. The MIT Press, 2005.

[2] R. Mahler, "Divergence detectors for multitarget tracking algorithms," in *Proceedings SPIE 8745, Signal Processing, Sensor Fusion, and Target Recognition XXII*, 2013.

[3] S. Reuter, "Multi-object tracking using random finite sets," Ph.D. dissertation, Ulm University, 2014.

[4] J. Hoffman and R. 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.

[5] 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.

[6] P. Barrios, G. Naqvi, M. Adams, K. Leung, and F. Inostroza, "The cardinalized optimal linear assignment (cola) metric for multi-object error evaluation," in *Proceedings of the 18th International Conference on Information Fusion*, 2015, pp. 271–279.

[7] R. Mahler, *Statistical Multisource-Multitarget Information Fusion.* Artech House Inc., Norwood, 2007.

[8] J. Mullane, B.-N. Vo, M. Adams, and B.-T. Vo, *Random Finite Sets for Robot Mapping and SLAM - New Concepts in Autonomous Robotic Map Representations*, ser. Springer Tracts in Advanced Robotics. Springer, 2011, vol. 72.

[9] F. Dellaert, D. Fox, W. Burgard, and S. Thrun, "Monte carlo localization for mobile robots," in *IEEE International Conference on Robotics and Automation (ICRA)*, vol. 2, 1999, pp. 1322–1328.

[10] P. Jensfelt, D. Austin, O. Wijk, and M. Andersson, "Feature based condensation for mobile robot localization," in *IEEE International Conference on Robotics and Automation (ICRA)*, vol. 3, 2000, pp. 2531–2537.

[11] H. Deusch, J. Wiest, S. Reuter, D. Nuss, M. Fritzsche, and K. Dietmayer, "Multi-sensor self-localization based on maximally stable extremal regions," in *IEEE Intelligent Vehicles Symposium*, 2014, pp. 555–560.

[12] J. Munkres, "Algorithms for the assignment and transportation problems," *Journal of the Society for Industrial and Applied Mathematics*, vol. 5, no. 1, pp. 32–38, 1957.

[13] F. Kunz, D. Nuss, J. Wiest, H. Deusch, S. Reuter, F. Gritschneder, A. Scheel, M. Stuebler, M. Bach, P. Hatzelmann, C. Wild, and K. Dietmayer, "Autonomous driving at ulm university: A modular, robust, and sensor-independent fusion approach," in *Intelligent Vehicles Symposium*, 2015, pp. 666–673.

[14] Y. Bar-Shalom and T. Fortmann, *Tracking and Data Association.* Academic Press, Inc., 1988.