Maximum Likelihood Combining of Stochastic Maps

Brandon Jones, Mark Campbell, and Lang Tong
Cornell University, Ithaca, NY 14850, USA
{bmj34,mc288,lt35}@cornell.edu

Abstract— The problem of combining stochastic maps obtained by independent agents is considered. Using the generalized likelihood ratio statistics, the problem of matching triangles that correspond to common landmark observations in different stochastic maps is formulated as a bipartite matching problem with generalized likelihood ratio statistics. From the matched triangles between each map, the maximum likelihood combining of stochastic maps is generated. It is shown that the generalized likelihood ratio statistic and the maximum likelihood combining of maps can be computed in closed form, which makes the proposed algorithm a scalable solution to matching and combining stochastic maps with a large number of landmarks.

Index terms: Stochastic maps, simultaneous localization and mapping (SLAM), maximum likelihood combining.

I. INTRODUCTION

The problem of combining stochastic maps from two or more independent agents arises naturally in autonomous robotics. By stochastic map we refer to the noisy estimates of landmark locations produced by a robot during its exploration. For example, when two robots explore an area independently, each constructs a stochastic map of the explored region under its own coordinate system by solving the *simultaneous localization and mapping* (SLAM) problem [1][2][3]. Although the maps of two separate robots that individually conduct SLAM are in general different (as they may explore different areas of the environment), some of the landmarks in the explored regions may have been observed by both robots. And if the two robots meet during their exploration, they may wish to construct a more informative and accurate *joint map* from their individual maps. The challenge of combining individual maps arises from the fact that the two robots have no common coordinate systems, the landmarks common to both maps are unknown, and their sensors are noisy.

If two robots each has a stochastic map with n estimated landmarks, there are total 2^n possible ways of matching one map with the other. In principle, the optimal way of combining the maps is to determine the best match among the 2^n possibilities under a certain criterion of "best fit." Such an approach is of course not tractable even when n is relatively small. For maps with hundreds of landmarks, exhaustive matching of maps is not an option.

A. Summary of Results

This paper presents a computationally tractable approach to the problem of combining stochastic maps with unknown

coordinate systems possibly containing unknown common landmarks. Our approach is based on the maximum likelihood principle and partitions the optimal map combining problem into two separate steps: *matching* and *combining*. The matching step identifies the common map (if any) between the two stochastic maps by matching a subset of estimated landmarks in one map with those of a different map. Once the common landmarks are identified, the conceptually simpler combining step finds the best translation and rotation operations to align the maps in an optimal fashion.

Finding the best match of subsets of estimated landmarks is combinatorial and nontrivial. Our contribution is a formulation of this problem as a bipartite matching of Delaunay triangles with Groth representation. Such a bipartite matching problem can be solved efficiently by a linear program, which yields a scalable solution. In the proposed approach, the cost function used in measuring the quality of matching is derived from the generalized likelihood ratio test (GLRT). To this end, we obtain the closed form expression for the generalized likelihood ratio that is trivial in computation.

In this paper, we propose a maximum likelihood (or *least squares* under the Gaussian noise assumption) approach to combining common landmarks. As in the case of computing the generalized likelihood ratio statistics, the optimal rotation and translation operations can also be obtained in closed form, making the combining step also trivial in computation.

B. Related work

The problem considered in this paper has application to many different fields, such as robotics, image processing, and pattern recognition. Most relevant to this work is that of Groth [4], who considered the problem of identifying stars in images from a star tracker with those of an existing catalog. Groth proposed a special technique based on an exhaustive matching of triangles. The algorithm presented in this paper is inspired by Groth's approach in that we also formulate the common landmark identification problem as one of matching of triangles. Our contribution, however, lies in the formulation of bipartite matching of the Delaunay graphs, the use of generalized likelihood ratio statistics, and the closed form computations of test statistics and map combining parameters.

The matching part of our approach can be viewed as as a special form of data association [5], [6], [7]. The standard data association problem in the context of single robot SLAM is to match a measurement with a set of existing landmarks. This problem has been considered by many authors, but

perhaps the most relevant work is the linear programming approach proposed by Zhang, Xie and Adams [8]. Our matching approach also involves linear programing, but in quite a different context and using different cost functions.

There is a great deal of existing work on obtaining the optimal alignment of two images once the common parts of the images are identified. To our knowledge, no closed-form solution in terms of elementary functions existed in the literature, although solutions based on the singular value decompositions (SVD) have been obtained [9]. See also [10], [11]. Perhaps the most closely related work is that of Thrun and Liu [12] due to its direct application to mobile robot navigation. In their work, an iterative hill-climbing approach is proposed as a solution to the correspondence and alignment problems.

II. BIPARTITE MATCHING

We now consider the problem of identifying common landmarks in the maps of two separate robots. The maps of the robots are referred to as map p and map q. The identification process includes three parts: (i) Delaunay triangulation and Groth representation; (ii) generalized likelihood ratio test (GLRT) statistics; (iii) bipartite matching of triangles.

A. Delaunay triangulation and Groth representation

We assume that there are two robots, p and q. Robot p has the stochastic map X_p which is a random vector including locations of all its estimated landmarks. The stochastic map X_q for robot q is similarly defined. Each robot can then generate its own Delaunay triangulation, from which it obtains a list of triangles. Each triangle is made of three vertices that are placed into a six dimensional vector Y following the Groth representation of triangles [4]. Notice that Delaunay triangulation is almost surely unique but the noisy versions of the same map may have quite different Delaunay graphs.

B. GLRT statistics

For a triangle in map p with Groth vector representation Y_p and a triangle in map q with Groth vector representation Y_q, we consider the following binary hypothesis H_0 vs. H_1. Under hypothesis H_0, Y_p and Y_q do not correspond to the same three landmarks in truth. Under the Gaussian measurement assumption, we have

$$H_0 : \begin{bmatrix} Y_p \\ Y_q \end{bmatrix} \sim \mathcal{N}\left(\begin{bmatrix} \mu_p \\ \mu_q \end{bmatrix}, \begin{bmatrix} \Sigma_p & \\ & \Sigma_q \end{bmatrix}\right) \quad (1)$$

where we assume that $\mu_p, \mu_q \in R^6$ are unknown location vectors and $\Sigma_p = \Sigma_q = \sigma^2 I_6$ are noise covariance matrices with known σ. (We use I_k to denote a $k \times k$ identity matrix.) Under hypothesis H_1, Y_p and Y_q are noisy measurements of the same three landmark in truth but the actual locations of the three landmarks are deterministic unknowns. In particular,

$$H_1 : \begin{bmatrix} Y_p \\ Y_q \end{bmatrix} \sim \mathcal{N}\left(\begin{bmatrix} \mu \\ R(\theta)\mu + Ft \end{bmatrix}, \begin{bmatrix} \Sigma_p & \\ & \Sigma_q \end{bmatrix}\right) \quad (2)$$

where $\mu \in R^6$ is the unknown vector of the three landmark locations (in Groth representation), θ the unknown rotation of map q, $t \in R^2$ the unknown translation of map q, and $F = [I_2 \ I_2 \ I_2]^T \in R^{6\times 2}$. Again we assume $\Sigma_p = \Sigma_q = \sigma^2 I_6$ with known σ.

The generalized likelihood ratio test (GLRT) is given by

$$\Lambda(y_p, y_q) = \frac{\max\limits_{\mu, t, \theta} f_1(y_p, y_q; \mu, t, \theta)}{\max\limits_{\mu_p, \mu_q} f_0(y_p, y_q; \mu_p, \mu_q)} \underset{H_0}{\overset{H_1}{\gtrless}} \tau \quad (3)$$

where τ is selected to control the level of false alarm. The likelihood function under H_0 is given by

$$f_0(y_p, y_q; \mu_p, \mu_q) \propto \exp\left\{-\frac{1}{2\sigma^2} J_0(\mu_p, \mu_q; y_p, y_q)\right\}$$

where

$$J_0(\mu_p, \mu_q; y_p, y_q) = ||y_p - \mu_p||^2 + ||y_q - \mu_q||^2. \quad (4)$$

The optimization under H_0 (3) gives $\mu_p^* = y_p$ and $\mu_q^* = y_q$.

The optimization under H_1 in (3), on the other hand, involves a more complicated likelihood function

$$f_1(y_p, y_q; \mu, t, \theta) \propto \exp\left\{-\frac{1}{2\sigma^2} J_1(\mu, t, \theta; y_p, y_q)\right\}$$

where

$$J_1(\mu, t, \theta; y_p, y_q) = ||y_p - \mu||^2 + ||y_q - R(\theta)\mu - Ft||^2. \quad (5)$$

Since J_1 is non-convex, the ability to obtain a tractable solution is essential to our objective of having a computationally scalable solution. The following theorem provides a closed-form solution to the optimization under H_1.

Theorem 1: The ML solution (θ^*, μ^*, t^*) under H_1 is given by

$$\theta^* = \text{sgn}(\beta)\left[\cos^{-1}\left(\frac{\alpha}{\sqrt{\alpha^2 + \beta^2}}\right) - \pi\right] \quad (6)$$

$$t^* = \frac{1}{3}F^T\left[y_q - R(\theta^*)y_p\right] \quad (7)$$

$$\mu^* = \frac{1}{2}R^T(\theta^*)\left[Q^+ R(\theta^*)y_p + Q^- y_q\right] \quad (8)$$

where

$$Q^+ = I_6 + \frac{1}{3}FF^T, Q^- = I_6 - \frac{1}{3}FF^T, \quad (9)$$

$$I_c = \begin{bmatrix} 1 & 0 \\ 0 & 1 \end{bmatrix} \quad \text{and} \quad I_s = \begin{bmatrix} 0 & -1 \\ 1 & 0 \end{bmatrix} \quad (10)$$

$$\alpha = -y_q^T(I_3 \otimes I_c)Q^- y_p \quad (11)$$

$$\beta = -y_q^T(I_3 \otimes I_s)Q^- y_p \quad (12)$$

$$\gamma = \frac{1}{2}y_p^T Q^- y_p + \frac{1}{2}y_q^T Q^- y_q. \quad (13)$$

With the above theorem, the generalized likelihood ratio statistics, the left hand side of (3), can be computed in closed form.

C. The ML Matching of Bipartite Graph

We now formulate the problem of identifying matching triangles as a bipartite matching problem. Given the set of Delaunay triangles $\{y_p^i : i = 1, \cdots, N_p\}$ from map p and $\{y_q^i : i = 1, \cdots, N_q\}$ from map q (in Groth representation), we consider the following linear program

$$\text{maximize} \quad \sum_{i=1}^{N} \sum_{j=1}^{N} \Lambda(y_p^i, y_q^j) x_{ij} \tag{14a}$$

$$\text{subject to} \quad \sum_{i=1}^{N} x_{ij} = 1, \quad j = 1, \ldots, N \tag{14b}$$

$$\sum_{j=1}^{N} x_{ij} = 1, \quad i = 1, \ldots, N \tag{14c}$$

$$\text{and} \quad 0 \leq x_{ij} \leq 1 \tag{14d}$$

where $\Lambda(y_p^i, y_q^j)$ is the generalized likelihood ratio statistic computed from Theorem 1 and $N = \max(N_p, N_q)$. It is well known that the solution of the above linear program are integer $(0-1)$ valued x_{ij} solutions. If $x_{ij} = 1$, we consider triangle y_p^i and y_q^j are tentatively matched. Thus the above linear program provides a list of tentatively matched triangles, one from each stochastic map. (When $N_p \neq N_q$, there will be unmatched triangles.) From this list, we then select a confirmed list of matches by applying a threshold test on the likelihood ratio statistics. Those triangle pairs with sufficiently high GLRT are considered matched and declared as corresponding to common landmarks. Those pairs with low GLRT are considered independently observed and not part of the common map.

The probabilities of missed detection and false alarm of course depend on the threshold τ used in the GLRT test. Because the matched landmarks are used to combine stochastic maps as described in the next section, it is important to choose τ sufficiently high in order to avoid false positive declarations.

III. Maximum likelihood map combining

From the set of matched triangles that pass the GLRT, we obtain the set $\mathcal{K}$ of matched landmarks and the corresponding estimates $\{y_p^i : i \in \mathcal{K}\}$ in map p and $\{y_q^i : i \in \mathcal{K}\}$ in map q. Using Theorem 1 with a generalization to n landmarks, we can then solve the best (posterior) common landmarks μ^*, the best (posterior) rotation θ^*, and the best (posterior) translation t^*. The maximum likelihood combined map is then given by

$$\mu_{\text{ML}}^* = \begin{bmatrix} \mu^* \\ \mu_{p|0}^* \\ R^T(\theta^*)(\mu_{q|0}^* - Ft^*) \end{bmatrix} \tag{15}$$

where $\mu_{p|0}^*$ is the ML estimate of landmarks under H_0 in map p and $\mu_{q|0}^*$ the ML estimate of landmarks under H_0 in map q.

IV. Simulation

Two 100 landmark maps with additive zero-mean Gaussian noise and 75 common landmarks are shown in Fig. 1. The threshold τ is chosen to minimize the number of false positive declarations. In each case, the signal-to-noise ratio is $10 \log(\bar{\sigma}^2 / 2\sigma^2)$ where $\bar{\sigma}^2$ is the sample variance of the edges lengths in the ground truth Delaunay graph and σ^2 is the variance of the additive noise. The example provides an idea of how the performance of our approach degrades with noise. With increasing SNR, the result indicates an increase in missed detections (incorrect declarations of H_0) and a decrease in true positives (correct declarations of H_1). At an SNR of 40 dB, there are 23 true positives (illustrated in green) and 83 missed detections (illustrated in blue). At 30 dB, there are 5 true positives and 62 missed detections. In each case, zero false positives were observed.
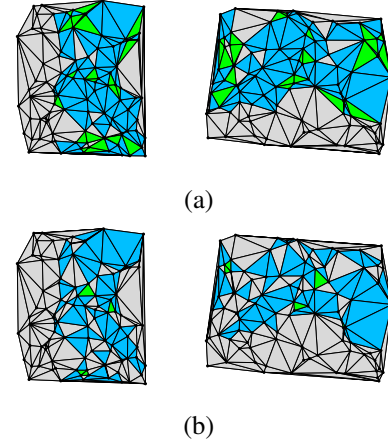

(a)

(b)

Fig. 1. Matching example with 100 landmarks in each map. Correct declarations of H_1 are illustrated by green triangles. Incorrect declarations of H_0 (missed detections) are illustrated in blue. Gaussian noise is added to each map with an SNR of (a) 40 dB and (b) 30 dB.

The combined maps are shown below in Fig 2. This result indicates that the maximum likelihood combined map is consistent with ground truth in the absence of false positives (we note, however, that false positives may produce undesired results in the combined map).
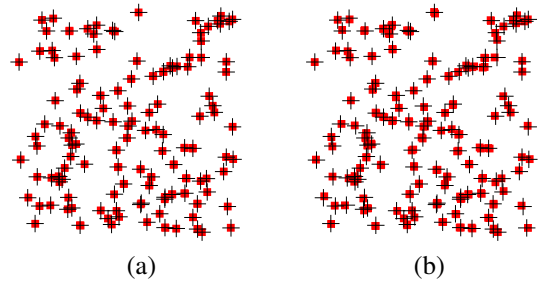

(a)   (b)

Fig. 2. Maximum likelihood combining. The SNR associated with the combined maps are (a) 40 dB and (b) 30 dB. This result demonstrates that duplicate entries may be introduced in the combined map due to missed detections (an issue that may be resolved using standard data association techniques). Cross hairs indicate ground truth and the red squares are entries of the combined map.

Fig. 3 shows the performance of our approach for repeated experiments of overlap and noise scenarios. In each scenario, the performance of our approach degrades gracefully with increasing SNR. The percentages of 75% and 50% common landmarks in each map are considered. At 75%, the results for the 10 and 20 landmark scenarios are relatively the same. However at 50%, a degrade in performance is observed. This result suggests that greater overlap may be required to maintain performance as the number of landmarks in each map is increased.
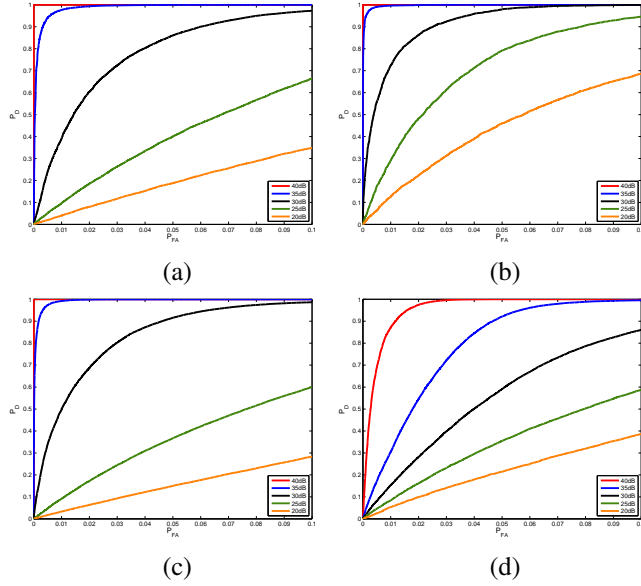

(a) (b) (c) (d)

Fig. 3. Monte Carlo receiver operating characteristics. The following scenarios are considered: (a)10 landmarks in each map with 75% overlap, (b) 10 landmarks in each map with 50% overlap, (c) 20 landmarks in each map with 75% overlap and (d) 20 landmarks in each map with 50% overlap. Each curve is computed based on 10,000 experiments with various SNR.

V. CONCLUSION

We considered in this paper the problem of combining stochastic maps from independent agents who explore possibly overlapping areas. The proposed optimization framework is based on the maximum likelihood principle. While our solution is suboptimal, the proposed approach has significant computational advantage as a scalable algorithm for the maximum likelihood combining of maps.

REFERENCES

[1] R. Smith, M. Self and P. Cheeseman, "Estimating uncertain spatial relationships in robotics," *Autonomous Robot Vehicles*, pp. 167–193, Springer-Verlag, 1990.
[2] G. Dissanayake, P. Newman, H.F. Durrant-Whyte, S. Clark and M. Csobra, "A solution to the simultaneous localisation and mapping (SLAM) problem," *IEEE Transactions on Robotics and Automation*, Vol. 17, No. 3, pp. 229–241, 2001.
[3] H. Durrant-Whyte and T. Bailey, "Simultaneous localization and mapping (SLAM): Part I," *IEEE Robotics and Automation Magazine*, Vol. 13, No. 2, pp. 99–110, 2006.
[4] E. Groth, "A pattern-matching algorithm for two-dimensional coordinate lists," *The Astronomical Journal*, Vol. 91, No. 5, pp. 1244–1248, 1986.
[5] T. Bailey and H. Durrant-Whyte, "Simultaneous localization and mapping (SLAM): Part II," *IEEE Robotics and Automation Magazine*, Vol. 13, No. 3, pp. 108–117, 2006.
[6] J. Neira and J. Tardós, "Data association in stochastic mapping using the joint compatibility test," *IEEE Transactions on Robotics and Automation*, Vol. 17, No. 6, pp. 890–897, 2001.
[7] Y. Bar-Shalom and T. Fortman, *Tracking and data association*, Academic Press, Boston, MA; 1988.
[8] S. Zhang, L. Xie and M. Adams, "An efficient data association approach to simultaneous localization and map building," *The International Journal of Robotics Research*, Vol. 24, No. 1, pp. 49-60, 2005.
[9] K. Arun, T. Huang and S. Blostein, "Least-squares fitting of two 3D point sets," *IEEE Transactions on Pattern Analysis and Machine Intelligence*, Vol. 9, No. 5, pp. 698–700, 1987.
[10] B. Horn, "Closed-form solution of absolute orientation using unit quaternions," *Journal of the Optical Society of America A*, Vol. 4, No. 4, pp. 629–642, 1987.
[11] S. Umeyama, "Least squares estimation of transformation parameters between two point sets," *IEEE Transactions on Pattern Analysis and Machine Intelligence*, Vol. 13, No. 4, pp. 376–380, 1991.
[12] S. Thrun and Y. Liu, "Multi-robot SLAM with sparse extended information filters," *Proceedings of the 11th Int. Symposium of Robotics Research*. Sienna, Italy: Springer 2003.