Maximum Likelihood Fusion of Stochastic Maps

Brandon M. Jones, *Member, IEEE*, Mark Campbell, *Member, IEEE*, and Lang Tong, *Fellow, IEEE*

***Abstract*—The fusion of independently obtained stochastic maps by collaborating mobile agents is considered. The proposed approach includes two parts: generalized likehood ratio matching and maximum likelihood alignment. In particular, an affine invariant hypergraph model is constructed for each stochastic map and a bipartite matching via a linear program is used to establish landmark correspondence between stochastic maps. A maximum likelihood alignment procedure is proposed to estimate rotation, translation and scale parameters in order to construct a global map of the environment. A main feature of the proposed approach is its scalability with respect to the number of landmarks: the matching step has polynomial complexity and the maximum likelihood alignment solution is obtained in closed form. Experimental validation of the proposed fusion approach is performed using the Victoria Park experimental benchmark.**

***Index Terms*—Data association, data fusion, hypothesis testing, maximum likelihood estimation, mobile robot navigation.**

I. Introduction

THE general nature of the problem under consideration is to construct a global map of landmarks from the individual efforts of collaborating agents that operate outside of a global frame of reference. Each agent independently builds a vector of estimated landmark locations referred to as a *stochastic map* [1], [2]. Constructing a combined global map within a common reference frame from the individual maps of the agents is referred to as a problem of *fusion of stochastic maps*. Intuitively, the problem has the interpretation of a mathematical jigsaw puzzle: the stochastic maps are the disoriented pieces and the sought after global map is the completed puzzle.

A benchmark scenario based on the Victoria Park dataset is illustrated in Fig. 1. The satellite image shows the ground truth environment from which the stochastic maps are obtained. Trees (landmarks) located in the park are mapped from the global reference frame of the environment to the individual local reference frames of the mobile agents. Each agent thus has an independent, but partial model of the explored environment. The shared objective of the agents is to reconstruct the state of nature from the sensor measurements independently obtained by each agent, which is a common problem encountered in signal processing [3]–[7], computer vision and robotics.

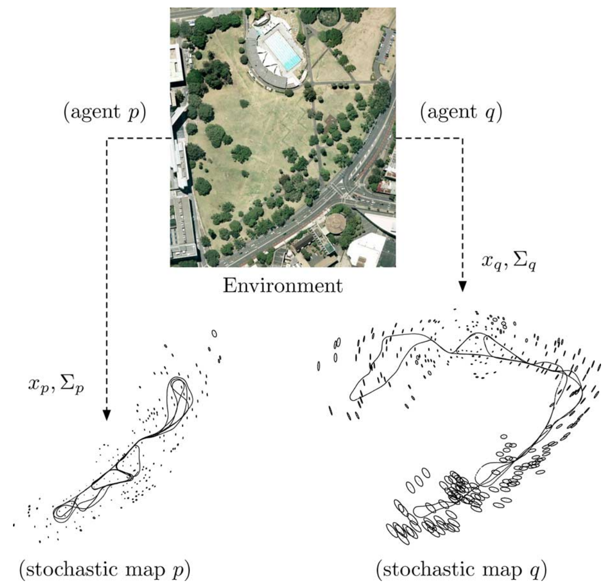

Fig. 1. A data fusion problem involving stochastic maps. The stochastic maps of agent p (bottom left) and agent q (bottom right) illustrate landmark locations estimated from the sensor observations of two mobile agents operating outside of a global reference frame. Uncertainties in estimation are indicated by ellipses and the path of exploration is shown by dotted lines. Estimation of a common global map from the individual stochastic maps, each obtained in a separate coordinate system, requires inferring common landmarks in addition to determining a common frame of reference.

Stochastic maps are obtained by independent agents using various estimation techniques. In robotics, the solution to the *simultaneous localization and mapping (SLAM)* problem provides an agent with a stochastic map of the environment as a model of landmark locations (see [1], [2], [8]–[10] and the references therein). The focus of this paper, however, is on the fusion – rather than building – of stochastic maps obtained in separate coordinate systems. Our starting point is at the individual stochastic maps, which are made available to a fusion agent for the construction of a global map.

The fusion problem with multiple sensor observations is challenging for several reasons, one being that the problem contains both discrete and continuous parts [11], [12]. In order to construct a global map, the fusion agent must first identify common landmarks residing in two separate maps. Using the

Manuscript received May 04, 2013; revised October 11, 2013 and December 24, 2013; accepted January 12, 2014. Date of publication February 04, 2014; date of current version March 19, 2014. The associate editor coordinating the review of this manuscript and approving it for publication was Prof. Francesco Verde. This work was supported in part by the Army Research Office under grant W911NF-10-1-0419 and the National Science Foundation under award CCF-1018115. A portion of this work was presented at the 49th Annual Allerton Conference on Communication, Control, and Computing, Monticello, IL, USA, September 2011, and at the SPIE Defense and Security Symposium, Baltimore, MD, USA, May 2012.

B. M. Jones and L. Tong are with the School of Electrical and Computer Engineering, Cornell University, Ithaca, NY 14853 USA (e-mail: bmj34@cornell.edu; ltong@ece.cornell.edu).

M. Campbell is with the Sibley School of Mechanical and Aerospace Engineering, Cornell University, Ithaca, NY 14853 USA (e-mail: mc288@cornell.edu).

Color versions of one or more of the figures in this paper are available online at http://ieeexplore.ieee.org.

Digital Object Identifier 10.1109/TSP.2014.2304435

earlier jigsaw analogy, the solver has to first identify common edges in order to match the individual pieces. Prior to exchanging stochastic maps, the agents are assumed to operate with no prior knowledge concerning the common landmarks (i.e., the common trees when considering the Victoria Park example) that are contained within the individual maps. The problem of matching common landmarks is of a combinatorial nature in general, which eliminates exhaustive search as an option for large maps.

Even if common landmarks between two maps have been identified, the agents are faced with the *alignment* problem of determining not only the best landmark estimates of common and uncommon landmarks contained by noisy maps obtained in separate coordinate systems, but also to determine the spatial parameters of rotation, translation and scale. Describing this again in terms of the earlier jigsaw analogy: not only are the pieces disoriented, but the edges are also imprecise (which makes it harder to see how the pieces fit together). The alignment optimization is continuous in nature, but is also nonlinear and non-convex in general.

A. Related Work

The data fusion problem considered in this paper has been studied in various forms. Zhou *et al.* [3] derived a two-step iterative optimization procedure for estimating the locations of common targets using range and azimuth measurements obtained by two separate radar sensors. Each sensor observes the locations of targets within a reference frame related by a known displacement. The framework proposed in this paper considers the more general case that the sensor reference frames are related by an unknown displacement, rotation and scale as part of a nonlinear least squares optimization, the solution of which is obtained in closed form. Thrun and Liu [11] proposed an SR-tree (Sphere/Rectangle-tree) search [13] in consideration of the matching problem. Common landmark correspondences and rotation-translation parameters are found using an iterative hill climbing approach to match triplet combinations formed within a small radius of the landmarks in each map. The radius forming the feature vectors of the SR-tree, however, would need to be adaptive in order to generalize to different environments and scales. Estimates of common landmarks are determined separately by a collapsing operation performed on matched landmarks in information form (see Grime and Durrant-Whyte [14], as well as Sukkarieh *et al.* [15], for further reading on fusion using information filtering). Julier and Uhlmann [16] introduced the *covariance intersection* algorithm as an approach to the data fusion problem. Their algorithm uses a convex combination of state information to achieve data fusion, but has the limitation that the input data must be of the same dimension (which is often not the case of stochastic maps built within different regions of exploration). Tardós *et al.* [17], and later Castellanos *et al.* [18], proposed *map joining* as a technique to enable an individual mobile robot to construct a global stochastic map based on a sequence of local maps. The approach is related to this paper by considering the sequence of local maps as being obtained from separate robots, but requires knowledge of a base reference to construct a global map.

Williams *et al.* [19] considered the fusion problem by providing parameter estimates of the relative rotation and translation between global and local maps. The expressions are derived by observing the geometry of the landmarks within each map. Our approach is distinct from [19] in that the geometry of the landmarks is incorporated in a nonlinear least squares solution based on the maximum likelihood principle. Several authors such as Zhou and Roumeliotis [20], Andersson and Nygards [21], Benedettelli *et al.* [22] and Aragues *et al.* [23] considered rendezvous approaches to the alignment problem. Rendezvous approaches, however, are somewhat restrictive as the agents are required to be in close proximity.

The matching approach of this paper is motivated by the work of Groth [24] and Ogawa [25]. Groth proposed one of the earliest matching algorithms in the context of astronomical point patterns, where a list of star measurements are matched against a known star catalog. In the proposed approach, structured point triplets referred to simply as triangles are used to match the measurements against the catalog. The Groth triangle convention is also incorporated in our approach, however the matching approach of Groth is not practical for large maps since all possible combinations of triangles are considered. An alternative approach was proposed by Ogawa [25], which instead incorporated *Delaunay triangulations* [26] to address the star matching problem. This paper therefore uses Delaunay triangulations with triangles that follow the Groth convention as a graphical model for matching stochastic maps.

B. Summary of Results and Organization

A maximum likelihood framework is proposed for the construction of a global map from local stochastic maps. The proposed approach includes (i) a landmark matching approach referred to as *generalized likelihood ratio (GLR) matching* and (ii) a least squares approach for jointly estimating rotation, translation, scale and common landmark locations referred to as *maximum likelihood (ML) alignment*. A Gaussian likelihood function is presented as the main proxy for deriving the procedures of each step.

Matching is a step that is performed in the absence of a global frame of reference, which requires a technique that is affine invariant. To this end, the original stochastic maps are represented as directed hypergraphs constructed from Delaunay triangulations. The hyperedges of each directed hypergraph are constructed from directed Delaunay triangles that follow the Groth convention, which leads to an affine invariant approach for determining common landmarks. The GLR matching algorithm uses a generalized likelihood ratio as a matching metric in order to obtain globally optimal landmark correspondences from the solution of a bipartite matching problem. The GLR metric is computed in closed form and the bipartite matching is solved in polynomial time as a solution to a linear program.

Once common landmarks are identified, the solution to the alignment problem of determining rotation, translation, scaling and common landmark locations between two stochastic maps is computed from the determined matched landmarks. The main contribution is a closed-form solution to the alignment problem as nonlinear non-convex optimization, which makes optimal alignment trivial to obtain computationally.

The remainder of the paper is organized as follows. The problem formulation and models used throughout the paper are provided in Section II. While the alignment and matching steps share common likelihood functions, the maximum likelihood alignment problem is presented first in Section III in order to introduce the proposed solution for closed form computations. The problem of determining common landmarks is treated in Section IV, where we present the GLR matching approach. Numerical examples and simulations are provided in Section V. The conclusion is given in Section VI and is followed by an appendix of proofs.

II. Model and Problem Formulation

A. Ground Truth Model of Landmarks

A landmark is represented by a vector in $\mathbb{R}^2$ under a specific coordinate system. Two collaborating agents p and q each estimate the locations of landmarks within a local frame of reference. The coordinate systems of p and q are related by a scaling parameter $s > 0$, a rotation with parameter $\theta \in [-\pi, \pi]$ and a translation parameter $t \in \mathbb{R}^2$. Specifically, if $\mu \in \mathbb{R}^2$ is the location of a landmark under coordinate system p, the landmark location under coordinate system q is then

$$\mu' = sr(\theta)\mu + t, \;\; r(\theta) \triangleq \begin{bmatrix} \cos\theta & -\sin\theta \\ \sin\theta & \cos\theta \end{bmatrix}.$$

In the general case of m landmark locations $\mu \in \mathbb{R}^{2m}$, again in coordinate system p, the representation in coordinate system q is of the form

$$\mu' = sR(\theta)\mu + Ft$$

where $R(\theta) \triangleq I_m \otimes r(\theta)$ is the rotation matrix in block diagonal form and I_m is an $m \times m$ identity matrix (the symbol $\otimes$ is the Kronecker product operator). The matrix $F \triangleq e_m \otimes I_2$, with e_m being an m-vector with all entries equal to 1, applies the translation t to each landmark in the map.

In this paper, without loss of generality, the ground truth of agent p is defined in the coordinate system of a ground truth vector of the form $u = (\mu^T, v_p^T, v_q^T)^T$, specified as follows: common landmarks observed by both agents are contained by $\mu \in \mathbb{R}^{2n}$, where n is the number of common landmarks, and landmarks observed exclusively by agent p and exclusively by agent q are contained by the vectors v_p and v_q, respectively. Agent p and agent q thus observe ground truth of the form $u_p = (\mu^T, v_p^T)^T$ and $u_q = (\mu^T, v_q^T)^T$, respectively, with the representation of u_q in coordinate system q given by

$$\begin{aligned} u_q' &= sR(\theta)u_q + Ft \\ &= s\begin{pmatrix} R_1(\theta) & \\ & R_0(\theta) \end{pmatrix}\begin{pmatrix} \mu \\ v_q \end{pmatrix} + \begin{pmatrix} F_1 \\ F_0 \end{pmatrix} t \end{aligned} \quad (1)$$

where the subscripts 1 and 0 indicate the partition the map into common and uncommon parts.

In addition to the inherent uncertainty of stochastic maps, the challenge of fusion as it relates to constructing a combined map is that the parameters $\{s, t, \theta\}$ and the common landmarks observed by both agents are usually also unknown. The ultimate goal of fusion is thus to estimate the combined map from the stochastic maps of the individual agents without prior knowledge of the model parameters $\{\mu, s, t, \theta\}$.

III. Maximum Likelihood Alignment

This section describes a maximum likelihood approach for constructing a global map of landmarks from stochastic maps obtained in separate coordinate systems. We begin by describing a Gaussian model for the maps and propose a closed form solution to the maximum likelihood alignment problem of estimating the parameters $\{\mu, v_p, v_q, s, t, \theta\}$ under the assumption that the common landmarks between the maps are known (the more general case of unknown common landmarks is considered in Section IV).

A. Matched Gaussian Maps

Let the random vectors X_p and X_q represent the noisy observations obtained by agents p and q, respectively. Prior to fusion, data is collected in the separate coordinate systems of the agents (i.e., X_p and X_q reside in coordinate systems p and q, respectively). The statistical model of *matched Gaussian maps* is given as

$$X_p = u_p + W_p \quad (2)$$
$$X_q = sR(\theta)u_q + Ft + W_q \quad (3)$$

where $W_p \sim \mathcal{N}(0, \sigma_p^2 I)$ and $W_q \sim \mathcal{N}(0, \sigma_q^2 I)$ are independent zero-mean additive Gaussian noise vectors. In order to separate the matching and alignment problems, an assumption is made that the common landmarks in both maps are known. The process of obtaining such a matching, however, is nontrivial and combinatorial in general (see Section IV for an affine invariant procedure for determining common landmarks).

B. Likelihood Decomposition and Closed Form Solution

Estimators of the parameters $\{\mu, v_p, v_q, s, t, \theta\}$ are derived by considering the likelihood function of the combined global map given by

$$L(\mu, v_p, v_q, s, t, \theta) \triangleq \eta \exp -\frac{1}{2} J(\mu, v_p, v_q, s, t, \theta) \quad (4)$$

where η is a normalizing constant and the function J, with unknown parameters as its arguments, is defined as

$$J(\mu, v_p, v_q, s, t, \theta) \triangleq \frac{1}{\sigma_p^2}\|x_p - u_p\|^2 + \frac{1}{\sigma_q^2}\|x_q - sR(\theta)u_q - Ft\|^2 \quad (5)$$

with $x_p = (x_p^{1T}, x_p^{0T})^T$ and $x_q = (x_q^{1T}, x_q^{0T})^T$, corresponding to the structure of u_p and u_q, respectively. By partitioning the problem into common and uncommon parts, it immediately follows that (5) decomposes as

$$J(\mu, v_p, v_q, s, t, \theta) = J_0(v_p, v_q, s, t, \theta) + J_1(\mu, s, t, \theta) \quad (6)$$

where J_0 is the squared error function parameterized by the uncommon landmarks v_p and v_q, in addition to the transform parameters $\{s, t, \theta\}$, specified as

$$J_0(v_p, v_q, s, t, \theta) \triangleq \frac{1}{\sigma_p^2}\|x_p^0 - v_p\|^2 + \frac{1}{\sigma_q^2}\|x_q^0 - sR_0(\theta)v_q - F_0 t\|^2 \quad (7)$$

and J_1 is the squared error function parameterized by the vector μ of common landmarks, in addition to $\{s, t, \theta\}$, specified as

$$J_1(\mu, s, t, \theta) \triangleq \frac{1}{\sigma_p^2}\|x_p^1 - \mu\|^2 + \frac{1}{\sigma_q^2}\|x_q^1 - sR_1(\theta)\mu - F_1 t\|^2. \quad (8)$$

This decomposition is exploited to minimize the combined error function J by minimizing J_0 and J_1 separately as follows. Let the solution $\{\mu^*, v_p^*, v_q^*, s^*, t^*, \theta^*\}$ be the global maximum of the likelihood function L so that

$$J(\mu^*, v_p^*, v_q^*, s^*, t^*, \theta^*) = \min_{\mu, v_p, v_q, s, t, \theta} J(\mu, v_p, v_q, s, t, \theta). \quad (9)$$

It follows that if $\{\hat{\mu}, \hat{s}, \hat{t}, \hat{\theta}\}$ is the global minimum of J_1 specified as

$$(\hat{\mu}, \hat{s}, \hat{t}, \hat{\theta}) = \operatorname*{argmin}_{\mu, s, t, \theta} J_1(\mu, s, t, \theta), \quad (10)$$

then the optimal (least squares) parameters corresponding to $\{\mu, s, t, \theta\}$ are given by $\mu^* = \hat{\mu}$, $s^* = \hat{s}$, $t^* = \hat{t}$, $\theta^* = \hat{\theta}$ since

$$J_0\left(x_p^0, \frac{1}{s}R_0^T(\theta)(x_q^0 - F_0 t), s, t, \theta\right) = 0 \quad (11)$$

for any $\{\mu, s, t, \theta\}$ with the decomposition $J = J_0 + J_1$. It immediately follows that the least squares estimators of v_p and v_q are given by

$$v_p^* = x_p^0, \quad v_q^* = \frac{1}{\hat{s}}R_0^T(\hat{\theta})(x_q^0 - F_0\hat{t}) \quad (12)$$

respectively. Thus the combined map $u^* = (\mu^{*T}, v_p^{*T}, v_q^{*T})^T$ is obtained by minimizing the error function J_1, which is a nonlinear and non-convex function of the parameters $\{\mu, s, t, \theta\}$. However, as specified by the following theorem, the structure of the problem admits a closed form solution for obtaining the global minimum.

Theorem 1 (Closed Form MLE): The ML estimators of the parameters $\{\mu, s, t, \theta\}$ are given by the following expressions.

1) The MLE of the rotation parameter θ is

$$\theta^* = \operatorname{sgn}(\beta)\left[\cos^{-1}\left(\frac{\alpha}{\sqrt{\alpha^2 + \beta^2}}\right) - \pi\right] \quad (13)$$

where $\operatorname{sgn}(\cdot)$ is the signum function. The coefficients α and β are given by

$$\alpha = -x_q^{1T}(I_n \otimes I_c)Qx_p^1 \quad (14)$$

$$\beta = -x_q^{1T}(I_n \otimes I_s)Qx_p^1 \quad (15)$$

respectively, where $Q = I_{2n} - F_1(F_1^T F_1)^{-1}F_1^T$ with n being the number of common landmarks (see Appendix for the constant matrices I_c and I_s).

2) The MLE of the scale factor s is

$$s^*(\theta^*) = \frac{x_p^{1T} Q R_1^T(\theta^*) x_q^1}{x_p^{1T} Q x_p^1} \quad (16)$$

denoted hereafter as s^* for simplicity.

3) The MLE of the translation t is

$$t^*(s^*, \theta^*) = (F_1^T F_1)^{-1} F_1^T\left[x_q^1 - s^* R_1(\theta^*) x_p^1\right] \quad (17)$$

denoted hereafter as t^*.

4) The MLE of the common landmarks μ is

$$\mu^*(\theta^*) = \phi_p^* x_p^1 + \phi_q^* x_q^1 \quad (18)$$

denoted hereafter as μ^*. The matrix gains ϕ_p^* and ϕ_q^* are given by

$$\phi_p^* = I_{2n} - \frac{(\sigma_p s^*)^2}{\sigma_p^2 + \sigma_q^2} Q \quad (19)$$

$$\phi_q^* = s^* \frac{\sigma_p^2}{\sigma_p^2 + \sigma_q^2} Q R_1^T(\theta^*) \quad (20)$$

respectively.

Proof: See Appendix. ■

Theorem 1 specifies a closed form solution to the ML alignment problem using the realizations of matched Gaussian maps as data. An important note, however, is that $Q = 0_{2\times 2}$ when $n = 1$, meaning $n > 1$ common landmarks are required to compute the solution of Theorem 1.

IV. Generalized Likelihood Ratio Matching

In a general mapping scenario, the ground truth structure observed by the agents is unknown. In particular, if the first two entries of $X_p = x_p$ correspond to the particular landmark, then the first two entries of $X_q = x_q$ correspond to a different landmark in general (and likewise with the remaining entries). Common landmarks in this case are identified by applying a matching procedure to x_p and x_q with consideration that the stochastic maps are obtained in separate coordinate systems related by s, t and θ. The matching procedure proposed in this section is based on the use of landmark triplets referred to as *triangles*, which requires that the maps of each agent contain at least three landmarks.

A. Directed Hypergraph Model

Triangles are constructed from the maps of each agent by following a direction convention used in the star-pattern matching approach of Groth [24]. In particular, given three landmark locations $y_a, y_b, y_c \in \mathbb{R}^2$, the Groth representation of a directed triangle is $y = (y_a^T, y_b^T, y_c^T)^T$, which is a vector in $\mathbb{R}^6$ with entries that follow the inequality

$$\|y_a - y_b\| < \|y_b - y_c\| < \|y_c - y_a\| \tag{21}$$

under the assumption that no two triangle edges have the same length. This convention, which is invariant to changes in rotation and translation, is used to construct a *directed hypergraph* model given by the pair $G = (\mathcal{V}, \mathcal{E})$, where $\mathcal{V}$ is a set of *vertices* (e.g., landmarks) and $\mathcal{E} = \{E_1, E_2, \ldots, E_m\}$ is a collection of *directed hyperedges* $E_i \subseteq \mathcal{V}$ consisting of ordered subsets of $\mathcal{V}$. In particular, each hyperedge E_i with $i = 1, 2, \ldots, m$ is such that $|E_i| = 3$ with an ordering that follows the inequality (21). Each agent constructs a directed hypergraph model of the form $G_p = (\mathcal{V}_p, \mathcal{E}_p)$ and $G_q = (\mathcal{V}_q, \mathcal{E}_q)$ using Delaunay triangulations [26] constructed from the landmarks of maps p and q, respectively.

B. Hypothesis Testing and Bipartite Matching

Determining common landmarks from the directed triangles of G_p and G_q is considered as a binary hypothesis testing problem. Under hypothesis H_0, the agents observe the ground truth directed triangles $\nu_p, \nu_q \in \mathbb{R}^6$, which contain a maximum of two landmarks in common. Under hypothesis H_1, the agents observe a common directed triangle $\delta \in \mathbb{R}^6$ within their respective coordinate systems. In this way, H_0 is the hypothesis of uncommon triangles and H_1 is the hypothesis of common triangles. The mathematical models of H_0 and H_1 are given by

$$H_0 : \begin{bmatrix} Y_p \\ Y_q \end{bmatrix} \sim \mathcal{N}\left(\begin{bmatrix} \nu_p \\ \nu_q \end{bmatrix}, \begin{bmatrix} \sigma_p^2 I & \\ & \sigma_q^2 I \end{bmatrix}\right) \tag{22}$$

$$H_1 : \begin{bmatrix} Y_p \\ Y_q \end{bmatrix} \sim \mathcal{N}\left(\begin{bmatrix} \delta \\ sR(\theta)\delta + Ft \end{bmatrix}, \begin{bmatrix} \sigma_p^2 I & \\ & \sigma_q^2 I \end{bmatrix}\right) \tag{23}$$

respectively. The appropriate matching hypothesis (i.e., H_0 or H_1) for the directed triangle data $Y_p = y_p$ and $Y_q = y_q$ is initially unknown. Given the realizations y_p and y_q from the stochastic maps of p and q, respectively, the matching hypothesis is determined using a *generalized likelihood ratio test (GLRT)* of the form

$$\Lambda(y_p, y_q) = \frac{\max_{\delta,s,t,\theta} L_1(\delta, s, t, \theta)}{\max_{\nu_p,\nu_q,s,t,\theta} L_0(\nu_p, \nu_q, s, t, \theta)} \underset{H_0}{\overset{H_1}{\gtrless}} \tau \tag{24}$$

where L_k are likelihood functions under H_k, with $k \in \{0, 1\}$, and the threshold τ is selected to control the level of false alarm. The likelihood statistic $\Lambda(y_p, y_q)$ is easily computed by applying Theorem 1. The performance of the approach in the presence of noise, as illustrated by the receiver operating characteristic (ROC) curves of Fig. 2, is of interest due to the uncertain nature of stochastic maps. As illustrated in the figure, the performance of the approach degrades gracefully with increasing levels of noise (the signal-to-noise ratio, or SNR, is discussed in Section V).

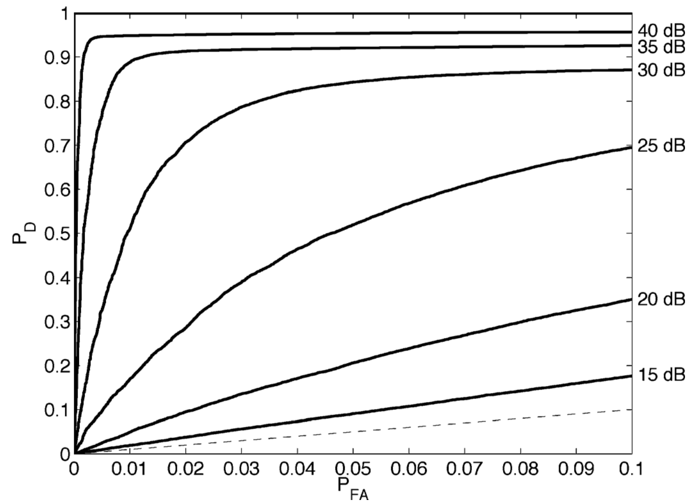

Fig. 2. Monte Carlo performance of GLRT likelihood statistic. Receiver operating characteristic (ROC) curves, illustrated above, show the performance of detecting triangle matches at various levels of SNR. Each of the curves are plots of the probability of detecting a match (P_{D}) versus the probability of a false alarm (P_{FA}). The dashed line in the lower region of the figure indicates the performance of a random guess.

Applying the GLRT enables the determination of common triangles from data, but not in a one-to-one fashion as required to produce a consistent combined map. If triangles $i \in \mathcal{P}$, with $\mathcal{P} = \{1, 2, \ldots, |\mathcal{E}_p|\}$, is denoted as y_p^i and $j \in \mathcal{Q}$, with $\mathcal{Q} = \{1, 2, \ldots, |\mathcal{E}_q|\}$, is denoted as y_q^j, then triangle matches are determined in a one-to-one fashion by formulating triangle matching as an *assignment problem* that seeks to

$$\text{maximize} \sum_{i=1}^{m}\sum_{j=1}^{m} f_{ij}(y_p^i, y_q^j) a_{ij} \tag{25}$$

$$\text{subject to} \sum_{i=1}^{m} a_{ij} = 1, \quad j = 1, \ldots, m \tag{26}$$

$$\sum_{j=1}^{m} a_{ij} = 1, \quad i = 1, \ldots, m \tag{27}$$

$$\text{and } a_{ij} \in \{0, 1\}, \tag{28}$$

where $m = \max(|\mathcal{E}_p|, |\mathcal{E}_q|)$ and the function f_{ij} used in the objective function is given by

$$f_{ij}(y_p^i, y_q^j) = \begin{cases} \Lambda(y_p^i, y_q^j), & i \in \mathcal{P} \text{ and } j \in \mathcal{Q} \\ 0, & \text{otherwise.} \end{cases} \tag{29}$$

An assignment obtained by the integer program is indicated by $a_{ij} = 1$, meaning triangle $i \in \mathcal{P}$ is assigned to triangle $j \in \mathcal{Q}$. A decision of $a_{ij} = 0$ indicates that triangles i and j are not assigned and the constraints (26) and (27) ensure that the assignments are one-to-one. The structure of the assignment problem

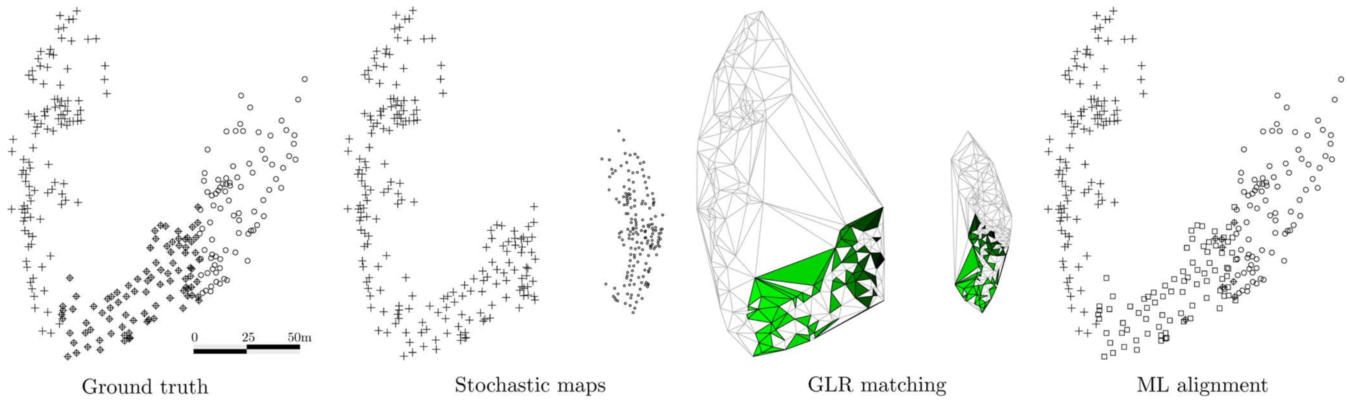

Fig. 3. Illustration of maximum likelihood fusion using the Victoria Park benchmark. (Ground truth) Ground truth locations observed by p and q are indicated by crosses (+) and circles (∘), respectively. (Stochastic maps) The maps of each agent are generated using a Gaussian noise model parameterized by s, t and θ. (GLR matching) Common landmarks are determined using a hypergraph representation to determine common directed triangles, illustrated above using color gradients to indicate matched triangles. (ML alignment) The combined map is computed using maximum likelihood estimation. Common landmarks found by the approach are indicated by squares (□).

allows for the use of standard linear programming routines by relaxing the integer constraints to $a_{ij} \in [0,1]$. The solution of the resulting linear program is indicated by the assignment set

$$\mathcal{A} \triangleq \left\{(i \in \mathcal{P}, j \in \mathcal{Q}) : a^*_{ij} = 1\right\} \tag{30}$$

which includes one-to-one assignments of directed triangles to be identified as belonging to H_0 or H_1. The GLRT (24) provides a statistical approach for determining a matching hypothesis, but (as indicated in Fig. 2) the performance of the test degrades with increasing noise. A robust detection scheme in the presence of uncertainty is to accept the assignments such that the MLEs $\{s^*_{ij}, t^*_{ij}, \theta^*_{ij}\}$ form a consensus [27]. Maximum likelihood estimation of the parameters $\{\mu, s, t, \theta\}$ is then accomplished by applying Theorem 1 to the landmarks of the triangle assignments under H_1 (i.e., common landmarks).

V. Numerical Examples and Simulations

The fusion approach of this paper constructs a global map of landmarks in two main steps referred to as GLR matching and ML alignment. This section provides an example of the ML fusion approach using the Victoria Park dataset and evaluates the performance of the matching and alignment steps in simulation.

Consider a Delaunay triangulation constructed from the ground truth landmarks observed by agent p and agent q with edge lengths $\{\ell_i : i = 1, 2, \ldots, m\}$. By modeling the variance of the stochastic maps as $\sigma_p^2 = \sigma_q^2 = \sigma^2$, the signal-to-noise ratio (SNR) in decibels follows as

$$\mathrm{SNR}_{\mathrm{dB}} = 10 \log \frac{\sigma_s^2}{\sigma_n^2} \tag{31}$$

with a signal variance of $\sigma_s^2 = \frac{1}{m}\sum_{i=1}^{m} \ell_i^2$ computed from the ground truth points and a noise variance of $\sigma_n^2 = 2\sigma^2$ since the zero-mean Gaussian noise is additive to the landmarks rather than to the edge lengths directly. The discussion of SNR in the remainder of the paper is in reference to (31).

A. Fusion Example (Victoria Park Benchmark)

An illustration of the proposed ML fusion approach is shown in Fig. 3. The ground truth vector u containing 285 landmark coordinates is obtained by applying the sparse local submap joining filter (SLSJF) proposed by Huang *et al.* [28] to the Victoria Park dataset. A simple overlap model is constructed by sorting the coordinates of the ground truth in order of increasing x-coordinates and partitioning u into two overlapping vectors u_p (containing 200 of the "top" entries of the sorted vector u) and u_q (containing 176 of the "bottom" entries of u). The resulting overlap of $n = 91$ landmarks serves as a model of the common landmarks, relative to ground truth, that are encountered along the paths independently explored by each agent within the Victoria Park environment. The stochastic maps of each agent are then generated using the additive noise model discussed in Section III at an SNR of 40 dB. Directed hypergraphs (Section IV-A) representing the maps of p and q contain $|\mathcal{E}_p| = 381$ and $|\mathcal{E}_q| = 337$ directed triangles, respectively, with an overlap of 144 triangles as a result of the ground truth construction and Gaussian noise model mentioned above. As illustrated in the figure, the spatial transformation of the map of agent q is parameterized by $\{s, t, \theta\}$ with a scale factor, rotation (in radians) and translation (in meters) given by $s = 0.5$, $\theta = 0.7854$ and $t = (150, 20)^T$, respectively. Common landmarks are determined using the GLR matching approach of Section IV, which are used to compute the closed form MLEs $s^* = 0.5000$, $\theta^* = 0.7851$ and $t^* = (150.0079, 19.9949)^T$ by applying Theorem 1. Once the maps of each agent are represented within a common frame of reference, missed detections in the matching are easily found using common data association techniques such as nearest neighbor and maximum likelihood. Estimation of the combined map $u^* = (\mu^{*T}, v_p^{*T}, v_q^{*T})^T$ immediately follows from the ML alignment approach described in Section III.

B. Monte Carlo Performance

The performance of GLR matching is considered, followed by the performance of ML alignment. The GLR matching ap-

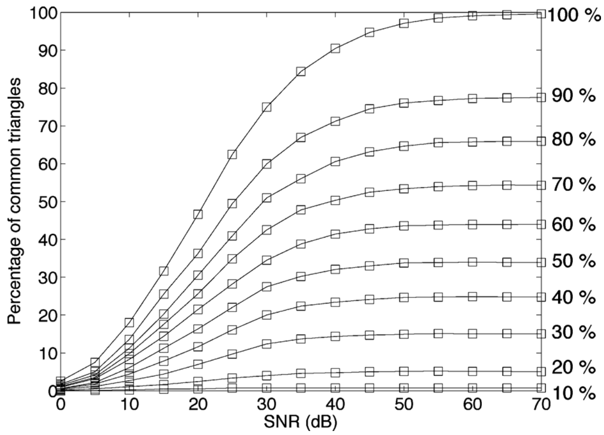

Fig. 4. Performance of hypergraph representation under various noise and landmark overlap. At each level of overlap (indicated by the percentages at the right), the percentage of common triangles found in the hypergraphs decreases with increasing noise and decreasing overlap.

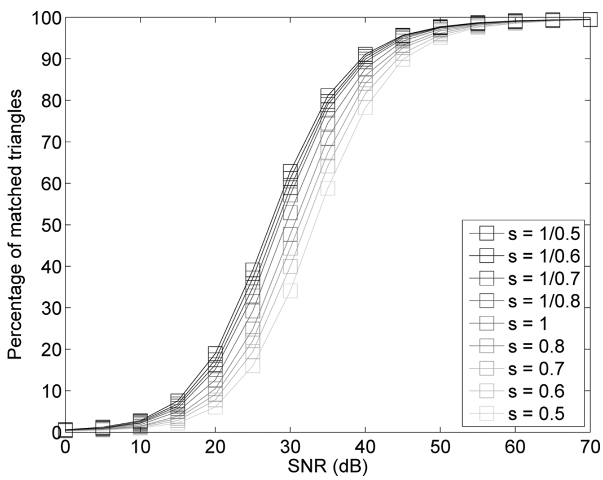

Fig. 5. Performance of GLR matching under various scale factors. In addition to the impact of noise is the consideration of the scale factor on the matching performance. As s tends to zero, landmarks become increasingly less distinguishable (as illustrated above by the decrease in performance for decreasing values of s).

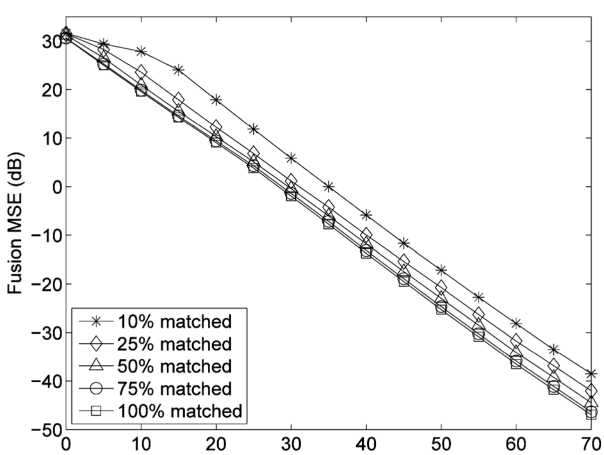

Fig. 6. Performance of ML alignment at various levels of triangle match detection. The figure above shows that the MSE performance exhibits a relatively low sensitivity to missed detections (e.g., the performance at 50% detection is comparable to the performance at 100%).

proach of Section IV represents the maps of agents p and q in the form of a directed hypergraph model. Fig. 2 shows the decline in performance of the GLR statistic used to determine common triangles between the hypergraph models for various SNR. In addition to the performance of the likelihood statistic is the noise and overlap performance of the hypergraph representation, illustrated by Fig. 4. As shown in the figure, the percentage of common triangles in the hypergraph representation for various levels of overlap decreases with decreasing SNR. In addition to noise is the change in performance due to landmark overlap. Consider the performance at the relatively high SNR of 70 dB. At 100% overlap, the hypergraph models of each agent contain all triangles in common (top right corner of Fig. 4). However, at 90%, 80% and 70% overlap, the figure shows that only 77.7188%, 66.1167% and 54.4552% of the hyperedges in each hypergraph represent common triangles, respectively, even at the relatively high SNR of 70 dB. The result thus shows that a gradual decline in the percentage of common triangles in the hypergraph representation can be expected to decrease with both decreasing SNR and decreasing overlap. The performance of the GLR matching applied to the directed hypergaph model under various scale factors and noise is shown in Fig. 5. As s tends to zero the landmarks become increasingly less distinguishable, resulting in a decrease in performance as shown in the figure for decreasing values of s.

The performance of ML alignment (Section IV) is shown in Fig. 6 for various levels of GLR matching performance. In addition to showing the trend in performance for various SNR, the figure also illustrates the robustness of the approach to missed detections. In particular, the figure shows that the MSE performance ranging from 50% detection to 100% is relatively the same for a wide range of noise. Fig. 6 thus indicates that relatively few matches are required to maintain the performance of the approach.

VI. Conclusion

This paper considered the problem of constructing a global map from the stochastic maps of collaborating mobile agents – fusion of stochastic maps. The problem can be formulated as a mixed integer and parameter estimation problem from which landmarks common to each agent are aligned under a global coordinate system. Under this framework, the optimal fusion of stochastic maps can be accomplished using the maximum likelihood principle. Unfortunately, however, the complexity of the true ML solution is prohibitive, which leads to a partitioning of the problem into two steps: (i) matching landmarks via bipartite hypergraph matching using a generalized likelihood ratio statistic as a matching metric and (ii) estimating the combined map within a common frame of reference using maximum likelihood estimation.

The main advantage of the proposed fusion approach, in spite of its suboptimality due to the separate treatment of the matching and alignment problems, is the comprehensive nature of the procedure: in spite of the individual stochastic maps being obtained in separate coordinate systems, a global map is found for various levels of noise and overlap with no prior knowledge of common landmarks. Monte Carlo simulations also show that the performance of the proposed approach degrades gracefully for various levels of match detection and scale. One way to further improve the approach would be to incorporate distortions due to changes in perspective (e.g., foreshortening), which is considered as future work.

APPENDIX

The constant 2×2 matrices I_c and I_s used in Theorem 1 are defined as

$$I_c \triangleq \begin{bmatrix} 1 & 0 \\ 0 & 1 \end{bmatrix} \text{ and } I_s \triangleq \begin{bmatrix} 0 & -1 \\ 1 & 0 \end{bmatrix}$$

respectively, with nonzero entries corresponding to the cosine and sine functions of the rotation matrix $r(\theta)$. In addition, the following lemma is used in the proof of Theorem 1.

Lemma 1: The matrix Q is an idempotent and symmetric matrix that commutes with a block diagonal matrix of the form $A = I_n \otimes B$, with $B \in \mathbb{R}^{2\times 2}$.

Proof of Lemma 1: The idempotence and symmetry properties are immediate from the structure of the matrix Q. The product of the matrix FF^T and the block diagonal matrix $A = I_n \otimes B$ is given by

$$\begin{aligned} FF^TA &= [(e_ne_n^T)I_n] \otimes [I_2B] \\ &= [I_n(e_ne_n^T)] \otimes [BI_2] \\ &= [I_n \otimes B][(e_ne_n^T) \otimes I_2] \\ &= AFF^T. \end{aligned}$$

Since $FF^TA = AFF^T$ and $F^TF = nI_2$, it follows that

$$\begin{aligned} QA &= [I_{2n} - F(F^TF)^{-1}F^T]A \\ &= A - \frac{1}{n}FF^TA \\ &= A - \frac{1}{n}AFF^T \\ &= A[I_{2n} - F(F^TF)^{-1}F^T] \\ &= AQ \end{aligned}$$

which proves that Q and A commute. ■

Proof of Theorem 1: Minimizing (8) with respect to (w.r.t.) μ leads to

$$\bar{\mu}(s,t,\theta) = \frac{\sigma_p^2\sigma_q^2}{\sigma_p^2+s^2\sigma_q^2}\left[\frac{1}{\sigma_p^2}x_p^1+\frac{1}{\sigma_q^2}sR_1^T(\theta)(x_q^1 - F_1t)\right] \quad (32)$$

and minimizing (8) w.r.t. t leads to

$$\bar{t}(s,\mu,\theta) = (F_1^TF_1)^{-1}F_1^T\left(x_q^1 - sR_1(\theta)\mu\right). \quad (33)$$

Using the evaluation $\mu = \bar{\mu}(\theta, t)$ in (33) results in the MLE of t as a function of θ given by

$$t^*(s,\theta) = (F_1^TF_1)^{-1}F_1^T\left(x_q^1 - sR_1(\theta)x_p^1\right). \quad (34)$$

Applying the evaluation $t = t^*(s,\theta)$ in (32) leads to the MLE of μ as a function of s and θ given by

$$\mu^*(s,\theta) = \phi_p(s,\theta)x_p^1 + \phi_q(s,\theta)x_q^1. \quad (35)$$

Using the symmetry and idempotence properties of the matrix Q (Lemma 1), it follows from the expression (35) that

$$\begin{aligned} \|x_p^1 - \mu^*(s,\theta)\|^2 &= \kappa_p\|x_q^1 - sR_1(\theta)x_p^1\|_{Q^TQ}^2 \\ &= \kappa_p\|x_q^1 - sR_1(\theta)x_p^1\|_Q^2 \end{aligned}$$

where $\kappa_p = \left(\frac{s\sigma_p^2}{\sigma_p^2+s^2\sigma_q^2}\right)^2$. Similarly, it follows from (34) and (35) that

$$\begin{aligned} \|x_q^1 - sR_1(\theta)\mu^*(\theta) - F_1t^*(\theta)\|^2 &= \kappa_q\|x_q^1 - sR_1(\theta)x_p^1\|_{Q^TQ}^2 \\ &= \kappa_q\|x_q^1 - sR_1(\theta)x_p^1\|_Q^2 \end{aligned}$$

where $\kappa_q = \left(1 - \frac{s^2\sigma_p^2}{\sigma_p^2+s^2\sigma_q^2}\right)^2$. Using these simplifications to define

$$\begin{aligned} J_1^*(s,\theta) &\triangleq \frac{1}{2\kappa}J_1(\mu^*, s, t^*, \theta) \\ &= \frac{1}{2}\|x_q^1 - sR_1(\theta)x_p^1\|_Q^2 \end{aligned} \quad (36)$$

where $\kappa = \frac{1}{\sigma_p^2}\kappa_p + \frac{1}{\sigma_q^2}\kappa_q$ and expanding the norm in the right hand side (RHS) of (36) as

$$\begin{aligned} \|x_q^1 - sR_1(\theta)x_p^1\|_Q^2 = s^2x_p^{1T}R_1^T(\theta)QR_1(\theta)x_p^1 + x_q^{1T}Qx_q^1 \\ -2sx_q^{1T}QR_1(\theta)x_p^1, \end{aligned} \quad (37)$$

it follows from Lemma 1 that the first term on the RHS of (37) reduces to

$$\begin{aligned} s^2x_p^{1T}R_1^T(\theta)QR_1(\theta)x_p^1 &= s^2x_p^{1T}R_1^T(\theta)R_1(\theta)Qx_p^1 \\ &= s^2x_p^{1T}Qx_p^1 \end{aligned} \quad (38)$$

so that from (36), (37) and (38), $J_1^*(\theta)$ reduces to

$$J_1^*(s,\theta) = \frac{1}{2}(s^2x_p^{1T}Qx_p^1+x_q^{1T}Qx_q^1-2sx_q^{1T}R_1(\theta)Qx_p^1) \quad (39)$$

which leads to

$$s^*(\theta) = \frac{x_p^{1T}QR_1^T(\theta)x_q^1}{x_p^{1T}Qx_p^1} \quad (40)$$

when minimized w.r.t. s. Notice in the last term on the RHS of (39) that

$$\begin{aligned} x_q^{1T}R_1(\theta)Qx_p^1 &= x_q^{1T}\left[R_c(\theta) + R_s(\theta)\right]Qx_p^1 \\ &= x_q^{1T}R_c(\theta)Qx_p^1 + x_q^{1T}R_s(\theta)Qx_p^1 \end{aligned} \quad (41)$$

where $R_c(\theta) = (I_n \otimes I_c)\cos(\theta)$ and $R_s(\theta) = (I_n \otimes I_s)\sin(\theta)$, meaning that

$$-2x_q^{1T}R_1(\theta)Qx_p^1 = 2\alpha\cos(\theta) + 2\beta\sin(\theta)$$

where $\alpha = -x_q^{1T}(I_n \otimes I_c)Qx_p^1$ and $\beta = -x_q^{1T}(I_n \otimes I_s)Qx_p^1$, from which it immediately follows that

$$J_1^*(s,\theta) = s\alpha\cos(\theta) + s\beta\sin(\theta) + \gamma(s) \quad (42)$$

where $\gamma(s) = \frac{1}{2}\left(s^2 x_p^{1T} Q x_p^1 + x_q^{1T} Q x_q^1\right)$. By virtue of the sinusoidal form (42), it follows that $J_1^*(s,\theta)$ has a unique minimum for $\theta \in [-\pi, \pi]$ given by

$$\begin{aligned}\theta^* &= \operatorname{sgn}(s\beta)\left[\cos^{-1}\left(\frac{s\alpha}{\sqrt{(s\alpha)^2+(s\beta)^2}}\right)-\pi\right]\\ &= \operatorname{sgn}(\beta)\left[\cos^{-1}\left(\frac{\alpha}{\sqrt{\alpha^2+\beta^2}}\right)-\pi\right]\end{aligned}$$

for $s > 0$, which leads to the MLEs of s, t and μ given by (16), (17) and (18), respectively. ■

ACKNOWLEDGMENT

The authors would like to thank J. Guivant, J. Nieto and E. Nebot for providing the Victoria Park dataset as a benchmark. The satellite image of Fig. 1 was captured using Google Maps with a latitude-longitude coordinate of $(-33.886577, 151.192061)$. Many thanks are also extended to the anonymous reviewers for the various comments relating to issues of scale, landmark overlap and experimental setup.

REFERENCES

[1] R. Smith, M. Self, and P. Cheeseman, "A stochastic map for uncertain spatial relationships," in *Proc. 4th Int. Symp. on Robot. Res.*, Cambridge, MA, USA, 1988, pp. 467–474.

[2] R. Smith, M. Self, and P. Cheeseman, "Estimating uncertain spatial relationships in robotics," in *Autonomous Robot Vehicles*. New York, NY, USA: Springer-Verlag, 1990, vol. 8, pp. 167–193.

[3] Y. Zhou, H. Leung, and P. Yip, "An exact maximum likelihood registration algorithm for data fusion," *IEEE Trans. Signal Process.*, vol. 45, no. 6, pp. 1560–1573, 1997.

[4] R. Viswanathan and P. Varshney, "Distributed detection with multiple sensors I: Fundamentals," *Proc. IEEE*, vol. 85, no. 1, pp. 54–63, 1997.

[5] R. Blum, S. Kassam, and H. Poor, "Distributed detection with multiple sensors II: Advanced topics," *Proc. IEEE*, vol. 85, no. 1, pp. 64–79, 1997.

[6] J. Chamberland and V. Veeravalli, "Decentralized detection in sensor networks," *IEEE Trans. Signal Process.*, vol. 51, no. 2, pp. 407–416, 2003.

[7] R. Niu, B. Chen, and P. Varshney, "Fusion of decisions transmitted over Rayleigh fading channels in wireless sensor networks," *IEEE Trans. Signal Process.*, vol. 54, no. 3, pp. 1018–1027, 2006.

[8] M. Dissanayake, P. Newman, S. Clark, H. Durrant-Whyte, and M. Csorba, "A solution to the simultaneous localization and map building (SLAM) problem," *IEEE Trans. Robot. Autom.*, vol. 17, pp. 229–241, Jun. 2001.

[9] H. Durrant-Whyte and T. Bailey, "Simultaneous localization and mapping (SLAM): Part I," *IEEE Robot. Autom. Mag.*, vol. 13, no. 2, pp. 99–110, 2006.

[10] H. Durrant-Whyte and T. Bailey, "Simultaneous localization and mapping (SLAM): Part II," *IEEE Robot. Autom. Mag.*, vol. 13, no. 3, pp. 108–117, 2006.

[11] S. Thrun and Y. Liu, "Multi-robot SLAM with sparse extended information filers," presented at the 11th Int. Symp. Robot. Res., Sienna, Italy, 2003.

[12] D. Hall and J. Llinas, "An introduction to multisensor data fusion," *Proc. IEEE*, vol. 85, no. 1, pp. 6–23, 1997.

[13] N. Katayama and S. Satoh, "The SR-tree: An index structure for high-dimensional nearest neighbor queries," in *Proc. ACM SIGMOD Int. Conf. Manage. Data (SIGMOD'97)*, New York, NY, USA, 1997, pp. 369–380.

[14] S. Grime and H. Durrant-Whyte, "Data fusion in decentralized sensor networks," *Contr. Eng. Practice*, vol. 2, no. 5, pp. 849–863, 1994.

[15] S. Sukkarieh, E. Nettleton, J.-H. Kim, M. Ridley, A. Goktogan, and H. Durrant-Whyte, "The ANSER project: Data fusion across multiple uninhabited air vehicles," *Int. J. Robot. Res.*, vol. 22, pp. 505–539, Jul. 2003.

[16] S. Julier and J. Uhlmann, "A non-divergent estimation algorithm in the presence of unknown correlations," in *Proc. Amer. Control Conf.*, 1997, vol. 4, pp. 2369–2373.

[17] J. Tardós, J. Neira, P. Newman, and J. Leonard, "Robust mapping and localization in indoor environments using sonar data," *Int. J. Robot. Res.*, vol. 21, no. 4, pp. 311–330, 2002.

[18] J. Castellanos, R. Martinez-Cantin, J. Tardós, and J. Neira, "Robocentric map joining: Improving the consistency of EKF-SLAM," *Robot. Autonom. Syst.*, vol. 55, no. 1, pp. 21–29, 2007.

[19] S. B. Williams, G. Dissanayake, and H. F. Durrant-Whyte, "Towards multi-vehicle simultaneous localisation and mapping," in *Proc. IEEE Int. Conf. Robot. Autom. (ICRA)*, 2002, pp. 2743–2748.

[20] X. Zhou and S. Roumeliotis, "Multi-robot SLAM with unknown initial correspondence: The robot rendezvous case," in *Proc. Int. Conf. Intell. Robots Syst.*, Oct. 2006, pp. 1785–1792.

[21] L. Andersson and J. Nygards, "C-SAM: Multi-robot SLAM using square root information smoothing," in *Proc. Int. Conf. Robot. Autom.*, 2008, pp. 2798–2805.

[22] D. Benedettelli, A. Garulli, and A. Giannitrapani, "Multi-robot SLAM using M-Space feature representation," in *Proc. IEEE Conf. Decision Control*, 2010, pp. 3826–3831.

[23] R. Aragues, J. Cortes, and C. Sagues, "Distributed consensus algorithms for merging feature-based maps with limited communication," *Robot. Autonom. Syst.*, vol. 59, no. 3-4, pp. 163–180, 2011.

[24] E. Groth, "A pattern-matching algorithm for two-dimensional coordinate lists," *Astronom. J.*, vol. 91, no. 5, pp. 1244–1248, 1986.

[25] H. Ogawa, "Labeled point pattern matching by Delaunay triangulation and maximal cliques," *Pattern Recogn.*, vol. 19, no. 1, pp. 35–40, 1986.

[26] B. Delaunay, "Sur la sphère vide," *Izv. Akad. Nauk SSSR, Otdelenie Matematicheskikh i Estestvennykh Nauk*, vol. 7, pp. 793–800, 1934.

[27] M. A. Fischler and R. C. Bolles, "Random sample consensus: A paradigm for model fitting with applications to image analysis and automated cartography," *Commun. ACM*, vol. 24, no. 6, pp. 381–395, 1981.

[28] S. Huang, Z. Wang, and G. Dissanayake, "Sparse local submap joining filter for building large-scale maps," *IEEE Trans. Robot.*, vol. 24, no. 5, pp. 1121–1130, 2008.

Brandon M. Jones (M'13) received the B.S. degree in electrical engineering from Southern University and A&M College, Baton Rouge, LA, in 2005 and the M.S. and Ph.D. degrees in electrical and computer engineering from Cornell University, Ithaca, NY, in 2013.

His current research interests are in the areas of optimization, statistical inference, and autonomous systems.

Dr. Jones is a recipient of the 2007 NASA Harriett G. Jenkins Fellowship and the Cornell University Sloan Fellowship.

Mark Campbell (M'00) received the B.S. degree in mechanical engineering from Carnegie Mellon University, and the M.S. and Ph.D. degrees in control and estimation from the Massachusetts Institute of Technology in 1993 and 1996, respectively.

He is currently a Professor and the S. C. Thomas Sze Director of the Sibley School of Mechanical and Aerospace Engineering at Cornell University. During 2005–2006, he was Visiting Scientist at the Insitu group, and an ARC International Fellow at the Australian Centre of Field Robotics. His research interests are in the areas of autonomous systems, with a focus on probabilistic models of human decision making, nonlinear estimation theory, cooperative vehicle control and estimation, and sensor fusion.

Dr. Campbell received Best Paper awards from AIAA Propulsion and GNC conferences, and Frontiers in Education conference, and teaching awards from

Cornell University, University of Washington, and ASEE. He has been recognized from NASA for his modeling and control work on the MACE Experiment, flown on STS-67. He has been an invited speaker for the NAE Frontiers in Engineering Symposium and NAS Kavli Frontiers of Science Symposium. He has served as Associate Editor for the IEEE TRANSACTIONS ON AEROSPACE AND ELECTRONIC SYSTEMS and the AIAA *Journal of Guidance, Control and Dynamics*.

Lang Tong (S'87–M'91–SM'01–F'05) received the B.E. degree in automation from Tsinghua University, Beijing, China, in 1985, and the M.S. and Ph.D. degrees in electrical engineering in 1987 and 1991, respectively, from the University of Notre Dame, Notre Dame, IN.

He is the Irwin and Joan Jacobs Professor in Engineering at Cornell University, Ithaca, NY. He is also the Cornell site direction of the Power System Engineering Research Center (PSERC). His current research focuses on inference, optimization, and economic problems in energy and power systems. He was a Postdoctoral Research Affiliate at the Information Systems Laboratory, Stanford University, in 1991. He was the 2001 Cor Wit Visiting Professor at the Delft University of Technology and had held visiting positions at Stanford University and the University of California at Berkeley.

Dr. Tong received the 1993 Outstanding Young Author Award from the IEEE Circuits and Systems Society, the 2004 Best Paper award from IEEE Signal Processing Society, and the 2004 Leonard G. Abraham Prize Paper Award from the IEEE Communications Society. He is also a coauthor of seven student paper awards. He received Young Investigator Award from the Office of Naval Research. He was a Distinguished Lecturer of the IEEE Signal Processing Society.