

Distributed Navigation with Unknown Initial Poses and Data Association via Expectation Maximization

Vadim Indelman*, Nathan Michael†, and Frank Dellaert‡

We present a novel approach for multi-robot distributed and incremental inference over variables of interest, such as robot trajectories, considering the initial relative poses between the robots and multi-robot data association are both unknown. Assuming robots share with each other informative observations, this inference problem is formulated within an Expectation-Maximization (EM) optimization, performed by each robot separately, alternating between inference over variables of interest and multi-robot data association. To facilitate this process, a common reference frame between the robots should first be established. We show the latter is coupled with determining multi-robot data association, and therefore concurrently infer both using a separate EM optimization. This optimization is performed by each robot starting from several promising initial solutions, converging to locally-optimal hypotheses regarding data association and reference frame transformation. Choosing the best hypothesis in an *incremental* problem setting is in particular challenging due to high sensitivity to *measurement aliasing* and possibly insufficient amount of data. Selecting an incorrect hypothesis introduces outliers and can lead to catastrophic results. To address these challenges we develop a model-selection based approach to choose the most probable hypothesis, while resorting to Chinese Restaurant Process to represent statistical knowledge regarding hypothesis prior probabilities. We evaluate our approach in real-data experiments.

I. Introduction

Distributed inference is a key capability in multi-robot autonomous systems that is of interest in a variety of problem domains, including surveillance, tracking, localization and mapping. Cooperatively inferring variables of interest, such as robot trajectories, observed objects and tracked targets, results in higher levels of performance, flexibility and robustness to failure. The research community has been addressing different aspects of this problem, considering both centralized and decentralized frameworks.

To facilitate cooperative inference it is essential to establish a common reference frame and world model between the robots, so that these can communicate with each other relevant information and correctly interpret it for their needs. While each of these problems has been previously addressed assuming the other problem is solved, only few attempts have been made to solve the two problems *simultaneously*: determining a common reference frame between the robots, and resolving data association between measurements (e.g. images or laser scans) acquired by different robots.

* Faculty of Aerospace Engineering, Technion - Israel Institute of Technology, Haifa 32000, Israel. Part of this research was performed while the author was with the Georgia Institute of Technology.

† Robotics Institute, Carnegie Mellon, Pittsburgh, PA 15213, USA.

‡ Institute for Robotics and Intelligent Machines (IRIM), Georgia Institute of Technology, Atlanta, GA 30332, USA.

Solving these two coupled problems is important as it will enable the robots, scattered in a complex, previously unknown environment, to establish collaboration without requiring any prior knowledge or infrastructure. For example, starting from some arbitrary guess as to where each robot is and by sharing measurements of onboard sensors, each robot will be capable of inferring the trajectories of other robots in the group, as shown in Figures 1 and 2.




Figure 1: Multi-robot candidate correspondences \mathcal{F}^r between (a) red and blue, and (b) green and blue robots in the real-world dataset **D1**. As a common reference frame is not yet established, robot initial poses are set to arbitrary values.

Multi-robot data association is a key challenge that shares some similarities with loop closure detection in the single-robot case. Incorrect data association can lead to catastrophic deterioration in performance and should be avoided at all costs; the robotics community has been indeed very active in the last two years in developing robust graph optimization techniques [22, 23, 31, 32] to address this crucial aspect.

Multi-robot data association has recently become an active research area as well [9, 19, 21, 26], with the same sensitivity to incorrect correspondences as in the single-robot case. This problem, however, becomes more complicated when the initial relative poses between the robots are unknown. Without a common reference frame, how can the robot decide what information to share with each other? Given the calculated multi-robot constraints based on this shared information, how to determine the inlier correspondences? Addressing this problem requires reasoning about multi-robot data association and initial relative poses concurrently.

In this paper we develop a distributed and incremental multi-robot approach that allows a group of robots to simultaneously establish a common reference frame and resolve multi-robot data association on-the-fly. We show that establishing the initial relative pose transformations between different robots is essential to facilitate collaboration between robots: attempting to solve the multi-robot PoseSLAM problem with unknown multi-robot data association *and* unknown initial relative pose transformations leads in all multi-robot correspondences being considered as outliers and therefore rejected. We therefore first infer these transformations and only then address the full multi-robot PoseSLAM problem.

Our approach is based on the observation that by analyzing the distribution of multi-

robot relative pose constraints it is possible to estimate the transformation between the robot reference frames and identify the inliers in these constraints: we show that this distribution is clustered for inlier and scattered for outlier correspondences, as long as there is no perceptual aliasing. Based on this insight we develop an expectation-maximization (EM) [27] approach to efficiently perform this inference by each of the robots independently, starting from several initial guesses and resulting in different locally-optimal solutions.

Choosing the correct solution is a key challenge, as a wrong decision will effectively introduce outliers to the graph optimization. This is particularly true when information is received *incrementally*, as required for online operation in real autonomous robotic systems: in this case, one needs also to decide whether *sufficient* amount of information has been received to perform this decision reliably.

Furthermore, perceptual aliasing presents additional challenges, as matching observations from different, but similar in appearance environments (such as two corridors) often result in reliable and consistent transformations that erroneously indicate the two environments are the same. These outliers form *clusters* of their own that compete with the inliers cluster, leading to the problem of choosing the right cluster among several candidates.

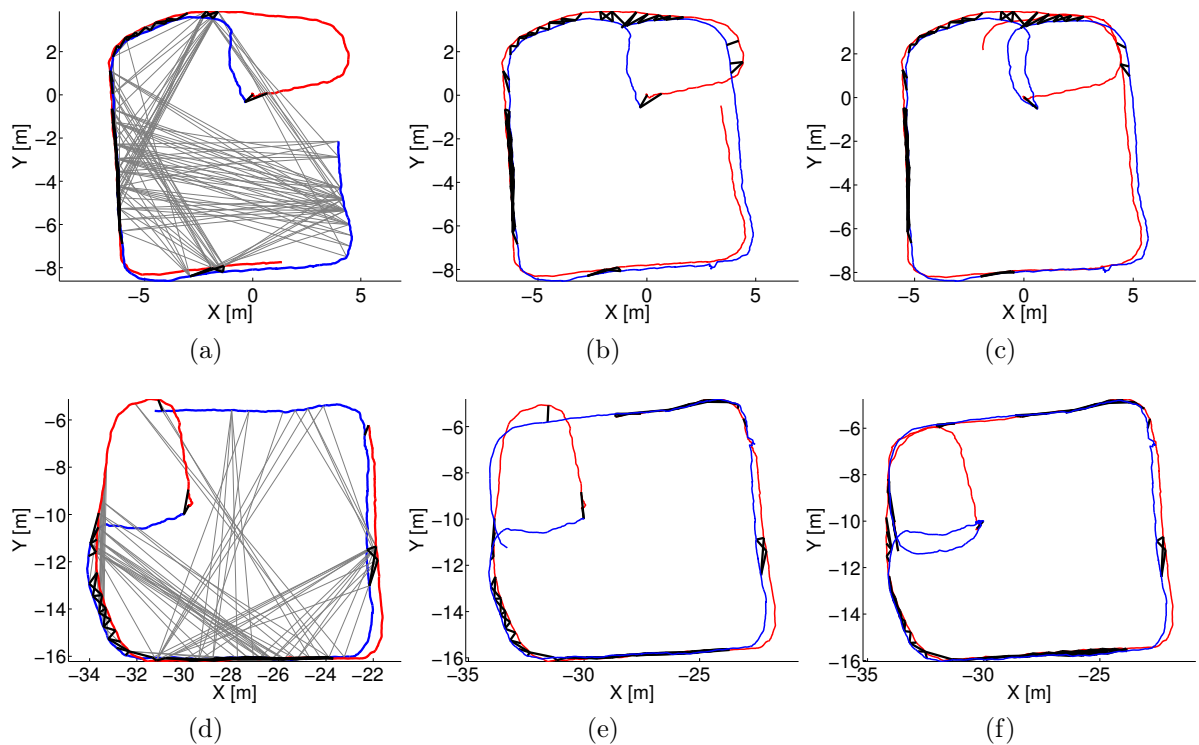


Figure 2: Robots' estimates of each other trajectories expressed in local frame of each robot over time. (a)-(c) Red robot; (d)-(f) Blue robot. Figures (a) and (d) show the identified inlier (black) and outlier (gray) correspondences after a common reference frame has been established.

We consider this challenging problem and frame it within a *model selection* framework [5], developing a probabilistically-sound approach for selecting the most probable cluster. Moreover, we address the question whether there *is* a correct cluster given the information available to each robot thus far, as the robots might have not observed the

same environment yet. We approach this problem by modeling the prior probability for each cluster using the Chinese restaurant process (e.g., [4,33]) that allows to disambiguate this decision-making as more information is accumulated.

Consequently, this paper makes the following contributions: (i) development of a new approach for distributed incremental multi-robot inference with unknown multi-robot data association and initial relative poses; (ii) Model selection for identifying the correct solution among several candidates; (iii) Determining if sufficient information has been obtained to disambiguate potential perceptual aliasing; (iv) Performance evaluation in realistic simulation and in real-world indoor experiments.

II. Related Work

Distributed cooperative localization and multi-robot simultaneous localization and mapping has been extensively studied over the last decade. Many of these research efforts, that can be classified into Full-SLAM (e.g. [11, 15, 16, 30]) and Pose-SLAM (e.g. [3, 17, 18, 20, 30, 34]) approaches, assume the initial relative poses between the robots are known and multi-robot data association has been externally established.

Several approaches have been developed to operate also when the initial relative poses between the robots are unknown, still assuming perfect multi-robot data association is given. These approaches, including [1, 6, 35], often use direct relative-pose observations between the robots and assume the identity of the observed robots is known. Relaxing this assumption has been recently investigated in [13]. Multi-robot Pose-SLAM using also *indirect* relative pose constraints, corresponding to different robots observing the same environment, has been mainly researched under the assumption of perfect data association and an established common reference frame between the robots (e.g. [17, 18, 20, 25]).

Determining data association is often treated as a pre-processing separate step from inference over the robot states. Cunningham et al. [9] developed a RANSAC [12] approach for distributed data association, and Montijano et al. [26] compute global correspondences by identifying inconsistent data association in a distributed framework. While these approaches do not require a common reference frame between the robots, they operate within the Full-SLAM paradigm, explicitly performing inference also over the observed structure (e.g. 3D points, objects) in addition to robot states.

In contrast, our approach is formulated in Pose-SLAM framework, which has computational advantages over Full-SLAM. Consistency-based outlier rejection has been recently also developed for multi-robot Pose-SLAM [21] assuming direct relative-pose observations, still assuming a common reference frame between different robots is known.

To be resilient to outliers overlooked by data association approaches, the robotics community has been recently focusing on robust graph optimization techniques [22, 23, 31, 32]. These new approaches in particular aim to be robust to loop-closures outliers, as current state of the art methods for generating loop closure constraints (e.g. FAB-MAP place recognition [7]) are not error-free. However, the majority of robust graph optimization approaches are developed for the single robot case, typically not considering the multi-robot case. In a previous recent work [19], we investigated the multi-robot case, considering a centralized and batch setting and without considering perceptual aliasing aspects. An incremental and distributed multi-robot configuration, that is the focus of this paper, has not been investigated thus far to the best of our knowledge.

III. Problem Formulation and Approach

We consider a group of R robots deployed to collaboratively operate in an unknown environment, *initially unaware* of each other. Our objective is for each robot r to estimate its own trajectory X^r (current and past poses) and additional variables of interest, such as the trajectories of other robots, in a distributed incremental framework. Such a capability is important for multi-robot cooperation in numerous applications; additionally, it allows the robots to extend their sensing horizon and establish a common world model, observed so far by the *entire* group. Although the latter is not explicitly inferred, it can be always recovered from the estimated poses and sensor observations (see, e.g. [14]).

In a centralized problem setting, this problem corresponds to calculating the maximum a posteriori (MAP) estimate

$$\hat{\mathcal{X}}^{\mathcal{R}} = \arg \max_{\mathcal{X}^{\mathcal{R}}} p(\mathcal{X}^{\mathcal{R}} | \mathcal{Z}^{\mathcal{R}}), \quad (1)$$

where $\mathcal{R} \doteq \{1, \dots, R\}$, and $\mathcal{X}^{\mathcal{R}}$ and $\mathcal{Z}^{\mathcal{R}}$ represent the trajectories and the observations of all the robots in the group.

We assume the robots share with each other observations of different parts of the environment, which, *given* data association, can be used to formulate constraints between appropriate poses of different robots. These constraints become part of the underlying graph structure of $p(\mathcal{X}^{\mathcal{R}} | \mathcal{Z}^{\mathcal{R}})$ and are essential for facilitating collaboration between the robots.

Assuming a common reference frame is known, the robots can identify and share with each other only observations of areas that are likely to be observed by other robots. However, in lack of a common reference frame, it is not obvious what information should the robots communicate with each other, as each robot r represents its trajectory X^r in its own local frame. Consequently, data association becomes much more challenging because of the high number of outlier correspondences, and as shown in the sequel, is coupled with establishing a common reference frame between the robots.

A distributed setting further complicated this problem, as each robot r has only access to $\mathcal{Z}^r \in \mathcal{Z}^{\mathcal{R}}$, comprising its own observations and observations shared by other robots. Letting $\mathcal{X}^r \subseteq \mathcal{X}^{\mathcal{R}}$ represent the trajectory of robot r and additional variables of interest, in our case, the appropriate trajectories of other robots (as defined in the sequel), the inference solved by each robot r is

$$\hat{\mathcal{X}}^r = \arg \max_{\mathcal{X}^r} p(\mathcal{X}^r | \mathcal{Z}^r). \quad (2)$$

Below we develop an EM-based approach for solving this problem in an incremental and distributed setting with unknown multi-robot data association and initial relative poses. Since this approach requires the initial relative poses to be first roughly determined, we develop in Section B a method that allows each robot to infer the latter and multi-robot data association simultaneously. As will be seen, the information available to each robot in an incremental problem setting at a given time often supports *several* possible solutions. The problem then turns into choosing the most probable solution and identifying whether sufficient amount of information has been accumulated to perform this reliably. We show this is particularly important in the incremental setting in the presence of *measurement aliasing*. In Section V we address this crucial step, formulating the problem within a model-based selection paradigm.

IV. Incremental Distributed Framework

In this section we develop a Bayesian formulation for the inference problem (2) to be solved by each robot r . First, however, we discuss our approach for information sharing between different robots, when a common reference frame is unknown. Throughout the paper we use the notation $x_i^r \in X^r$ to represent the pose of robot r at time t_i , expressed in its local reference frame.

We assume each robot r shares at each time t_k its current measurement z_k^r , if it is informative, and also tracks all these informative measurements $\{z_i^r\}$ over time. Additionally, it shares its current pose estimate \hat{x}_k^r and local measurements from its onboard sensors. We assume these measurements and information are shared at some basic frequency, or upon moving sufficient distance (for example, 0.5 meters in our implementation). Sharing only local measurements guarantees *consistent* multi-robot estimates [8] (additional approaches include [2, 18]). While this may seem as a lot of information to share, in practice one can represent all these measurements by a single relative pose constraint with an appropriate covariance (see e.g. [8]).

Any robot r that receives a measurement $z_k^{r'}$ from some robot r' , generates candidate correspondences by matching $z_k^{r'}$ with its own informative measurements $\{z_i^r\}$. Each such successful match $u_{k,l}^{r',r}$ between $z_k^{r'}$ and $z_l^r \in \{z_i^r\}$ represents a (relative-pose) constraint involving the poses $x_k^{r'}$ and x_l^r , with $l \leq k$. We denote by \mathcal{F}^r the set of multi-robot data association, that is available to robot r , where each individual data association $(r', k, l) \in \mathcal{F}^r$ represents the constraint $u_{k,l}^{r',r}$.

An examples of the multi-robot candidate correspondences set \mathcal{F}^r is shown in Figure 1. The figure illustrates the candidate correspondences in \mathcal{F}^r between the blue robot and other robots (green and red). Since the initial relative poses between the robots are unknown, these transformations were set to *arbitrary* values, i.e. the initial pose of each robot was chosen arbitrarily.

Observe that many of these correspondences in \mathcal{F}^r are outliers. One may argue that outliers can be directly identified and rejected by matching algorithms, such as RANSAC-based fundamental matrix estimation in the case of image observations or ICP matching in the case of laser measurements. However, this argument is only partially true: while these algorithms are capable of accurate relative pose estimation given observations of a *common* scene, identifying the fact that two given observations were acquired from *different* parts of the environment is a much more challenging task. In the single robot case, a related aspect is establishing loop closures; also here current techniques cannot guarantee outlier-free data associations and approaches for robust graph optimization are actively investigated in recent years. Similarly, in the multi-robot framework that we consider herein, multi-robot data association cannot be assumed outlier-free.

Moreover, perceptual aliasing will often result in *numerous* false data associations. In particular, when fed with two laser scans from different yet *similar* parts of the environment (e.g. corridors, hallways), ICP will typically produce some relative pose estimate, with a reasonable uncertainty covariance and number of matched points between the two scans. Not only this estimate is completely wrong, but also *similar* erroneous estimates will be obtained for all such outliers, making it difficult to identify these are all outliers. We address this crucial aspect in detail in Section V.

We can now formulate the joint probability distribution (2) given all the measurements

available to robot r :

$$p(\mathcal{X}^r | \mathcal{Z}^r) \doteq p\left(X^r, \mathcal{X}^{\mathcal{R} \setminus \{r\}} | Z^r, \mathcal{Z}^{\mathcal{R} \setminus \{r\}}, \mathcal{Z}_{local}^{\mathcal{R} \setminus \{r\}}\right), \quad (3)$$

where the set $\mathcal{X}^{\mathcal{R} \setminus \{r\}}$ represents all the poses $x_i^{r'}$ of robots $r' \in \mathcal{R}$ that contributed at least one correspondence to \mathcal{F}^r . Accordingly, $\mathcal{Z}^{\mathcal{R} \setminus \{r\}}$ represents the actual measurements that participate in those constraints, and $\mathcal{Z}_{local}^{\mathcal{R} \setminus \{r\}}$ the local observations of these robots, both shared by these robots. For brevity, we collect all these measurements into the set \mathcal{Z}^r and also let \mathcal{X}^r represent all the random variables in Eq. (3):

$$\mathcal{Z}^r \doteq Z^r \cup \mathcal{Z}^{\mathcal{R} \setminus \{r\}} \cup \mathcal{Z}_{local}^{\mathcal{R} \setminus \{r\}}, \quad \mathcal{X}^r \doteq X^r \cup \mathcal{X}^{\mathcal{R} \setminus \{r\}}. \quad (4)$$

Given the multi-robot data association \mathcal{F}^r , and the appropriate constraints $u_{k,l}^{r',r}$, the joint pdf (3) can be expressed as

$$p(\mathcal{X}^r | \mathcal{Z}^r) \propto p(X^r | Z^r) p\left(\mathcal{X}^{\mathcal{R} \setminus \{r\}} | \mathcal{Z}_{local}^{\mathcal{R} \setminus \{r\}}\right) \prod_{(r',k,l) \in \mathcal{F}^r} p\left(u_{k,l}^{r',r} | x_k^{r'}, x_l^r\right) \quad (5)$$

where $p(X^r | Z^r)$ is only a function of local observations Z^r . Since we assume no a priori knowledge about the environment and the initial pose of the robots, the reference frame of each robot is set *arbitrarily* an value.

As the robots express their local trajectories with respect to *different* reference systems, the measurement likelihood term in Eq. (5) is

$$p\left(u_{k,l}^{r',r} | x_k^{r'}, x_l^r\right) \propto \exp\left(-\frac{1}{2} \left\| \text{err}\left(u_{k,l}^{r',r}, x_k^{r'}, x_l^r\right) \right\|_{\Sigma}^2\right), \quad (6)$$

with

$$\text{err}\left(u_{k,l}^{r',r}, x_k^{r'}, x_l^r\right) \doteq u_{k,l}^{r',r} \ominus h\left(x_k^{r'}, x_l^r\right), \quad (7)$$

and

$$h\left(x_k^{r'}, x_l^r\right) \doteq x_k^{r'} \ominus \left(T_r^{r'} \oplus x_l^r\right), \quad (8)$$

where we use the notation \ominus in $a \ominus b$ to express b locally in the frame of a , and \oplus to represent the compose operator [24].

In Eq. (8), $T_r^{r'}$ represents a transformation between the reference frames of robots r and r' . Observe that if this transformation was known, performing inference over Eq. (5) while identifying and rejecting outliers could be addressed by robust graph optimization techniques [22,23,31,32]. In contrast, we consider multi-robot data association and initial relative poses are *both* unknown.

A. Multi-Robot Data Association and Localization via EM

Instead of assuming multi-robot data association is given, we introduce a latent binary variable $j_{k,l}^{r',r}$ for each correspondence $(r',k,l) \in \mathcal{F}^r$ and use the convention that this correspondence is an inlier if $j_{k,l}^{r',r} = \text{inlier}$ and accordingly outlier when $j_{k,l}^{r',r} = \text{outlier}$. Denoting all the latent variables representing data association between robot r and other robots by \mathcal{J}^r and considering it to be part of the inference, the probabilistic formulation (5) turns into:

$$p(\mathcal{X}^r, \mathcal{J}^r | \mathcal{Z}^r) \propto p(X^r | Z^r) p\left(\mathcal{X}^{\mathcal{R} \setminus \{r\}} | \mathcal{Z}_{local}^{\mathcal{R} \setminus \{r\}}\right) \prod_{(r',k,l) \in \mathcal{F}^r} p\left(j_{k,l}^{r',r}\right) p\left(u_{k,l}^{r',r} | x_k^{r'}, x_l^r, j_{k,l}^{r',r}\right). \quad (9)$$

We let Σ_{in} and Σ_{out} to represent the covariances corresponding to inlier and outlier distributions, respectively, with $\Sigma_{in} \ll \Sigma_{out}$. The probability $p\left(u_{k,l}^{r',r} | x_k^{r'}, x_l^r, j_{k,l}^{r',r}\right)$ in Eq. (9) can be evaluated, for both $j_{k,l}^{r',r} = inlier$ and $j_{k,l}^{r',r} = outlier$, using Eq. (6).

Assuming for the moment the initial relative poses $T_r^{r'}$ are known, we can jointly infer the robot trajectories and multi-robot data association via Eq. (9). In particular, the MAP estimate over robot states is given by

$$\hat{\mathcal{X}}^r = \arg \max_{\mathcal{X}^r} p(\mathcal{X}^r | \mathcal{Z}^r) = \arg \max_{\mathcal{X}^r} \sum_{\mathcal{J}^r} p(\mathcal{X}^r, \mathcal{J}^r | \mathcal{Z}^r). \quad (10)$$

However, since the above optimization is computationally expensive, we resort to an Expectation-Maximization approach, a single iteration of which can be formulated as

$$\hat{\mathcal{X}}_{(i+1)}^r = \arg \max_{\mathcal{X}^r} \int p\left(\mathcal{J}^r | \hat{\mathcal{X}}_{(i)}^r, \mathcal{Z}^r\right) \log \left[p\left(\mathcal{X}^r | \hat{\mathcal{J}}_{(i)}^r, \mathcal{Z}^r\right) \right] d\mathcal{J}^r, \quad (11)$$

where the notation (i) represents an iteration number.

We note that when the initial relative pose $T_r^{r'}$ between any two robots r and r' is unknown, performing inference over Eqs. (11) or (9) is doomed to fail: since the transformation $T_r^{r'}$ is unknown and can only be arbitrarily set, each candidate multi-robot data association $(r', k, l) \in \mathcal{F}^r$ with a corresponding constraint $u_{k,l}^{r',r}$ will typically result in high discrepancy between $u_{k,l}^{r',r}$ and the prediction $h(x_k^{r'}, x_l^r)$ from Eq. (8), i.e. high errors $err\left(u_{k,l}^{r',r}, x_k^{r'}, x_l^r\right)$. These high errors will be obtained *both for inlier and outlier correspondences*.

Since $\Sigma_{in} \ll \Sigma_{out}$, the probability $p\left(u_{k,l}^{r',r} | x_k^{r'}, x_l^r, j_{k,l}^{r',r}\right)$ will be higher for $j_{k,l}^{r',r} = outlier$ than $j_{k,l}^{r',r} = inlier$, *regardless* if the correspondence $(r', k, l) \in \mathcal{F}^r$ is an inlier or outlier in practice. As a result, attempting to infer both multi-robot data association and robot trajectories will lead to *all* candidate correspondences in \mathcal{F}^r to be identified as outliers. It is for this reason that initial relative poses *must* be first estimated so that the error in Eq. (7) could be used to distinguish between inlier and outlier correspondences.

This observation is illustrated in Figure 3 for the candidate correspondences between the blue and red robots shown in Figure 1a. The figure summarizes the errors (7) for all such correspondences evaluated using *ground truth* value for the initial relative pose transformation $T_r^{r'}$ (Figure 3a) and the *inferred* transformation $\hat{T}_r^{r'}$ at each time step. For each case, the errors are separately shown for inlier and outlier correspondences, determined from ground truth data - the filled areas represent the minimum and maximum of these errors. As seen, when using the true transformation $T_r^{r'}$, the errors substantially differ for inliers and outliers, and therefore they can be distinguished from each other. On the other hand, as evaluating Eq. (7) using the arbitrarily chosen robot reference frames, results in high errors and, more importantly, the inlier and outlier error levels *overlap each other* and therefore the inlier and outliers cannot be easily distinguished. Only after this transformation is correctly established, the errors drop and a natural segmentation into inliers and outliers arises (right area in Figure 3b).

Consequently, we propose first to infer the transformations $T_r^{r'}$ and only then we proceed to infer robot trajectories via the EM optimization (11). In the following sections we describe our approach for estimating the transformations $T_r^{r'}$ without assuming multi-robot data association is given.




Figure 3: Distribution of inlier and outlier correspondences error over time evaluated using (a) ground truth and (b) inferred initial relative pose transformation $T_r^{r'}$. The shown filled areas represent the maximum and minimum norm of translation error in the *actual* inlier and outlier correspondences as determined from ground truth data. Using ground truth $T_r^{r'}$ (a) the inliers can be easily distinguished from outliers. In contrast, (b) shows that for arbitrary value of $T_r^{r'}$ the errors for inlier and outlier correspondences overlap, and only after estimating the transformation $T_r^{r'}$ (around pose index 140), inliers and outliers become distinguishable.

B. Distributed Inference over Initial Relative Pose via EM

Our approach is based on the following *key observation*: given local robot trajectories, *each* candidate multi-robot correspondence $(r', k, l) \in \mathcal{F}^r$, regardless if it is inlier or outlier, suggests a solution for the transformation $T_r^{r'}$. However, *only* the inlier correspondences will produce similar transformations, while those calculated from outlier correspondences will typically disagree amongst each other, *unless* these outliers are caused by measurement aliasing.

We illustrate this concept in a synthetic toy example in Figure 4. Ground truth robot trajectories are given in Figure 4a; candidate matches between two of the robots (red and blue) are shown in Figure 4b, with robot initial poses set to arbitrary values. Many of these correspondences are outliers - in this simple example we assume there is no measurement aliasing. The distribution of the transformation $T_r^{r'}$, calculated for each correspondence, is shown in Figure 4c. One can clearly observe the cluster corresponding to the correct transformation.

Real-world scenarios, however, often exhibit some level of measurement aliasing, which often leads to *multiple* clusters and further complicate the identification of the correct transformation $T_r^{r'}$. Examples are given throughout the paper (e.g. Figures 9 and 10); our approach addresses this challenge as we discuss in the sequel (Section V).

Based on the above key observation, we now formulate an EM-based optimization that allows each robot to recover the initial relative pose with other robots (i.e. establish a common reference frame) in a distributed manner.

The MAP estimate of an initial relative pose between robot r and each robot $r' \in$

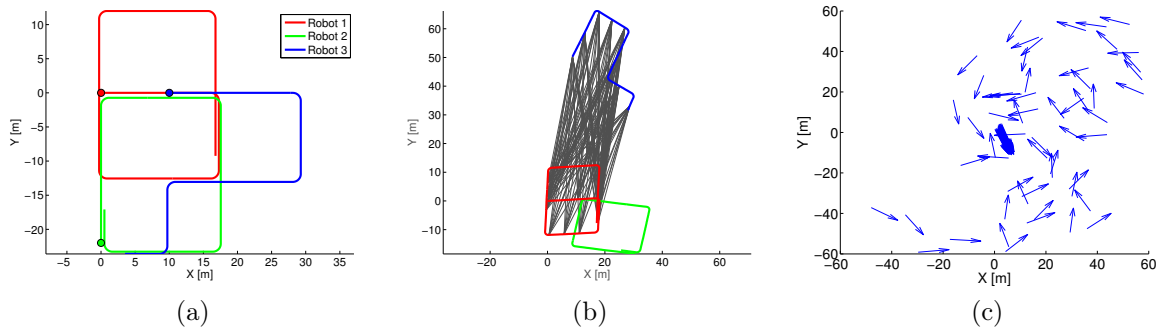


Figure 4: Synthetic toy example illustrating proposed concept *without* considering perceptual aliasing aspects. (a) ground truth trajectories for three robots; (b) candidate matches between blue and red robots. Robot initial poses are unknown and set to arbitrary values. (c) Distribution of the transformation $T_r^{r'}$ between blue and red robots, calculated for each candidate match shown in (b). Planar case is considered (x, y axes correspond to position, arrow direction to orientation). The cluster corresponds to inliers and to correct transformation.

$\mathcal{R} \setminus \{r\}$ can be written as

$$\hat{T}_r^{r'} = \arg \max_{T_r^{r'}} p \left(T_r^{r'} | \hat{\mathcal{X}}^r, \mathcal{Z}^r \right), \quad (12)$$

where, $\hat{\mathcal{X}}^r$ and \mathcal{Z}^r are defined in Eq. (4). This inference can be solved by each robot r in a distributed fashion, based on the available observations \mathcal{Z}^r .

Similar to Eq. (10), we introduce the latent variables \mathcal{J}^r to this inference problem

$$\hat{T}_r^{r'} = \arg \max_{T_r^{r'}} \sum_{\mathcal{J}^r} p \left(T_r^{r'}, \mathcal{J}^r | \hat{\mathcal{X}}^r, \mathcal{Z}^r \right), \quad (13)$$

and re-write the above equation in EM framework:

$$\hat{T}_r^{r'} = \arg \max_{T_r^{r'}} \int p \left(\mathcal{J}^r | \hat{T}_r^{r'}, \hat{\mathcal{X}}^r, \mathcal{Z}^r \right) \log \left[p \left(T_r^{r'} | \hat{\mathcal{J}}^r, \hat{\mathcal{X}}^r, \mathcal{Z}^r \right) \right] d\mathcal{J}^r, \quad (14)$$

where the iteration number was omitted (see Eq. (11)). The underlying EM equations for Eq. (14) are given in the Appendix.

Solving Eq. (14) involves a non-linear optimization that converges to a *local* minima given an initial guess for a solution. As the optimization problem (14) is highly non-convex, an initial solution in the vicinity of the global minima is required to guarantee convergence to the *correct* transformation $T_r^{r'}$.

How to obtain such an initial solution? We recall the key observation from the beginning of this section and using the available robot local trajectories $\hat{\mathcal{X}}^r$ we calculate an initial relative pose for each correspondence $(r', k, l) \in \mathcal{F}^r$ resulting in a set $\mathcal{T}_r^{r'}(\mathcal{F}^r)$. We then analyze the distribution of this set and perform a very naïve clustering for each element in $T_r^{r'}$ separately: we simply choose k most common values that are sufficiently far away from each other. We use $k = 5$ and consider two entries to be sufficiently apart if the difference is above the corresponding std in Σ_{in} : 0.5 meters for position and 0.05 radians for rotation. Instead of considering all the possible combinations for each axis as

initial solution, i.e. k^n where n is the dimensionality of $T_r^{r'}$, we discard any entries that do not have a nearby initial relative pose in $\mathcal{T}_r^{r'}(\mathcal{F}^r)$.

This procedure generates several initial solutions for $T_r^{r'}$; running the EM optimization (14) on each one of these solutions produces locally-optimal initial relative poses $\hat{T}_r^{r'}$. We merge different hypotheses that converge to the same value $\hat{T}_r^{r'}$. Each of these estimates is *optimal* given the corresponding partitioning of the multi-robot data association $\hat{\mathcal{J}}^r$ into inliers I and outliers O , as calculated in the E step in EM, see Eq. (14).

We shall now use the term *hypothesis* $h \doteq \{I, O\}$, with $I \cup O = \mathcal{J}^r$, to represent each such partition and collect all these hypotheses into the set \mathcal{H} . Note that each hypothesis $h \in \mathcal{H}$ leads to an estimate of $T_r^{r'}$ according to

$$\hat{T}_r^{r'}(h) = \arg \max_{T_r^{r'}} p\left(T_r^{r'} | h, \hat{\mathcal{X}}^r, \mathcal{Z}^r\right). \quad (15)$$

Figure 9b illustrates the distribution of relative pose constraints for each correspondence $(r', k, l) \in \mathcal{F}^r$ (planar case is considered: x and y axes correspond to position, arrow direction represents orientation). The transformations $\hat{T}_r^{r'}(h)$ for each hypothesis $h \in \mathcal{H}$ are denoted by red bold arrows.

As there are often more than one hypothesis in \mathcal{H} , one has to choose which hypothesis $h^* \in \mathcal{H}$ to use and set the initial relative pose to $\hat{T}_r^{r'} = \hat{T}_r^{r'}(h^*)$ to be later used in the underlying optimization of Eq. (11). Recalling the discussion from Section A, identifying the *correct* hypothesis in \mathcal{H} is crucial, especially in the presence of perceptual aliasing. In the next section we develop a probabilistic method for choosing the most probable hypothesis, focusing in particular on the complexities introduced by the incremental framework that is essential for online operation.

V. Hypothesis Model-Based Selection

In this section we develop a probabilistic approach for choosing the most likely hypothesis, given the current information from the set \mathcal{H} . An incremental setting, in which information (i.e. robot trajectories, local observations and multi-robot correspondences) is obtained gradually, introduces a number of challenges. How to know sufficient amount of information has been accumulated to reliably estimate the initial relative poses? For example, robot trajectories and observed environments may initially not overlap (or not overlap at all), in which case all multi-robot correspondences are outliers.

Measurement aliasing, i.e. observations of similar but different environments complicate the problem even further: consider the robots start operating in two nearby somewhat *similar* environments (e.g., corridors). Each robot r shares informative measurements, as described in Section IV, match measurements transmitted from other robots against its own informative measurements and introduces the matches as candidate correspondences into the set \mathcal{F}^r . Since the two environments are similar but different, not only these matches will be outliers, but they will also be *consistent* with each other, i.e. suggesting the two environments are actually the same. This translates directly into a hypothesis $h = \{I, O\}$ with many inliers, all of which are actually (consistent) outliers.

Figure 10b shows such a case in a real-world scenario: the robots travel in two different corridors, however because these corridors are similar in appearance many consistent matches between scans from the two corridors are generated, leading to the cluster emphasized by red ellipse).

Note that even after the robots observe some areas in common, the correct hypothesis

(corresponding to true inliers) will have to compete with the consistent-outlier hypothesis. Making a wrong decision and choosing an incorrect hypothesis will lead to catastrophic results as outliers will be introduced into the multi-robot optimization (11). This is demonstrated Section VI-C (Figure 10), where as the result of choosing the wrong cluster (denoted by red ellipse), robot trajectories are incorrectly aligned (ground truth trajectories are shown in Figure 8a) and all the true inlier correspondences are considered as outliers.

We approach this challenging problem within a model selection framework, aiming to calculate the probability of each hypothesis h in the set \mathcal{H} : $p(h|\hat{\mathcal{X}}^r, \mathcal{Z}^r)$. This probability can be obtained by integrating over all possible values of the continuous random variable $T_r^{r'}$:

$$p(h|\hat{\mathcal{X}}^r, \mathcal{Z}^r) = \int_{T_{r_2}^{r_1}} p(T_r^{r'}, h|\hat{\mathcal{X}}^r, \mathcal{Z}^r) \quad (16)$$

Applying Bayes rule one can write

$$p(h|\hat{\mathcal{X}}^r, \mathcal{Z}^r) = \int_{T_{r_2}^{r_1}} c \cdot p(\mathcal{Z}^r|T_{r_2}^{r_1}, h, \hat{\mathcal{X}}^r) p(T_r^{r'}, h|\hat{\mathcal{X}}^r) \quad (17)$$

with a hypothesis-independent constant $c \doteq 1/p(\mathcal{Z}^r|\hat{\mathcal{X}}^r)$ that needs not be calculated. Factorizing the last term as

$$p(T_r^{r'}, h|\hat{\mathcal{X}}^r) = p(T_r^{r'}|h, \hat{\mathcal{X}}^r) p(h|\hat{\mathcal{X}}^r), \quad (18)$$

Eq. (17) turns into

$$p(h|\mathcal{Z}^r, \hat{\mathcal{X}}^r) = c \cdot p(h|\hat{\mathcal{X}}^r) p(\mathcal{Z}^r|h, \hat{\mathcal{X}}^r), \quad (19)$$

with

$$p(\mathcal{Z}^r|h, \hat{\mathcal{X}}^r) = \int_{T_r^{r'}} p(T_r^{r'}|h, \hat{\mathcal{X}}^r) p(\mathcal{Z}^r|T_r^{r'}, h, \hat{\mathcal{X}}^r). \quad (20)$$

To calculate hypothesis probability (17), each of the terms in Eqs. (19)-(20) needs to be evaluated: the hypothesis prior $p(h|\hat{\mathcal{X}}^r)$, relative-pose transformation prior $p(T_r^{r'}|h, \hat{\mathcal{X}}^r)$, and the measurement likelihood term $p(\mathcal{Z}^r|T_r^{r'}, h, \hat{\mathcal{X}}^r)$.

We proceed by first discussing the measurement likelihood term. Recalling the hypothesis definition, $h = \{I, O\}$, and writing down individual inlier and outlier observation terms we get

$$p(\mathcal{Z}^r|T_r^{r'}, h, \hat{\mathcal{X}}^r) = k(h) \prod_{i \in I} \exp\left(-\frac{1}{2} \left\| \text{err}(u_i, \hat{\mathcal{X}}^r, T_r^{r'}) \right\|_{\Sigma_{in}}^2\right) \cdot \prod_{o \in O} \exp\left(-\frac{1}{2} \left\| \text{err}(u_o, \hat{\mathcal{X}}^r, T_r^{r'}) \right\|_{\Sigma_{out}}^2\right), \quad (21)$$

where we used u_i and u_o to represent the appropriate relative pose constraints $u_{k,l}^{r',r}$ and $\text{err}(\cdot)$ is defined in Eq. (7). The hypothesis-dependent coefficient $k(h)$ is defined as

$$k(h) \doteq \prod_{i \in I} \frac{1}{\sqrt{|2\pi\Sigma_{in}|}} \prod_{o \in O} \frac{1}{\sqrt{|2\pi\Sigma_{out}|}}. \quad (22)$$

Considering $p(T_r^{r'}|h, \hat{\mathcal{X}}^r)$ can be expressed as the Gaussian

$$p(T_r^{r'}|h, \hat{\mathcal{X}}^r) = N(T_0, \Sigma_0). \quad (23)$$

with known parameters T_0 and Σ_0 (as discussed in Section B), the integrant $q(T_r^{r'})$ in Eq. (20),

$$q(T_r^{r'}) \doteq p(T_r^{r'}|h, \hat{\mathcal{X}}^r) p(Z|T_r^{r'}, h, \hat{\mathcal{X}}^r), \quad (24)$$

can be expressed using Eqs. (21)-(22) and (23). Since the involved distributions are all Gaussians, we can represent this expression by a single Gaussian

$$q(T_r^{r'}) = N(\hat{T}_r^{r'}, \Sigma_{MAP}), \quad (25)$$

where the MAP estimate is obtained by $\hat{T}_r^{r'} = \arg \max_{T_r^{r'}} q(T_r^{r'})$, and can be then used as a linearization point to calculate the covariance Σ_{MAP} .

Recall the relation

$$\int_x a \exp(-\|x - \hat{x}\|_{\Sigma_{MAP}}^2) = a \sqrt{|2\pi \Sigma_{MAP}|}, \quad (26)$$

where the notations \hat{x} and Σ_{MAP} represent, respectively, the MAP estimate and covariance of some random variable x .

Substituting Eqs. (24)-(25) back to Eq. (19) and using the above relation we get the following expression for the hypothesis probability:

$$p(h|\mathcal{Z}^r, \hat{\mathcal{X}}^r) = c \cdot p(h|\hat{\mathcal{X}}^r) k'(h) \sqrt{|2\pi \Sigma_{MAP}|}, \quad (27)$$

where $k'(h) \doteq k(h) \cdot \frac{1}{\sqrt{|2\pi \Sigma_0|}}$.

The remaining missing parts are the priors $p(h|\hat{\mathcal{X}}^r)$ and $p(T_r^{r'}|h, \hat{\mathcal{X}}^r)$. We now discuss each of these terms, and then elaborate on our approach for selecting the most likely hypothesis $h^* \in \mathcal{H}$, given the calculated hypotheses probabilities.

A. Hypothesis Prior $p(h|\hat{\mathcal{X}}^r)$

Hypothesis prior is an important component in deciding which hypothesis to choose from the set \mathcal{H} . While the measurement likelihood $p(\mathcal{Z}^r|T_r^{r'}, h, \hat{\mathcal{X}}^r)$ essentially prioritizes hypotheses with higher number of inliers (since $R_{in} \ll R_{out}$, see Eq. (22)), it does not address the question whether sufficient amount of information has been accumulated to make a decision. In other words, given a set \mathcal{H} there is always a hypothesis with the highest measurement likelihood (i.e. highest number of inliers); how to decide if that hypothesis is unambiguous and should be indeed chosen?

This aspect is particularly crucial in the *incremental* setting in the context of *measurement aliasing* that can lead to a dominant hypothesis of consistent outliers (see Section V): Note that each correspondence, regardless if it is considered inlier or outlier by a hypothesis, is the result of a high-quality match between two observations - in our case, ICP match between two laser scans. Therefore, the main reason for a correspondence to be an outlier *in practice*, i.e. highly-confident match between two scans from different

areas, is perceptual aliasing. Relying only on the measurement likelihood term, it is easy to mistakenly choose this incorrect hypothesis.

To address this crucial issue we argue the hypothesis prior, $p(h|\hat{\mathcal{X}}^r)$, can provide insight as to how likely is the hypothesis $h = \{I, O\}$ is in the first place. Given the robot local trajectories, one can reason about the partitioning of the set \mathcal{F}^r into inlier and outlier sets I and O , respectively, even without the actual measurements of the corresponding constraints. Additionally, we introduce into an additional hypothesis $h_0 = \{I_0, O_0\}$ into the set \mathcal{H} : this hypothesis, that we denote as the *null-hypothesis*, corresponds to perceptual aliasing, i.e. it represents the possibility that *all* of the correspondences are actually outliers ($I_0 \equiv \phi$).

The prior probabilities all hypotheses in \mathcal{H} , including the null-hypothesis h_0 , can now be calculated as we describe next. The most likely hypothesis h , according to the posterior $p(h|\hat{\mathcal{X}}^r, \mathcal{Z}^r)$ from Eq. (27) is then chosen *only* if its prior probability $p(h|\hat{\mathcal{X}}^r)$ is significantly dominant compared to all other hypotheses in \mathcal{H} . That way we wait until sufficient information is accumulated to disambiguate which hypothesis is the correct one.

The underlying basic assumption in our approach is that since the robots are operating in closed indoor environments, they will eventually observe common places. Each given candidate correspondence can therefore represent the same place, observed by two different robots, or two different places. However, the number of unique places is unknown ahead of time. To capture this probabilistically, we resort to the Chinese restaurant process (and the closely related Dirichlet process), see e.g. [4, 33], which has been recently used for topological mapping [29]. Next, we provide a brief overview of the Chinese restaurant process (Section V-A.1), and then in Section V-A.2 describe the details of our approach.

1. Chinese Restaurant Process

According to the Chinese Restaurant Process, the probability of observing a new place, after previously observing n unique places, is given by

$$\frac{\alpha}{\alpha + n - 1}, \quad (28)$$

where α is a concentration parameter. This parameter defines the extent in which repeated observations of the same place take place - larger α corresponds to higher probability of observing a *new* place; see Figure 5a that describes the evolution of the probability in Eq. (28) as a function of n and α . One can observe from the figure and from Eq. (28) that this probability decreases with n for a given α .

Based on Eq. (28), it is straightforward to show [29] that, given n observations, the probability of observing n unique places can be written as:

$$f(n) \doteq \frac{\alpha^n}{\prod_{j=1}^n (\alpha + j - 1)}. \quad (29)$$

Figure 5b illustrates this probability as a function of n and for different values of α . As seen, this probability decreases with n for a given α , corresponding to our above assumption that observing only new places becomes less likely as we make more observations.




Figure 5: Chinese restaurant process: (a) probability of observing a new place and (b) probability of observing n unique places for different values of α . See text for more details.

2. Hypothesis Prior Calculation via Chinese Restaurant Process

We will now make use of the above model (29) and develop an approach for quantifying the hypothesis prior for each $h \in \mathcal{H}$.

Our approach is based on the following interpretation of inlier and outlier correspondences, as determined by a given hypothesis $h = \{I, O\} \in \mathcal{H}$: an inlier $i \in I$ relating two different poses of robots r and r' indicates the two poses are close to each other (generating overlapping observations); accordingly, the involved two poses in an outlier correspondence $o \in O$ are considered as two different, unique, places.

More formally, we let m_{in} and m_{out} represent, respectively, the number of inliers and outliers in a given hypothesis $h = \{I, O\} \in \mathcal{H}$ and denote by m the overall number of candidate correspondences. Note that m is the same for all hypotheses in \mathcal{H} , and also $m = m_{in} + m_{out}$.

Given some hypothesis $h = \{I, O\} \in \mathcal{H}$, we approximately capture the number of unique places due to inlier correspondences as m_{in} , i.e.: each inlier correspondence is modeled to contribute a single unique place: as mentioned in Section IV, a new informative measurement is shared by each robot r only if it has moved certain distance since it shared the previous informative measurement. Therefore any two inlier correspondences $i_1, i_2 \in I$, even they involve sequentially shared measurements, will contain some unique information regarding the environment (since the robot moves between one shared measurement to another). Moreover, only a small portion of the measurements pairs (originating from different robots) are successfully matched and added as candidate correspondences. Consequently, the number of unique observed places due to inlier correspondences is modeled as m_{in} .

Similarly, we consider any two pairs of outlier correspondences $o_1, o_2 \in O$ to represent four different places, and thus approximately capture the number of unique observed places due to outliers as $2m_{out}$. We make yet another approximation at this point, and consider the unique places contributed by inliers and outliers do not overlap.

Based on the above, we model the number of unique places for each given hypothesis as

$$n(h) = 2m_{out}(h) + m_{in}(h), \quad (30)$$

where the parameters m_{out} and m_{in} are hypothesis-specific. We leave the investigation of methods to relax the above assumptions and capturing more accurately the number of unique places for each hypothesis to future work.

Observe that while the number of involved pairs of cells, m , is the *same* for all hypotheses in \mathcal{H} , the number of unique places, n , changes from one hypothesis to another in the range $n \in [m, 2m]$. In particular, n always assumes the highest value for the null-hypothesis ($m_{out} = m, m_{in} = 0$) and, on the other extreme, the lowest value for all-inliers hypothesis ($m_{out} = 0, m_{in} = m$), if such a hypothesis exists in \mathcal{H} .

The prior probability $p(h|\hat{\mathcal{X}}^r)$ of each hypothesis $h \in \mathcal{H}$ can now be calculated from Eq. (29) using the hypothesis-specific value for n that is calculated according to Eq. (30). We further normalize the prior probabilities of all hypotheses in \mathcal{H} by

$$\eta \doteq \sum_{h \in \mathcal{H}} f(n(h)) \quad (31)$$

to obtain a valid probability distribution:

$$p(h|\hat{\mathcal{X}}^r) = \frac{f(n(h))}{\eta}. \quad (32)$$

The above provides a mechanism to decide if sufficient information has been accumulated to indicate a certain hypothesis is dominant with respect to competing hypotheses while explicitly accounting for potential measurement aliasing via the introduction of the null-hypothesis.

In particular, for a *given* m and α , hypotheses with many inliers span different parts of robot trajectories. These hypotheses get higher probabilities since the number of unique places n is smaller for such hypotheses: the two extreme cases are hypotheses with only inlier or only outlier correspondences, corresponding, respectively, to $n = m$ and $n = 2m$. In general, the parameter n for some hypothesis is given by γm , with $1 \leq \gamma \leq 2$. One can verify from Eq. (29) that $f(\gamma m)$ decreases with γ (see also Figure 5b). The intuition here is that, at each given time (i.e. fixed m), a hypothesis whose inliers involve larger portions of robot trajectories (smaller n) is less prone to measurement aliasing over a hypothesis whose inliers only involve a small area (larger n).

How one determines if sufficient information has been obtained so that a decision regarding the most probable hypotheses can be made? To address this question we recall the Chinese restaurant model (29) and formulate the following lemma.

Lemma. *The ratio between hypotheses prior probabilities increases as more information is accumulated, i.e. as more unique places are observed (n increases).*

Proof. We consider any two different hypotheses $h_i, h_j \in \mathcal{H}$ and denote the number of unique places in these hypotheses by n_i and n_j , with $n_i, n_j \in [m, 2m]$. Without losing generality, we assume $n_i < n_j$ and therefore from the discussion above: $p(h_i|\hat{\mathcal{X}}^r) > p(h_j|\hat{\mathcal{X}}^r)$. Since both probabilities are calculated using the same normalization constant η (Eq. (31)): $f(n_i) > f(n_j)$.

Now, we can always find parameter $\beta \in [1, 2]$, such that $n_j = \beta n_i$ and therefore express the ratio $f(n_i)/f(n_j) \equiv f(n_i)/f(\beta n_i)$ via Eq. (29) as follows:

$$\begin{aligned} f(n_i)/f(\beta n_i) &\doteq \frac{\alpha^{n_i}}{\prod_{j=1}^{n_i} (\alpha + j - 1)} \cdot \frac{\prod_{j=1}^{\beta n_i} (\alpha + j - 1)}{\alpha^{\beta n_i}} = \\ &= \frac{1}{\alpha^{(\beta-1)n_i}} \prod_{j=n_i+1}^{\beta n_i} (\alpha + j - 1) = \frac{1}{\alpha^{(\beta-1)n_i}} \frac{(\alpha + \beta n_i - 1)!}{(\alpha + n_i)!}. \end{aligned}$$

Since $\beta \in [1, 2]$ and $\alpha > 1$, the above expression monotonically increases with n_i , as shown in Figure 6. However, increasing n_i corresponds to accumulating more information, i.e. adding new candidate correspondences to the set \mathcal{F} . Therefore, with time, the prior for any two given hypotheses becomes more distinguishable. \square

The implication of the above Lemma is that, as more information is accumulated, it becomes possible to disambiguate between the different hypotheses. The parameter α determines how fast this process is, and therefore this parameter can be considered as tuning parameter to be set according to prior knowledge regarding the environment expected size, if such knowledge exists. In particular, too small value of α will lead to premature down-weighting of the null hypothesis, thereby increasing sensitivity to measurement aliasing. On the other hand, larger values of α will require accumulating more information before the prior of the appropriate hypothesis becomes sufficiently dominant to facilitate reliable decision.

In our current implementation, we wait until the (normalized) prior of a hypothesis is $p(h|\hat{\mathcal{X}}^r) \geq 0.8$ before considering sufficient data exists to make a decision. The final decision is of course made based on the overall hypothesis probability $p(h|\hat{\mathcal{X}}^r, \mathcal{Z}^r)$, as described in Section V.

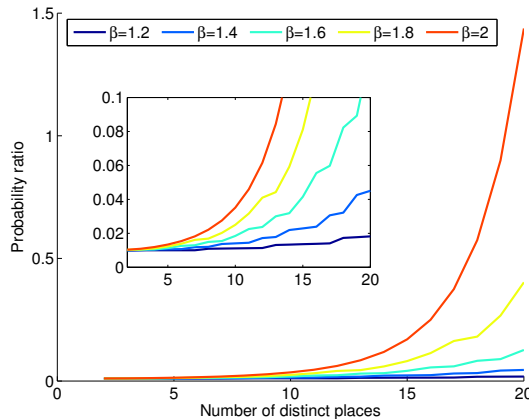


Figure 6: Probability ratio $f(n)/f(\beta n)$ with $1 \leq \beta \leq 2$.

This concept is illustrated in a toy example in Figure 7 that considers 3 different hypotheses that are detailed in Table 1: the null hypothesis (h_0), hypothesis with half of the correspondences inliers (h_1), and hypothesis with only inliers (h_2). The figure shows the evolution of the hypotheses prior as a function of m . One can clearly observe that the prior probability of the null-hypothesis decreases as m increases, and the prior for h_2 (all

inliers) gradually increases. The gap between the latter and the other two hypotheses (h_0 and h_1) drastically increases as well, corresponding to higher confidence level in choosing h_2 as a valid hypothesis candidate.




Figure 7: Prior probabilities of 3 different hypotheses h_0 , h_1 and h_2 (see Table 1) as a function of m . Larger m corresponds to accumulation of additional information. Parameter α is set to 50 in all cases.

h	Description	m_{in}	m_{out}	n
h_0	null-hypothesis	0	m	$2m$
h_1	half inliers, half outliers	$0.5m$	$0, 5m$	$1.5m$
h_2	all inliers	m	0	m

Table 1: Considered three different hypotheses in toy example from Figure 7. See text for more details.

B. Initial Relative Pose Prior $p(T_r^{r'} | h, \hat{\mathcal{X}}^r)$

Given the inlier and outlier correspondences in \mathcal{F}^r but without the actual measurements, our ability to say anything informative regarding the prior on the initial relative pose between the two robots, $p(T_r^{r'} | h, \hat{\mathcal{X}}^r)$, is very limited.

Although, in principle, one could consider about what values of $T_r^{r'}$ are unreasonable, we currently employ a basic approach: We assume that because of the way constraints are generated (see Section IV), on average, the relative-pose measurements $u_{k,l}^{r',r}$ are zero-biased and therefore consider each measurement likelihood term to be distributed according to $N(0, \Sigma_{high})$ where Σ_{high} is a high uncertainty covariance (we use 10 meters in position and 90 degrees in rotation). We therefore model $p(T_r^{r'} | h, \hat{\mathcal{X}}^r)$ as the Gaussian:

$$p(T_r^{r'} | h, \hat{\mathcal{X}}^r) = \prod_{(r',k,l) \in \mathcal{F}^r} N(0, \Sigma_{high}) \doteq N(T_0, \Sigma_0).$$

C. Choosing the Most Probable Hypothesis

Decision on the most probable hypothesis is made as follows. We calculate the probabilities Eq. (27) of only the most promising hypotheses, identified by their support (number of

inliers). The highest probability hypothesis h is then chosen only if it satisfies two conditions: (a) its posterior probability is sufficiently higher compared to all other hypotheses, (b) its prior probability is significantly dominant with respect to other hypotheses - in our implementation we require $p(h|\hat{\mathcal{X}}^r, \mathcal{Z}^r) \geq 0.8$. For numerical reasons, computation of $p(h|\mathcal{Z}^r, \hat{\mathcal{X}}^r)$ is performed in the log space.

VI. Results

The developed approach was implemented within the GTSAM optimization library [10], and evaluated in real-world experiments involving three quadrotors, equipped with laser scanners, operating in indoor environments. In this evaluation we are in particular interested in the following metrics:

- Estimation accuracy of initial relative poses $T_r^{r'}$ between any two robots r and r' .
- Statistics of inferred multi-robot data association: we evaluate the percentage of correctly identified inliers and outliers, as well as false negatives (inliers that were identified as outliers), and false positives (outliers that were identified as inliers).
- Robustness to measurement aliasing: we are interested in demonstrating our approach chooses the correct hypothesis in presence of measurement aliasing in real-world scenarios.
- Exhibit ability of each robot to infer its own and other robots' trajectories, once appropriate initial relative poses $T_r^{r'}$ are established.

As the initial relative poses between the robots are unknown, the initial pose of each robot was set to an arbitrary value (see Figure 9a). Each robot r executes the same algorithm, as discussed in Sections IV and V: it shares highly informative laser scans, receives scans from other robots and calculates relative pose constraints between these scans and its own informative scans using ICP. Saliency of a scan is calculated according to its auto-covariance [28]. These constraints are stored in a separate graph and appropriate entries are added to the multi-robot data association set \mathcal{F}^r . In all cases the parameter α (Section V-1) was set to 500. Unless otherwise specified, robot local observations do *not* include loop closure constraints.

A. Datasets

The presented approach was evaluated in three real-world indoor dataset, **D1**, **D2** and **D3**. Reference trajectories, color-coded according to robot number (red, green, blue), are shown in Figure 8 along with the laser scans of the first robot (red). In two of the datasets, **D1** and **D2**, the robots start operating from the same location. In the third dataset, **D3**, the robots start operating from different locations.

We proceed by presenting a detailed performance analysis for the first dataset, **D1**, and discuss results for the other two datasets in Section VI-D.

B. Dataset D1

In this dataset the robots start moving from the same position, with the red and blue robots moving counterclockwise, and green robot moving in clockwise direction (see

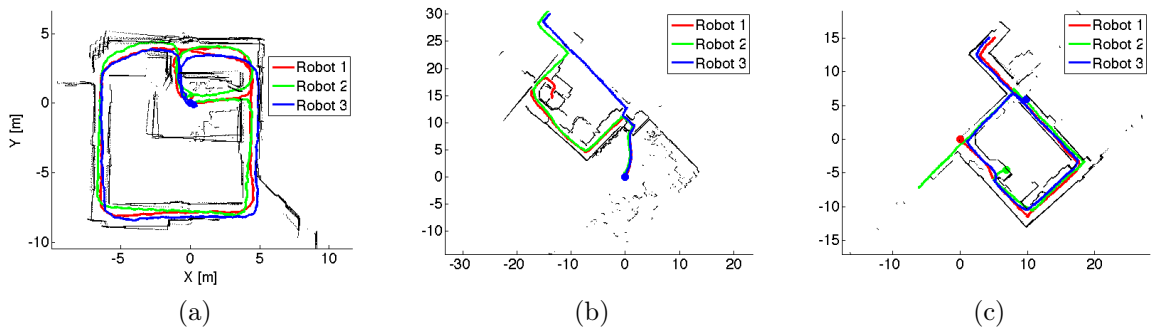


Figure 8: Ground truth for the datasets (a) **D1**, (b) **D2** and (c) **D3**. Robot initial positions are denoted by filled circle marks.

ground truth trajectories in Figure 8a. In contrast to the other datasets, for this dataset we *manually* identified loop-closures for each of the three robots.

Figure 9 details the algorithm steps considering establishing initial relative pose transformation $T_r^{r'}$ between the red and blue robots, i.e. $r = red$ and $r' = blue$. Figure 9a shows the candidate correspondences between the two robots that are available to the red robot, the set \mathcal{F}^r , just before the transformation $T_r^{r'}$ is estimated.

Recall the set $\mathcal{T}_r^{r'}(\mathcal{F}^r)$ represents the initial relative poses that are calculated for each of the candidate correspondences in \mathcal{F}^r (Section IV-B). Figure 9b shows the distribution of this set $\mathcal{T}_r^{r'}(\mathcal{F}^r)$, with each arrow representing a single planar transformation $T_r^{r'} \in \mathcal{T}_r^{r'}(\mathcal{F}^r)$ in terms of x-y coordinates and providing the orientation θ by the arrow angle.

The hypotheses probabilities of the identified hypotheses that correspond to the most dominant clusters in Figure 9b are shown in Figure 9c. The probability of the null-hypothesis h_0 is shown as well. The figure shows both the posterior probability $p(h|\hat{\mathcal{X}}^r, \mathcal{Z}^r)$ (upper part) and the prior probability $p(h|\hat{\mathcal{X}}^r)$ (bottom part); see Eqs. (27) and (32). The posterior probability is given in log-space and without the constant c . The threshold on hypothesis prior for selecting the most probable hypothesis (0.8) is shown in red dashed line (see Section V-C).

As seen, in this case there are five hypotheses in the set \mathcal{H} . Since the posterior for h_1 is higher compared to other hypotheses in \mathcal{H} and its prior is above the threshold, this hypothesis is selected: $h^* \doteq h_1$.

Inlier and outlier correspondences for the chosen hypothesis $h_1 = \{I, O\}$ and for another hypothesis (h_2) are shown in Figures 9e and 9f.

Once a hypothesis is chosen, it becomes possible for robot r (red robot) to perform the multi-robot EM optimization (11) based on all the multi-robot constraints in \mathcal{F}^r that involve robot r' (blue robot), and possibly other robots with an established common frame. The r' th robot poses in $\hat{\mathcal{X}}^r$ (see Eq. (4)), can now be expressed in the reference frame of robot r using $\hat{T}_r^{r'}(h^*)$. As a result, the inlier and outlier correspondences errors (7) no longer overlap each other (right part of Figure 3a), allowing our EM approach to successfully identify the inliers. A correspondence $(r', k, l) \in \mathcal{F}^r$ is identified as inlier if $p(j_{k,l}^{r',r}|\hat{\mathcal{X}}^r, \mathcal{Z}^r)$ approaches 1, see Eq. (11). As shown in Figure 9d, in this case our method successfully identified all inliers; in general the method tends to infer majority of the inliers. Furthermore, the method did not produce any false positive decisions (identified inliers which are outliers in practice).

The result of this EM optimization (11) is shown in Figure 2a, with the identified inliers indicated in black and outliers in gray. One can observe the robot trajectories are properly aligned to each other, compared to the reference trajectories in Figure 8a. Estimation errors of the common reference frame between the two robots are given in Table 2.

Dataset	r, r'	Red		Green		Blue	
Est. error:		$\ t\ $	$ \theta $	$\ t\ $	$ \theta $	$\ t\ $	$ \theta $
D1	Red	-	-	0.5	6.0	0.4	4.7
	Green	0.4	6.5	-	-	0.4	10.9
	Blue	0.4	4.9	0.3	10.4	-	-
D2	Red	-	-	0.1	1.4	0.3	1.8
	Green	0.1	1.2	-	-	0.4	3.2
	Blue	0.35	2.3	0.4	3.5	-	-
D3	Red	-	-	0.6	0.6	2.1	3.3
	Green	0.5	0.3	-	-	1.3	8.2
	Blue	2.7	3.4	2.9	8.6	-	-

Table 2: Initial relative pose estimation errors in the three real-world datasets **D1**, **D2** and **D3**. Estimation errors are reported in terms of norm of translation error ($\|t\|$) in meters, and absolute value of orientation error $|\theta|$ in degrees.

From that moment on, new multi-robot constraints generated between these two robots are directly added to the graph and optimized via Eq. (11). Results of this optimization for several time instances are shown in Figures 2a-2c (only the identified inliers are shown in the last 2 figures). Similar results are obtained by the blue robot ($r = 3$) for $r' = 1$, as shown in Figures 2d-2f.

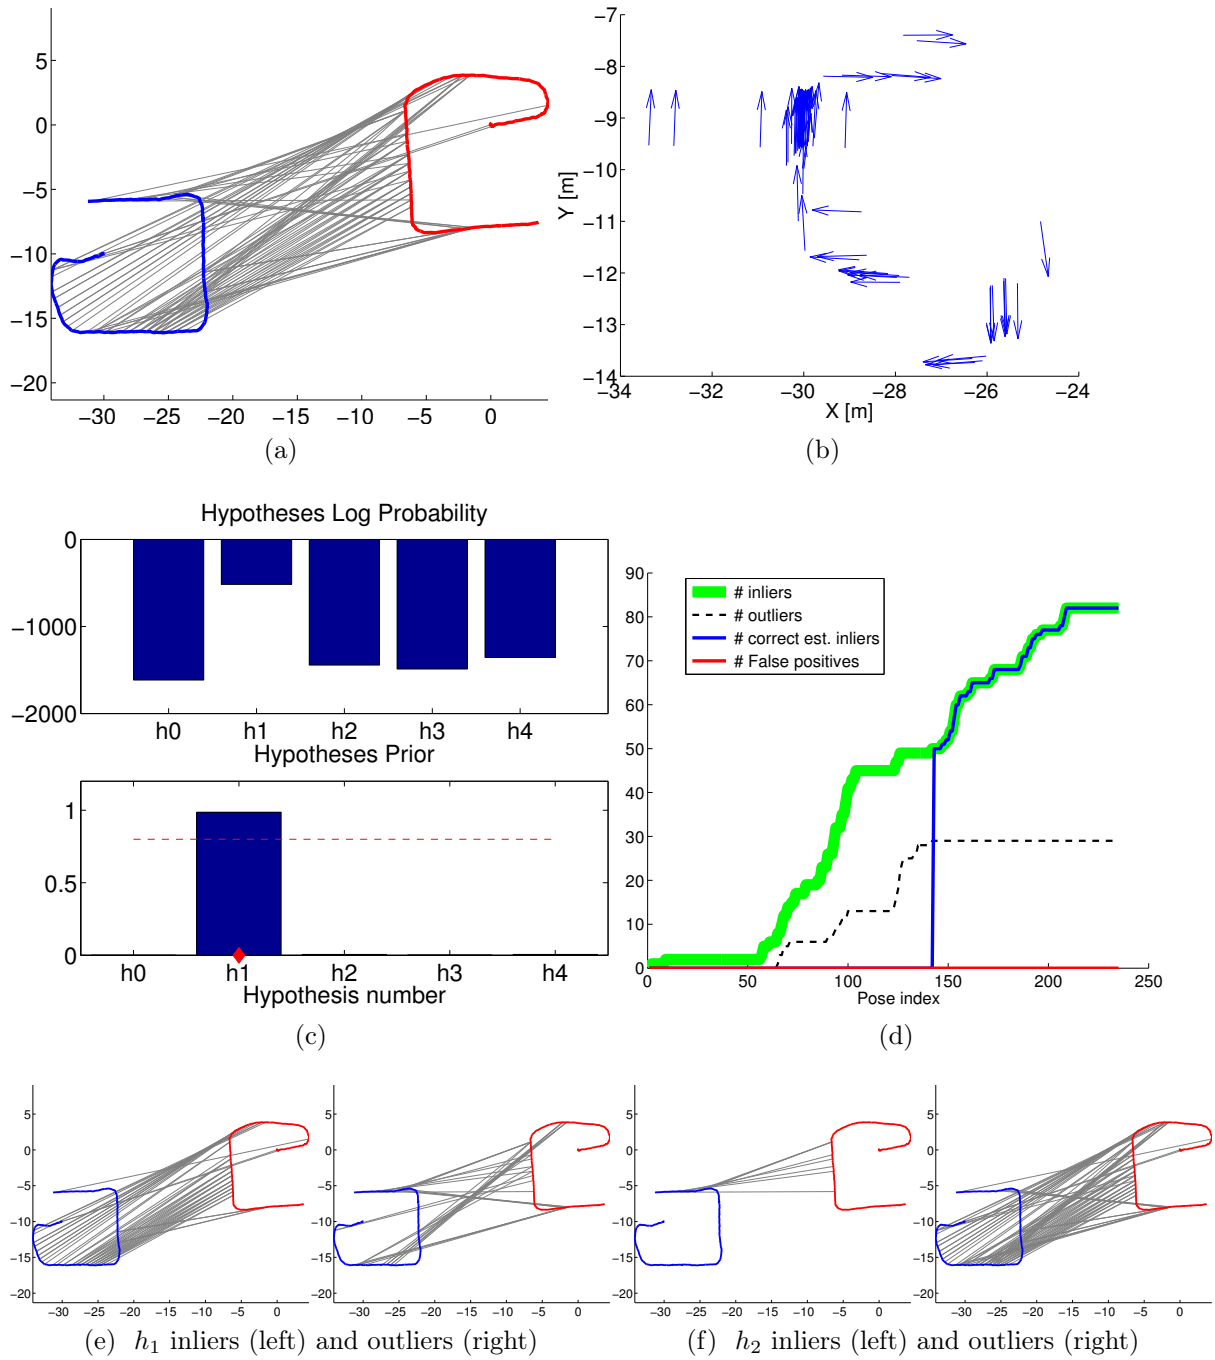


Figure 9: Dataset **D1**: (a) Candidate correspondences between the red and blue robots. Initial poses of these robots are set to arbitrary values. (b) Distribution of the set \mathcal{T}_r' (\mathcal{F}^r) for red and blue robots. (c) Hypothesis probability (in log-space) and prior. The hypothesis h_1 , corresponding to the dominant cluster in (b), is chosen. (d) Actual and inferred inliers and outliers. The method correctly identified all inliers. (e) and (f): Inlier and outlier correspondences of the chosen hypothesis h_1 and of another hypothesis (h_2).

C. Robustness to Measurement Aliasing

Establishing an initial relative pose and multi-robot data association for the green robot is more challenging. It travels in opposite direction (clockwise) with respect to the other two robots and therefore *does not observe the same areas* with these robots until the very end (top right area in Figure 8a). Moreover, since the robots operate in similar environments (e.g. two corridors), *measurement aliasing* causes laser scans from these different environments to be successfully matched^a suggesting the robots are actually observing the same environment. As a result, a cluster of *consistent outliers* is obtained, and only after the robots actually observe the same environments, the cluster corresponding to the correct transformation is formed.

Without using a hypothesis prior $p(h|\hat{\mathcal{X}}^r)$, there is nothing to prevent from choosing the consistent-outlier hypothesis, which can lead to catastrophic result as shown in Figure 10: Figure 10a and 10b show, respectively, the candidate correspondences \mathcal{F}^r between green and red robots and the equivalent distribution of the initial relative poses (the set $\mathcal{T}_r^{r'}$ (\mathcal{F}^r)). The cluster corresponding to *consistent outliers* (Section V) is indicated by a red ellipse. The inliers I of this consistent-outlier hypothesis $h = \{I, O\}$ are shown in Figure 10d on top of the robot trajectories expressed with arbitrary initial relative pose (since it is unknown); the same inliers are shown in Figure 10e on top of ground truth trajectories. As seen, all these correspondences are erroneously considered as inliers, as they relate between different areas. Using the chosen (incorrect) hypothesis and performing the EM optimization (11) produces the result given in Figure 10c, where the trajectories of the green and red robots are incorrectly aligned.

In contrast, incorporating the hypothesis prior $p(h|\hat{\mathcal{X}}^r)$ prevents making this incorrect decision: by introducing the null-hypothesis and modeling the probability of observed unique places for each hypothesis, as described in Section A, the priors for different hypotheses (including consistent-outlier and null hypotheses) compete with each other and only after accumulating more observations by covering additional areas, the prior corresponding to a hypothesis with the highest number of inliers exceeds the threshold and therefore is selected. However, by that time the robots have already observed common areas and the chosen hypothesis is indeed the correct one.

The decision when to consider sufficient information is accumulated, described by a hypothesis prior exceeding a threshold, depends on the value of the parameter α (Section 1), as it determines, given the competing hypotheses in \mathcal{H} , how fast the prior of a hypothesis becomes dominant. As discussed in Section 2, α should be set according to some problem-specific knowledge, such as expected size of the environment to be travelled: in our experience, it is better to use conservative (large) values of α , to reduce chances of making the incorrect decision.

We illustrate the time evolution of the hypotheses priors the green and red robots in Figures 11 and 12 ($\alpha = 500$ in all cases). Three different time instances, represented by indices 4778, 5690 and 6081, are considered. Figure 11a shows the candidate correspondences in these time instances on top of the evolving robot trajectories; Figure 11b details posterior (in log space) and prior probabilities of the identified hypotheses. As seen, the hypothesis h_1 has the highest posterior probability starting from time index 5690, but

^aWe consider two scans to be successfully matched if the following conditions are satisfied: (a) the ratio of nearby points between the first scan and the second scan, transformed using an ICP-calculated relative pose, is above a threshold; we use 0.95 for this threshold. (b) covariance of ICP-calculated relative pose is smaller than threshold (we use 0.4 meters for position and 4 degrees for rotation).

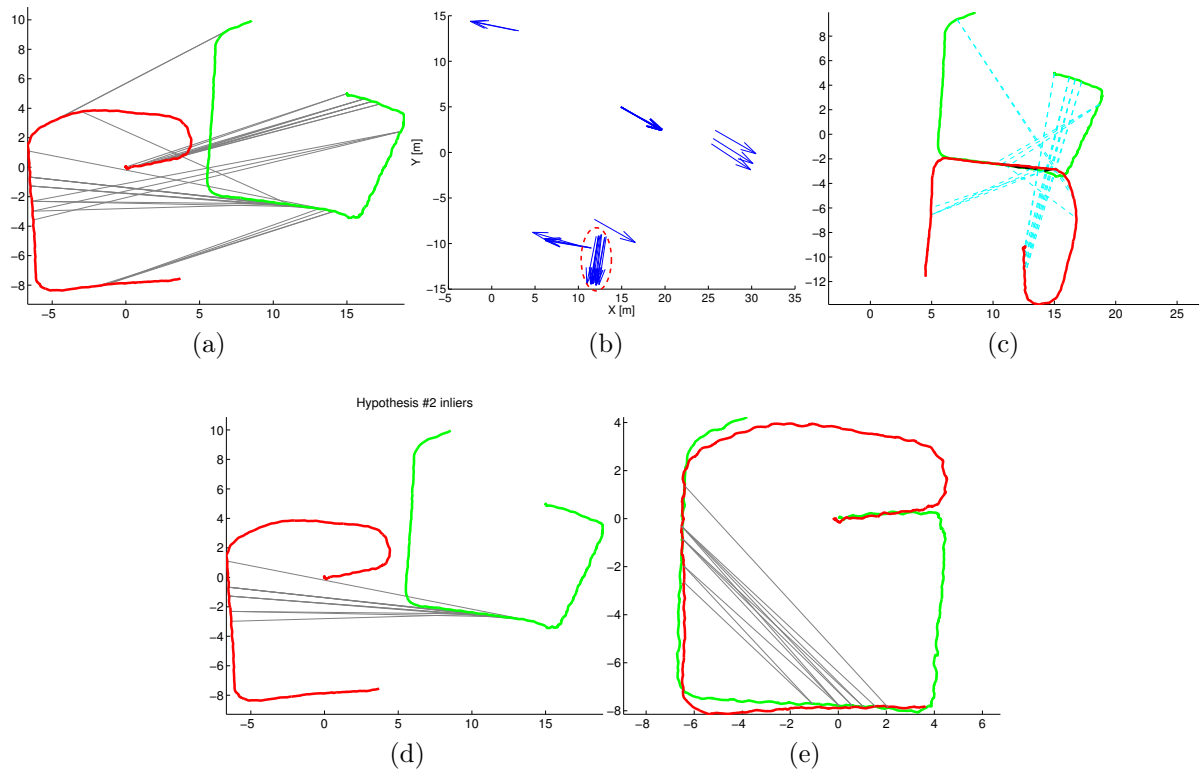


Figure 10: Dataset **D1**: Effect of *measurement aliasing* when *not* using hypothesis prior. (a)-(b) Candidate correspondences \mathcal{F}^r and distribution of set $\mathcal{T}_r^{r'}$ (\mathcal{F}^r) for green and red robots ($r = 2$, $r' = 1$). Cluster of consistent outliers is emphasized. (c) EM optimization result using incorrect hypothesis. Robot trajectories are erroneously aligned. (d) Inliers of the chosen hypothesis (all outliers in practice) drawn on top of robot trajectories with arbitrary initial pose, and (e) on top of ground truth trajectories.

its prior reaches the threshold only at time index 6081. Therefore, only at that time this hypothesis is chosen.

Figure 12 provides further details for the underlying process at time index 6081. The inlier and outlier correspondences for the null-hypothesis^b h_0 , the chosen hypothesis h_1 and another hypothesis h_2 are shown in Figures 12a-12e. Table 3 further details the different parameters described in Section 2 that are used for calculating the prior probability: number of inliers m_{in} and outliers m_{out} , total number of correspondences $m = m_{in} + m_{out}$, number of unique places n . Finally, the result of our EM optimization using the correctly chosen hypothesis is shown in Figure 12f.

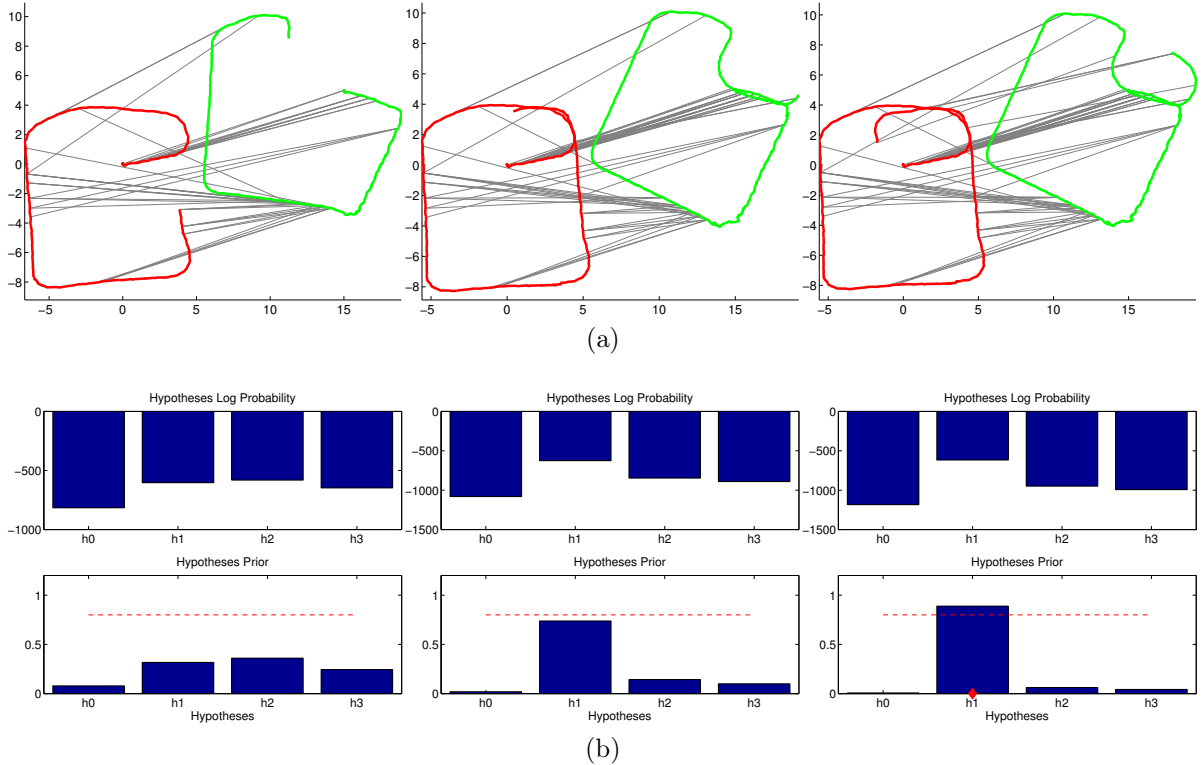


Figure 11: Dataset **D1** - Establishing common reference frame between green and red robots ($r = 2$, $r' = 1$): (a) Candidate correspondences \mathcal{F}^r between green and red robots at three different time instances: indices 4778, 5690 and 6081. (b) Hypotheses posterior (first row) and prior (second row) probabilities for the three time indices. Posterior probability is shown in Log space. First hypothesis represents the null-hypothesis. At time index 6081, the second hypothesis is selected because it has the highest posterior probability and its prior probability crosses the threshold (0.8). See Figure 12 for further details.

D. Datasets D2 and D3

The developed approach performs similarly well also in the datasets **D2** and **D3**. We therefore only summarize the results for these two datasets in this section.

Figure 13 describes the results for dataset **D2**: Figure 13a shows the candidate correspondences (the sets \mathcal{F}^r) between different pairs of robots. In this particular case, there was only one hypothesis (excluding the null-hypothesis) that indicated all these

^bOnly outliers are shown for the null-hypothesis, as by definition it does not contain any inliers.

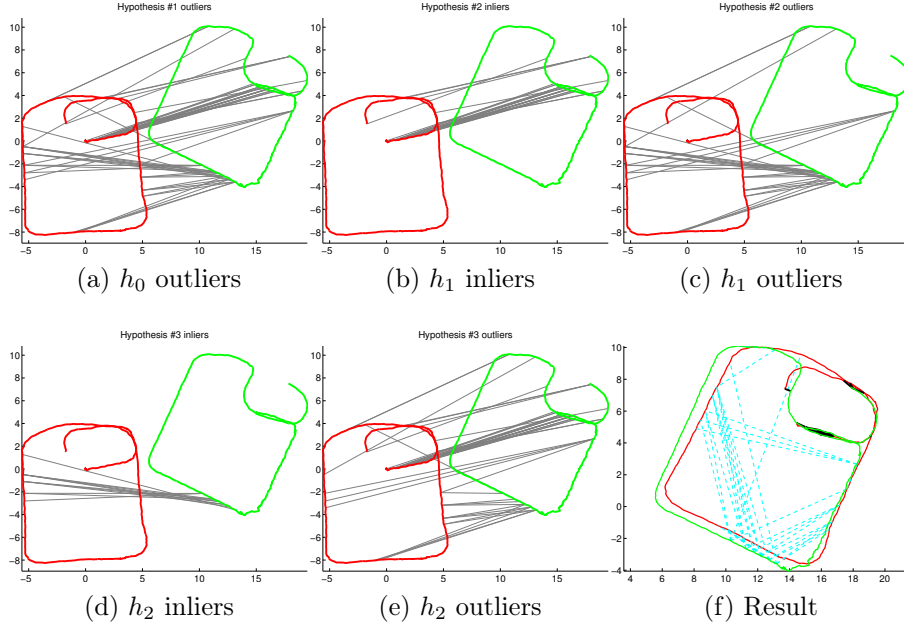



Figure 12: Details of prior probability calculation at time index 6081. (a) Outlier correspondences of the null-hypothesis h_0 (no inliers by definition). (b)-(c) Inlier and outlier correspondences of the chosen hypothesis (h_1). (d)-(e) Inlier and outlier correspondences of another hypothesis (h_2). (f) EM optimization result using the correctly chosen hypothesis (ground truth is given in Figure 8a). Inliers are denoted by black color.

h	m	m_{in}	m_{out}	n	$f(n)$	prior
h_0	58	0	58	116	$4e-6$	0.007
h_1		26	32	90	$5e-5$	0.89
h_2		11	47	105	$3e-5$	0.06
h_3		9	49	107	$2e-5$	0.04


Table 3: Number of involved unique places for the hypotheses shown in Figure 12, and the corresponding *unnormalized* prior given by $f(n)$ from Eq. (29). The normalized prior probabilities are shown in the last column as well as in Figure 11.

correspondences are inliers. The prior of this hypothesis and the null-hypothesis competed with each other, and after sufficient number of observations the prior of the former reached the pre-defined threshold (0.8) and therefore this hypothesis was chosen. Similarly to the basic-example from Figure 7, for a given m (number of correspondences) at each time step, the number of unique places for the chosen and null-space hypotheses are, respectively, m and $2m$.


Figure 13b shows the result of the EM optimization (Eq. (11)), for each of the three robots, right after choosing the hypothesis as described above. As seen, the algorithm correctly identified the robots start from the same location; estimation errors of initial relative poses between the robots are provided in Table 2. It now becomes possible for each robot to estimate its own trajectory, as well as trajectories of all other robots. These estimates are shown in Figure 13c - the trajectories are the same for all the three robots.



(a) blue and green robots (left); blue and red robots (middle); red and green robots (right).



(b) (left) red robot. Zoom-in on inliers is shown; (middle) green robot; (right) blue robot.



(c)

Figure 13: Dataset **D2**: (a) Robot trajectories with initial pose set to arbitrary values and multi-robot candidate correspondences. (b) Identified inlier correspondences and robot trajectories expressed in the same reference frame once the latter is established. See text for details.

Algorithm performance in the third dataset, **D3**, is summarized in Figure 14 and in

Table 2. In contrast to the first two datasets, **D1** and **D2**, in this dataset the robot start operating from different positions (see Figure 8c) and operate in similar (corridor like) environments. This is a challenging problem, since perceptual aliasing introduces consistent outliers, as discussed in the context of the dataset **D1**. Nevertheless, the proposed approach is capable of identifying the correct initial relative transformation between the robots: Figure 14a shows the set of candidate correspondences \mathcal{F}^r and the equivalent distribution of the initial relative poses (the set $\mathcal{T}_r^{r'}(\mathcal{F}^r)$) between the green and the blue robots. The latter exhibits several clusters, and the most dominant four clusters were identified. The probabilities of the corresponding hypotheses including the null-hypothesis are shown in Figure 14b, and the hypothesis h_1 is chosen. The result of the EM multi-robot optimization is shown in the right side of Figure 14b (the trajectory of the red robot is also shown since by that time, a common reference frame between the green and the red robot has been established). Comparing with ground truth trajectories (Figure 8c), one can observe the robot trajectories, expressed in green robot’s reference frame, are aligned well. Similar to dataset **D2**, since no loop closures were used, the trajectories drift over time (Figure 14c).

Observe, the drift in estimated trajectories with respect to ground truth in both datasets **D2** and **D3** (Figures 13c and 14c). These errors happen for two reasons: a) no single-robot loop closure constraints were used in these datasets (as opposed to dataset **D1**); and b) due to increased estimation errors, all multi-robot candidate matches are identified as outliers. As further discussed in Section VII, the latter is a limitation of the proposed approach and of state-of-the-art single-robot robust inference approaches.


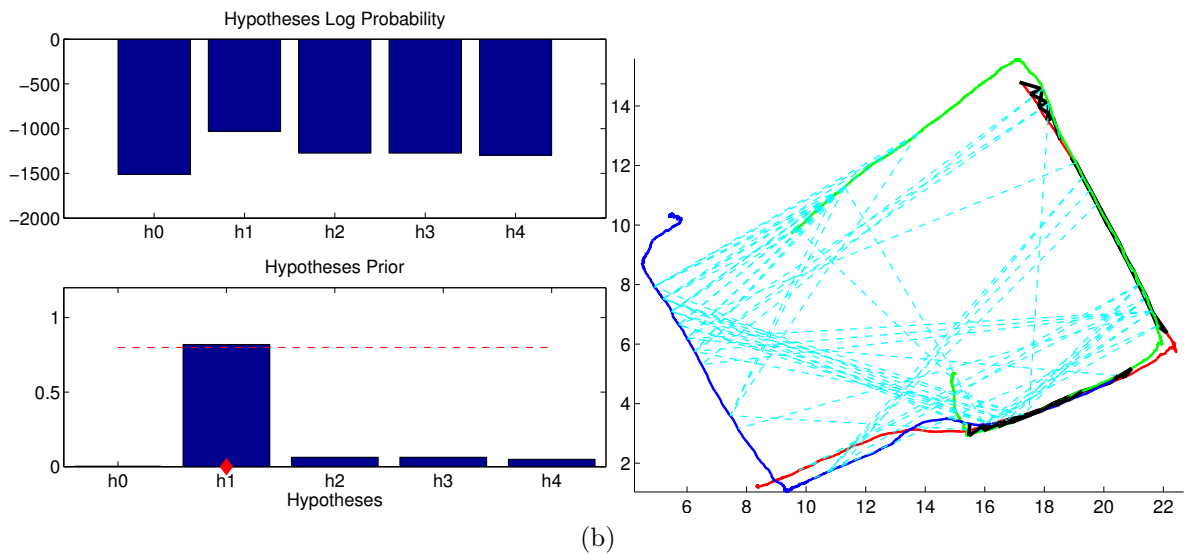

VII. Limitations

The main limitation is related to challenges in identifying multi-robot inlier constraints in cases where estimation errors in robot trajectories are significant. One can observe this is the situation in datasets **D2** and **D3**: while the common reference frames between the robots have been estimated *correctly* (Sections IV-B and V), after the robots proceed in their motion, the current implementation was *unable* to identify the appropriate candidate multi-robot constraints as inliers, and therefore estimation errors are not reduced (see, e.g. Figure 13c). In contrast, this is not the case for dataset **D1**, since in that dataset each of the robots was assumed to be capable of identifying its own loop-closures, resulting in reduced trajectory errors in each of the robots.

Current research focuses on addressing this challenge. In particular, a direction that is being considered is to infer multi-robot data association using EM formulation from Section IV-B, while extending the formulation to account for the uncertainty in robot trajectories.

VIII. Conclusions

We presented an approach for distributed and incremental inference by a group of collaborating robots that are initially unaware of each other’s position and without assuming multi-robot data association to be given. We formulate this problem within an EM framework that, starting from promising initial guesses, converges to a number of locally-optimal hypotheses regarding data association and reference frame transformations. Choosing the correct hypothesis is challenging in the incremental setting because of measurement aliasing and since there may be insufficient data to make this decision



(c) Estimated trajectories of red, green and blue robots by each of the robots.

Figure 14: Dataset **D3**: (a) candidate correspondences and the equivalent distribution of the initial relative poses between green and blue robots; (b) hypotheses (posterior and prior) probabilities and the result of EM multi-robot optimization (11) after hypothesis h_1 is chosen; (c) Robot estimated trajectories on top of ground truth. The figure shows results for all the three robots: for each robot, its own and other robots' estimates are shown. See text for details.

reliably. We address these challenges by developing a model-based selection approach for choosing the most probable hypothesis, while using Chinese Restaurant Process to model hypothesis prior probabilities. Our method is evaluated on real-world experiments, exhibiting accurate estimation of common reference frames and correctly determining multi-robot data association assuming estimation errors in robot local trajectories (poses) are not significant. Future work will focus on relaxing this assumption.

Appendix

This appendix details EM equations for recovering a local maxima of $T_r^{r'}$. Recall the MAP estimate is given by Eq. (12):

$$\hat{T}_r^{r'} = \arg \max_{T_r^{r'}} \sum_{\mathcal{J}^r} p \left(T_r^{r'}, \mathcal{J}^r | \hat{\mathcal{X}}^r, \mathcal{Z}^r \right).$$

Let us simplify notations and write the above equation as

$$\hat{T} = \arg \max_T \sum_{\mathcal{J}} p \left(T, \mathcal{J} | \hat{\mathcal{X}}, \mathcal{Z} \right). \quad (33)$$

The EM algorithm:

- E step calculates a lower bound on $p \left(T | \hat{\mathcal{X}}, \mathcal{Z} \right)$ via

$$Q \left(T | \hat{T}^{(i)} \right) \doteq \mathbb{E}_{\mathcal{J} | \hat{T}^{(i)}, \hat{\mathcal{X}}, \mathcal{Z}} \left[\log p \left(T, \mathcal{J} | \hat{\mathcal{X}}, \mathcal{Z} \right) \right] = \int p \left(\mathcal{J} | \hat{T}^{(i)}, \hat{\mathcal{X}}, \mathcal{Z} \right) \log p \left(T, \mathcal{J} | \hat{\mathcal{X}}, \mathcal{Z} \right) d\mathcal{J}. \quad (34)$$

- M step maximizes this lower bound:

$$\hat{T}^{(i+1)} = \arg \max_T Q \left(T | \hat{T}^{(i)} \right).$$

We now focus on the E step and write the underlying explicit equations. The joint pdf $p \left(T, \mathcal{J} | \hat{\mathcal{X}}, \mathcal{Z} \right)$ is

$$p \left(T, \mathcal{J} | \hat{\mathcal{X}}, \mathcal{Z} \right) \propto p \left(\mathcal{Z} | \hat{\mathcal{X}}, \mathcal{J}, T \right) p \left(\mathcal{J}, T | \hat{\mathcal{X}} \right).$$

Assuming \mathcal{J} and T are independent given robot poses, $p \left(\mathcal{J}, T | \hat{\mathcal{X}} \right) = p \left(T | \hat{\mathcal{X}} \right) p \left(\mathcal{J} | \hat{\mathcal{X}} \right)$, the above can be expressed in terms of individual candidate correspondences (see Eq. (9)):

$$p \left(\mathcal{J}, T | \hat{\mathcal{X}} \right) \propto p \left(T | \hat{\mathcal{X}}^r \right) \prod_{(r, r', k, l)} p \left(j_{k, l}^{r', r} | x_k^{r'}, x_l^r \right) p \left(u_{k, l}^{r', r} | x_k^{r'}, x_l^r, j_{k, l}^{r', r}, T \right).$$

Also here, we resort to simplified notations and re-write the above as

$$p \left(\mathcal{J}, T | \hat{\mathcal{X}} \right) \propto p \left(T | \hat{\mathcal{X}}^r \right) \prod_s p \left(j_s | x_{s1}, x_{s2} \right) p \left(z_s | x_{s1}, x_{s2}, j_s, T \right). \quad (35)$$

In particular, $\mathcal{J} \doteq \{j_1, \dots, j_n\}$.

Using Eq. (35), the lower bound $Q\left(T|\hat{T}^{(i)}\right)$ from Eq. (34) becomes

$$\begin{aligned} Q\left(T|\hat{T}^{(i)}\right) &= \int p\left(j_1, \dots, j_n|\hat{T}^{(i)}, \hat{\mathcal{X}}, \mathcal{Z}\right) \log p\left(T, \mathcal{J}|\hat{\mathcal{X}}, \mathcal{Z}\right) d\mathcal{J} \\ &= p\left(T|\hat{\mathcal{X}}^r\right) + \\ &+ \sum_s \int p\left(j_1, \dots, j_n|\hat{T}^{(i)}, \hat{\mathcal{X}}, \mathcal{Z}\right) \log [p\left(j_s|x_{s1}, x_{s2}\right) p\left(z_s|x_{s1}, x_{s2}, j_s, T\right)] d\mathcal{J}. \end{aligned} \quad (36)$$

At this point, we assume the following factorization of $p\left(j_1, \dots, j_n|\hat{T}^{(i)}, \hat{\mathcal{X}}, \mathcal{Z}\right)$ for each s :

$$\forall s \quad p\left(j_1, \dots, j_n|\hat{T}^{(i)}, \hat{\mathcal{X}}, \mathcal{Z}\right) = p\left(j_s|\hat{T}^{(i)}, \hat{\mathcal{X}}, \mathcal{Z}\right) p\left(j_1, \dots, j_{s-1}, j_{s+1}, \dots, j_n|\hat{T}^{(i)}, \hat{\mathcal{X}}, \mathcal{Z}\right).$$

This is equivalent to assuming the variables j_1, \dots, j_n are statistically independent conditioned on $\hat{T}^{(i)}, \hat{\mathcal{X}}, \mathcal{Z}$.

Plugging this into Eq. (36) results in

$$\begin{aligned} Q\left(T|\hat{T}^{(i)}\right) &= p\left(T|\hat{\mathcal{X}}^r\right) \\ &+ \sum_s \int p\left(j_s|\hat{T}^{(i)}, \hat{\mathcal{X}}, \mathcal{Z}\right) \log [p\left(j_s|x_{s1}, x_{s2}\right) p\left(z_s|x_{s1}, x_{s2}, j_s, T\right)] d\mathcal{J}. \end{aligned} \quad (37)$$

with

$$p\left(j_s|\hat{T}^{(i)}, \hat{\mathcal{X}}, \mathcal{Z}\right) = \frac{p\left(z_s|\hat{T}^{(i)}, x_{s1}, x_{s2}, j_s\right) p\left(j_s\right)}{p\left(z_s|\hat{T}^{(i)}, x_{s1}, x_{s2}\right)}. \quad (38)$$

Since j_s is binary, $j_s \in \{inlier, outlier\}$, the second term in Eq. (37) can be further expanded as:

$$\begin{aligned} &\sum_s p\left(j_s = inlier|\hat{T}^{(i)}, \hat{\mathcal{X}}, \mathcal{Z}\right) \log [p\left(j_s|x_{s1}, x_{s2}\right) p\left(z_s|x_{s1}, x_{s2}, T, j_s = inlier\right)] + \\ &p\left(j_s = outlier|\hat{T}^{(i)}, \hat{\mathcal{X}}, \mathcal{Z}\right) \log [p\left(j_s|x_{s1}, x_{s2}\right) p\left(z_s|x_{s1}, x_{s2}, T, j_s = outlier\right)]. \end{aligned}$$

The denominator in Eq. (38) is not actually evaluated - instead, only the nominator is calculated for both inlier and outlier models, and then normalized so that

$$p\left(j_s = inlier|\hat{T}^{(i)}, \hat{\mathcal{X}}, \mathcal{Z}\right) + p\left(j_s = outlier|\hat{T}^{(i)}, \hat{\mathcal{X}}, \mathcal{Z}\right) = 1.$$

References

- [1] L. Andersson and J. Nygard. C-SAM : Multi-robot SLAM using square root information smoothing. In *IEEE Intl. Conf. on Robotics and Automation (ICRA)*, 2008.
- [2] A. Bahr, M.R. Walter, and J.J. Leonard. Consistent cooperative localization. In *IEEE Intl. Conf. on Robotics and Automation (ICRA)*, pages 3415–3422, May 2009.
- [3] T. Bailey, M. Bryson, H. Mu, J. Vial, L. McCalman, and H. Durrant-Whyte. Decentralised cooperative localisation for heterogeneous teams of mobile robots. In *IEEE Intl. Conf. on Robotics and Automation (ICRA)*, 2011.

- [4] D. Blei, T. Griffiths, M. Jordan, and J. Tenenbaum. Hierarchical topic models and the nested chinese restaurant process. In *Advances in Neural Information Processing Systems (NIPS)*, volume 16, 2003.
- [5] K. P. Burnham and D. R. Anderson. *Model selection and multi-model inference: a practical information-theoretic approach*. Springer, 2002.
- [6] L. Carlone, M. Kaouk Ng, J. Du, B. Bona, and M. Indri. Rao-Blackwellized particle filters multi robot SLAM with unknown initial correspondences and limited communication. In *IEEE Intl. Conf. on Robotics and Automation (ICRA)*, pages 243–249, 2010.
- [7] M. Cummins and P. Newman. FAB-MAP: Probabilistic Localization and Mapping in the Space of Appearance. *Intl. J. of Robotics Research*, 27(6):647–665, June 2008.
- [8] A. Cunningham, V. Indelman, and F. Dellaert. DDF-SAM 2.0: Consistent distributed smoothing and mapping. In *IEEE Intl. Conf. on Robotics and Automation (ICRA)*, Karlsruhe, Germany, May 2013.
- [9] A. Cunningham, K. Wurm, W. Burgard, and F. Dellaert. Fully distributed scalable smoothing and mapping with robust multi-robot data association. In *IEEE Intl. Conf. on Robotics and Automation (ICRA)*, St. Paul, MN, 2012.
- [10] Frank Dellaert. Factor graphs and GTSAM: A hands-on introduction. Technical Report GT-RIM-CP&R-2012-002, Georgia Institute of Technology, 2012.
- [11] J. Fenwick, P. Newman, and J. Leonard. Cooperative concurrent mapping and localization. In *IEEE Intl. Conf. on Robotics and Automation (ICRA)*, volume 2, pages 1810–1817, 2002.
- [12] M. Fischler and R. Bolles. Random sample consensus: a paradigm for model fitting with application to image analysis and automated cartography. *Commun. ACM*, 24:381–395, 1981.
- [13] A. Franchi, G. Oriolo, and P. Stegagno. Mutual localization in multi-robot systems using anonymous relative measurements. *Intl. J. of Robotics Research*, 32:81–95, 2013.
- [14] R. I. Hartley and A. Zisserman. *Multiple View Geometry in Computer Vision*. Cambridge University Press, second edition, 2004.
- [15] A. Howard. Multi-robot simultaneous localization and mapping using particle filters. *Intl. J. of Robotics Research*, 25(12):1243–1256, 2006.
- [16] A. Howard, G. S. Sukhatme, and M. J. Matarić. Multi-robot mapping using manifold representations. *Proceedings of the IEEE - Special Issue on Multi-robot Systems*, 94(9):1360 – 1369, Jul 2006.
- [17] V. Indelman, P. Gurfil, E. Rivlin, and H. Rotstein. Distributed vision-aided cooperative localization and navigation based on three-view geometry. *Robotics and Autonomous Systems*, 60(6):822–840, June 2012.

- [18] V. Indelman, P. Gurfil, E. Rivlin, and H. Rotstein. Graph-based distributed cooperative navigation for a general multi-robot measurement model. *Intl. J. of Robotics Research*, 31(9), August 2012.
- [19] V. Indelman, E. Nelson, N. Michael, and F. Dellaert. Multi-robot pose graph localization and data association from unknown initial relative poses via expectation maximization. In *IEEE Intl. Conf. on Robotics and Automation (ICRA)*, 2014.
- [20] B. Kim, M. Kaess, L. Fletcher, J. Leonard, A. Bachrach, N. Roy, and S. Teller. Multiple relative pose graphs for robust cooperative mapping. In *IEEE Intl. Conf. on Robotics and Automation (ICRA)*, pages 3185–3192, Anchorage, Alaska, May 2010.
- [21] J. Knuth and P. Barooah. Outlier rejection for pose graph optimization. In *IEEE/RSJ Intl. Conf. on Intelligent Robots and Systems (IROS)*, 2013.
- [22] Y. Latif, C. D. C. Lerma, and J. Neira. Robust loop closing over time. In *Robotics: Science and Systems (RSS)*, 2012.
- [23] G. H. Lee, F. Fraundorfer, and M. Pollefeys. Robust pose-graph loop-closures with expectation-maximization. In *IEEE/RSJ Intl. Conf. on Intelligent Robots and Systems (IROS)*, 2013.
- [24] F. Lu and E. Milios. Globally consistent range scan alignment for environment mapping. *Autonomous Robots*, pages 333–349, Apr 1997.
- [25] L. Merino, J. Wiklund, F. Caballero, A. Moe, JR Martinez De Dios, P-E Forssen, K. Nordberg, and A. Ollero. Vision-based multi-uav position estimation. *Robotics & Automation Magazine*, 13(3):53–62, 2006.
- [26] E. Montijano, R. Aragues, and C. Sagues. Distributed data association in robotic networks with cameras and limited communications. *IEEE Trans. Robotics*, 29(6):1408–1423, 2013.
- [27] T. K. Moon. The expectation-maximization algorithm. *Signal processing magazine, IEEE*, 13(6):47–60, 1996.
- [28] J. Nieto, T. Bailey, and E. Nebot. Recursive scan-matching slam. *Robotics and Autonomous Systems*, 55(1):39–49, 2007.
- [29] A. Ranganathan and F. Dellaert. Online probabilistic topological mapping. *Intl. J. of Robotics Research*, 30(6):755–771, 2011.
- [30] S.I. Roumeliotis and G.A. Bekey. Distributed multi-robot localization. *IEEE Trans. Robot. Automat.*, August 2002.
- [31] N. Sunderhauf and P. Protzel. Towards a robust back-end for pose graph slam. In *IEEE Intl. Conf. on Robotics and Automation (ICRA)*, pages 1254–1261. IEEE, 2012.
- [32] Niko Sünderhauf and Peter Protzel. Switchable constraints vs. max-mixture models vs. rrr—a comparison of three approaches to robust pose graph slam. In *IEEE Intl. Conf. on Robotics and Automation (ICRA)*, 2013.

- [33] Y. W. Teh, M. Jordan, M. Beal, and D. Blei. Sharing clusters among related groups: Hierarchical dirichlet processes. In *Advances in Neural Information Processing Systems (NIPS)*, 2004.
- [34] J. M. Walls and R. M. Eustice. An exact decentralized cooperative navigation algorithm for acoustically networked underwater vehicles with robustness to faulty communication: Theory and experiment. 2013.
- [35] X. Zhou and S.I. Roumeliotis. Multi-robot SLAM with unknown initial correspondence: The robot rendezvous case. In *IEEE/RSJ Intl. Conf. on Intelligent Robots and Systems (IROS)*, pages 1785–1792. IEEE, 2006.