

A Hybrid Approach to the Simultaneous Localization and Mapping (SLAM) Problem

S. Chakravorty, R. Saha

Abstract

A hybrid Bayesian/ frequentist approach is presented for the Simultaneous Localization and Mapping Problem (SLAM). A frequentist approach is proposed for mapping with time varying robotic poses and is generalized to the case when the robotic pose is both time varying and uncertain. The SLAM problem is then solved in two steps: 1) the robot is localized with respect to a sparse set of landmarks in the map using a Bayes filter and a belief on the robot pose is formed, and 2) this belief on the robot pose is used to map the rest of the map using the frequentist estimator. The hybrid methodology is shown to have complexity linear in the map components, is robust to the data association problem and is provably consistent. The hybrid methodology is tested on several large maps in order to validate the theoretical guarantees.

I. INTRODUCTION

The problem of Simultaneous Localization and Mapping (SLAM) using mobile robotic agents is addressed in this paper. Autonomous agents such as military robots would be required to operate in highly uncertain and dynamic environments, such as a battle zone, with minimal or no external help such as GPS or human intervention. Other applications of such mobile autonomous systems range from disaster relief in areas such as the 9/11 WTC site, to undersea and planetary exploration. Hence, the development of true autonomy in such systems can have very wide military, societal and scientific impact. In order to be “truly autonomous”, any such mobile autonomous system needs to be capable of accomplishing two sub-tasks: 1) it needs to model its environment, along with its attendant uncertainties, as well as find its own location within the environment, using only the noisy observations from its on-board sensors, and 2) it needs to use the environment model to intelligently plan its actions within the environment such that it accomplishes its objectives in a robust and timely fashion. Problem 1 above is known as the Simultaneous Localization and Mapping (SLAM) problem, while the addition of planning to the problem of SLAM in Problem 2 results in the so-called Simultaneous Planning, Localization and Mapping (SPLAM) problem. Thus, the SLAM problem is a fundamental problem in enabling “true autonomy”. We propose a hybrid Bayesian/ frequentist approach to the SLAM problem. In the proposed method, the environment is split into a small set of sparse landmarks/ features and the rest of the (dense) environment. The philosophy behind the hybrid method amounts to the following two steps: 1) localize with respect to the landmarks, i.e., form a belief on the pose of the robotic system based on observations of the landmarks using a Bayes filter such as an EKF(**the Bayesian sub-problem**), and then, 2) map the rest of the environment based on the belief on the robotic pose using the frequentist mapping technique (**the frequentist sub-problem**). This formulation has linear complexity in the map components, is robust to the data association problem and provably consistent. In the following, we survey the state of the art in the SLAM literature.

There are two main categories of approaches to SLAM: recursive and trajectory based. In the recursive Kalman filter/ Information filter based approach [1]–[8], the map is appended to the filter as a parameter and the joint pose-map pdf is estimated using the Kalman recursion in either the covariance or the information matrix form. The Kalman filter scales quadratically as the size of the environment since the correlations between all the map elements need to be maintained in order for the map to be consistent [1] owing to the nature of the SLAM problem. The EKF SLAM based approach can only extract sparse maps since it is not robust to the data association problem in dense maps. In the SEIF (sparse extended information filter) based approach, the information form of the Kalman update is used and the sparsity of the information filter is used to obtain nearly constant time SLAM algorithms [7]. However, the SEIF based filter tends to get overconfident and the latest research has concentrated on alleviating this problem by either having exact sparsity by keeping track of the trajectory [8] or enforcing sparsity by solving an approximate problem involving “kidnapping and relocalizing” the robot [9]. In the approach outlined in this paper, due to the frequentist update of the map, the correlations between the various map elements need not be maintained for consistency, which is in contrast to the result in [1] which asserts that maintaining correlations is key to consistency. The difference is due to the frequentist problem formulation which obviates the need to keep track of the correlations between the various map components. In addition, it is also shown that the hybrid formulation is robust to the data association problem in dense maps. In our opinion, the hybrid methodology proposed can most fruitfully be used in conjunction with sophisticated EKF/ SEIF based SLAM algorithms, wherein the SEIF/ EKF solve the Bayesian sub-problem of the hybrid formulation, in order to scale to really large, dense environments such as cities, ocean floors and planetary terrain.

The RBPF based SLAM algorithms, which falls under the broader category of trajectory based SLAM, on the other hand keep track of the whole trajectory of the robot which decorrelates the observations of the various map components. These methods have become very popular over the past few years and efficient techniques have been developed to form sparse landmark based maps as well as dense occupancy grid maps [10]–[16]. Another trajectory based method, called Consistent Pose Estimation (CPE) [5] relies on maintaining a graph on the poses at which various scans of the map were made and then, optimizing the inter-node distances such that the likelihood of the observed data is maximized given the statistics of the observation process [17]–[19]. However, these trajectory based methods keep track of the whole vehicle trajectory and thus, their state space grows unbounded over time. The CPE methods also require multiple passes over the same data to solve the pose optimization problem and thus, are essentially an offline batch processing method. For long term SLAM in large environments, it is still necessary to truncate the data at some finite time in the past [5], [6] due to the constraints on the memory requirements, which may lead to the loss of consistency. In contrast, the method presented here is purely recursive and does not require the belief to be maintained over the entire vehicle pose history in order to decorrelate the map components, while being provably consistent.

It can be seen from the previous discussion that the main problem with the Bayesian formulation of the SLAM problem is the fact that computational complexity and consistency are at cross purposes, and alleviating one of the problems tends to worsen the other. In fact, starting with the important paper [1], the consistency of the SLAM algorithms, both RBPF and EKF/SEIF based, has increasingly come under scrutiny in recent years [9], [20]–[22]. It was found that the RBPF based filters tend to lose consistency as time increases because of their inability to forget the past [21], [23], and the EKF based methods need stabilizing noise for consistency [20]. Also, neglecting the weak off-diagonal elements in SEIF based methods can lead to inconsistent results [9]. The method proposed in this paper first localizes with respect to a sparse set of landmarks in the map using a Bayes filter such as an EKF/SEIF to obtain a belief on the pose of the robot, and then, maps the rest of the environment, based on the belief, using a frequentist approach. The frequentist part of the algorithm has complexity linear in the map and is provably consistent given that the Bayesian part of the problem is consistent. The complexity of the Bayesian part of the formulation can be kept under control owing to the sparseness of the set of landmarks/features and using suitable sophisticated feature-based SLAM methods. Further, the frequentist part of the formulation is robust to the data association problem, while the Bayesian part can be expected to be robust owing to the sparseness of the features/landmarks and thus, the hybrid formulation is robust to the data association problem.

A valid question at this point is: why use the hybrid method? Why not localize using the EKF/SEIF, and map the rest of the environment using a Bayesian method such as the occupancy grid OG method [24], as proposed in the DenseSLAM approach [25], [26]. The answer to this question is that it may be impossible to maintain consistency in the Bayesian approach without maintaining correlation between map components, even under the “first localize - then map” philosophy adopted here! A simple counterexample is provided at the end of Section 2 to prove this. The other methods close to the work presented here, in philosophy, is [27], and other Expectation-Maximization based methods [28], [29]. In these papers, a frequentist approach, the Baum-Welch algorithm [30], is used to find the ML estimate of a sparse set of landmarks and the vehicle trajectory, and then, this ML estimate is used to construct an OG map of the environment. However, the method used is an offline batch processing algorithm. The frequentist method proposed here is a recursive stochastic approximation algorithm and as such, different from the Baum Welch algorithm. Also, the Baum-Welch algorithm does not provide an estimate of the uncertainty in the map estimate, whereas the method proposed here provides a complete probabilistic description of the map, much like the OG maps, which can then be used for motion planning, exploration etc. The use of frequentist estimators based on recursive ML (RML) or recursive least squares (RLS) is standard practice in the Hidden Markov Model literature [31], and an application of this methodology in the SLAM context is made in the reference [23]. These methods usually need to evaluate the filter derivative which is an $O(N^2)$ operation, where N is the number of particles used to represent the pdf of the state. This is usually impractical and only through recent advances in the particle filtering community [22], the above operation can now be done with $O(N \log N)$ complexity. In contrast, the method presented here does not require the filter derivative, and if the Bayesian part of the hybrid formulation is implemented using a particle filter, the complexity of the frequentist algorithms is $O(N)$, where N is the number of particles used to represent the robot pose pdf. This is accomplished by using (i) a probabilistic description of the map as the parameter in the methodology, instead of the deterministic description common in general joint state-parameter estimation algorithms [22], [23], [31], (ii) utilizing the “first localize -then map” philosophy, and (iii) by exploiting the structure of the resulting problem to intuitively define a frequentist estimator that is provably consistent.

The rest of the document is organized as follows. In section 2, we present the hybrid approach to the SLAM problem. In section 3, we prove the consistency of the hybrid approach. In section 4, we present several experiments wherein large environments with multiple cycles are mapped using the hybrid methodology. Preliminary versions of this paper have been published in [32]–[34]. In references [32], [33], the pure mapping problem was addressed, i.e., there was no uncertainty in the pose of the robot while in [34], preliminary versions of the problems in this paper were addressed. In particular, this paper considers the problem of dense contiguous maps and addresses the issue of data association in such maps, in contrast to [34].

II. THE HYBRID BAYESIAN/ FREQUENTIST METHODOLOGY

A. Frequentist Mapping

Consider a single autonomous agent and let its state be denoted by the variable s (also sometimes called the robotic pose), and let the state of the environment be denoted by the variable $Q = \{q_1, \dots, q_M\}$, where q_k are components of the environment (for instance, these would be the individual grid cells in a grid cell decomposition of the environment). The state and the environment are assumed to be discrete-valued random variables. The environment is assumed to be stationary and uncorrelated, i.e.,

$$p^*(Q) = \prod_{i=1}^M p^*(q_i), \quad (1)$$

It can be anticipated that an overwhelmingly large part of most environments can be modeled in this fashion. In fact, any deterministic environment trivially satisfies the above assumptions. The probability of observing the i^{th} environmental component in the state \hat{q}_i , where \hat{q}_i can take one of D values, and given that it is observed from the pose s , is given by:

$$p(\hat{q}_i/s) = \sum_{q_1, \dots, q_N} p(\hat{q}_i/q_1, \dots, q_N, s) p^*(q_1) \dots p^*(q_N). \quad (2)$$

The above equation can be rewritten as:

$$p(\hat{q}_i/s) = \sum_{q_i} p^*(\hat{q}_i/q_i, s) p^*(q_i), \quad (3)$$

$$p^*(\hat{q}_i/q_i, s) = \sum_{q_1, \dots, q_{i-1}, q_{i+1}, \dots, q_N} p(\hat{q}_i/q_1, \dots, q_N, s) p^*(q_1) \dots p^*(q_{i-1}) p^*(q_{i+1}) \dots p^*(q_N). \quad (4)$$

The above may be compactly written in matrix form as the equation

$$\hat{P}_i(s) = A_i^*(s) P_i^*, \quad (5)$$

where the vector $\hat{P}_i(s)$ stacks the observation probabilities $p(\hat{q}_i/s)$, and the matrix $A_i^*(s)$ is the true observation model of the i^{th} component when observed from pose s . **The above equation is the fundamental equation for the frequentist approach and provides an avenue for estimating the true environmental probabilities P_i^* .** Suppose we make repeated observations of the i^{th} component from pose s . We could count the number of times that we observe the i^{th} component in its various states, and form a consistent estimate of the observation probability vector $\hat{P}_i(s)$ by averaging, i.e.,

$$\hat{P}_i(s) = E_z[1(\hat{q}_i/s, z)] \equiv E_z[c_i(s, z)] = \lim_N \frac{1}{N} \sum_{t=1}^N 1(\hat{q}_{i,t}/s, z_t). \quad (6)$$

In the above, given an observation z , the observation vector $c_i(s, z) = [1(\hat{q}_i/s, z)]$ ($1(\cdot)$ denotes the indicator function) enters a one into the \hat{q}_i component, and zero in every other component at time t (for instance, in the occupancy grid representation it will enter a 1 into the ‘‘occupied’’ entry if the grid cell is observed to be occupied, or a 1 into the ‘‘empty’’ entry otherwise). The above equation is correct due to the Law of Large Numbers. Then, using the knowledge of $A_i^*(s)$, we can obtain the true environmental probabilities P_i^* as

$$P_i^* = A_i^*(s)^{-1} \hat{P}_i(s). \quad (7)$$

Next, we may relax the assumption that the observations are made from the pose s and have that the observations are made from the time varying poses $\{s_t\}$, with true observation models $A_i^*(s_t)$. Again, if we keep track of the relative frequencies of observations of the i^{th} component in its various different states, then the estimate of the true probabilities P_i^* can be recovered asymptotically using a time averaged observation model as follows:

$$P_i^* = \bar{A}_i^{-1} \hat{P}_i, \quad (8)$$

$$\hat{P}_i = \frac{1}{N} \sum_{t=1}^N c_i(s_t, z_t), \quad (9)$$

$$\bar{A}_i = \frac{1}{N} \sum_{t=1}^N A_i^*(s_t). \quad (10)$$

In order to derive the above expression, note that if we interpret the frequency of seeing the i^{th} map component in its \hat{q}_i level during the course of the mapping experiment as a probability, and if we also interpret the frequency of the robot being in a

state s as a probability, then it follows using the simple rules of conditional probability that:

$$p(\hat{q}_i) = \sum_{q_i, s} p^*(\hat{q}_i/q_i, s) p^*(q_i) p(s) = \sum_{q_i} [\sum_s p^*(\hat{q}_i/q_i, s) p(s)] p^*(q_i). \quad (11)$$

Provided that the state s_t converges to some stationary distribution, the left hand side $p(\hat{q}_i)$ in the above equation is given by Eq. 9, and the matrix $[\sum_s p^*(\hat{q}_i/q_i, s) p(s)]$ is given by Eq. 10, and hence, the estimation equations for the time varying pose case follow. The true environmental probabilities can then be recovered recursively using the following estimator if $A^*(s_t)$ is positive definite (which is true under mild conditions).

Estimator E1:

$$P_{i,t} = \Pi_{\mathcal{P}} \{P_{i,t-1} + \gamma_t (c_i(s_t, z_t) - A_i^*(s_t) P_{i,t-1})\}, \quad (12)$$

where \mathcal{P} represents the space of probability vectors in \mathbb{R}^D , and $\Pi_{\mathcal{P}}(\cdot)$ denotes a projection onto this compact set. The sequence $\{\gamma_t\}$ is usually of the form $at^{-\alpha}$, $\alpha < 1$, where a and α are design parameters and greatly influence the convergence characteristics of the algorithm. However, there still remains the problem of using the “true” observation models $A_i^*(s)$ in order to form the estimates. This is highly unreasonable since it depends on the true map probabilities that we are trying to estimate. However, note that as the algorithm progresses, we have an estimate $P_i(t)$ of the map probabilities for the different components. This estimate can then be used in Eqs. 4 -5 to form the observation models $A_i(s)$ as an approximation of the true observation models. These models can be inferred from the model of the particular type of sensor used for sensing the environment [24].

At this point, a few more details regarding the projection operator $\Pi_{\mathcal{P}}$ is provided. Note that $p(q_i = 1) + \dots + p(q_i = D) = 1$ and all of the terms are positive since they are probabilities. Then, we may eliminate one of the probabilities by replacing $p(q_i = D)$ by $1 - p(q_i = 1) + \dots + p(q_i = D - 1)$ and having the constraint that every term is positive and their sum is less than one. This is easier understood in the 2-d case, i.e., when any map component can take one of only two values 0/1. In this case, $P(q_i = 0) = p_{i2} = 1 - P(q_i = 1)$. Denoting the probability $p(q_i = 1)$ by P_i , the constraint becomes that $0 \leq P_i \leq 1$. Thus, in that case, the map probabilities can be denoted by the vector $P = [P_1, \dots, P_M]$, comprised of the occupancy probabilities of each of the map components. Then, the above algorithm reduces to the following:

$$P_{i,t+1} = \Pi_{\mathcal{P}} [P_{i,t} + \gamma_t \{c_i(s_t, z_t) - p(\hat{q}_i = O/q_i = O, s_t) P_{i,t} - p(\hat{q}_i = 0/q_i = E, s_t) (1 - P_{i,t})\}], \quad (13)$$

where $c_i(s_t, z_t) = 1(\hat{q}_{i,t} = 0, s_t, z_t)$, i.e., whether the grid is occupied or not, and $\Pi_{\mathcal{P}}[\cdot]$ represents projection onto the interval $[0, 1]$. Many other possible projection mechanisms are possible but this is the method we use throughout this paper.

B. Frequentist Mapping with Uncertain Robotic Poses

In this section, we relax the assumption that the pose of the robot is known perfectly. Instead, we assume that we are given a belief, i.e., a probability distribution, on the possible poses of the robot. Given the belief on the pose of the robot, $b_t(s)$ at time t , and a reading z_t of the environment, the frequentist mapping method is now used to map the (dense) environment Q . However, it is immediately apparent that there is an inherent “data association” problem associated with the mapping problem in this scenario. The observation, \hat{q}_i , of an environmental component q_i is no longer certain, since it varies with the pose of the robot. Consider the simple situation illustrated in Fig. 1. The map component q_1 , given reading z_2 , is empty or occupied depending on whether the robot is at pose s_1 or s_2 respectively. Thus, given the uncertainty in the pose of the robot $b(s)$ and the reading of the environment z , the observation of the i^{th} component of the environment \hat{q}_i is given by the probability vector (derived using the rules of conditional probability, and Bayes rule)

$$c_i^*(b(s), z) \equiv [p(\hat{q}_i/b, z)] = \sum_s [1(\hat{q}_i/s, z)] \frac{p^*(z/s) b(s)}{p^*(z/b)}, \quad (14)$$

$$p^*(z/s) = \sum_{q_1, \dots, q_N} p(z/s, q_1, \dots, q_N) p^*(q_1) \dots p^*(q_N), \quad (15)$$

where $p^*(z/b) = \sum_s p^*(z/s) b(s)$ is the factor used to normalize $c_i(\cdot)$ and $p^*(z/s)$ is the true likelihood of the observation z given that it is made from pose s . In order to derive the above expression, note that using the theorem of total probability, $p(\hat{q}_i/b, z) = \sum_s p(\hat{q}_i/s, b, z) p(s/z, b)$. We can expand the term $p(s/z, b)$ using Bayes rule which gives us $p(s/z, b) = \frac{p^*(z/s) b(s)}{p^*(z/b)}$, and using the fact that $p(z/s, b) = p(z/s)$, Eq. (14) above follows.

As in the perfect pose information case, averaging over all observations z (which can be formed by a time average due to the Law of Large Numbers), allows us to estimate the probability of observing state \hat{q}_i given the belief state $b(s)$, i.e.,

$$p(\hat{q}_i/b) = E_z [c_i^*(b, z)] \approx \frac{1}{N} \sum_{t=1}^N c_i^*(b, z_t). \quad (16)$$

Note that the above probabilistic description of the observation solves the “data association” problem: we are no longer certain




Fig. 1. The problem of data association

if the observed value of the i^{th} map component is in its k^{th} level, instead we associate a probability with this observation. The probability of observing the map component q_i at level \hat{q}_i , given the belief on the pose $b(s)$ is also given by

$$p(\hat{q}_i/b) = \sum_s \left[\sum_{q_1, \dots, q_N} p(\hat{q}_i/q_1, \dots, q_N, s) p^*(q_1) \cdots p^*(q_N) \right] b(s), \quad (17)$$

which can be written in compact matrix form as follows:

$$\hat{P}_i(b) = A_i^*(b) P_i^*, \quad (18)$$

where $\hat{P}_i(b) = [p(\hat{q}_i/b)]$, and

$$A_i^*(b) = \sum_s A_i^*(s) b(s). \quad (19)$$

Note here that this equation is exactly analogous to the frequentist mapping equation 5, wherein the exact pose knowledge s has been replaced by the belief on the pose of the robot $b(s)$. The observation model $A_i^*(s)$ is replaced by the averaged observation model with the averaging being done with respect to the belief on the pose of the robot. Thus, similar to the case with perfect pose information, if we were to remain in the belief state $b(s)$ and make repeated observations of the i^{th} component of the environment, we would be able to recover the left hand side of the above Eq. 18, $\hat{P}_i(b)$, by averaging the (probabilistic) observations of the i^{th} component, $c_i(b, z_t)$ (cf. Eq. 16). Hence, the true environmental probabilities may be recovered asymptotically by inverting Eq. 18. Generalizing the situation to the case when we have a time-varying belief on the pose of the robot, $b_t(s)$, the true environmental probabilities can be estimated recursively using the following analog of frequentist estimator E1 .

Estimator E2

$$P_{i,t} = \Pi_{\mathcal{P}} \{ P_{i,t-1} + \gamma_t (c_i^*(b_t, z_t) - A_i^*(b_t) P_{i,t-1}) \}, \quad (20)$$

As in the pure mapping case, the estimator is actually run by using the current estimate of the true observation models/ observation likelihood. In other words, the above algorithm is run using $c_i(b_t, z_t, P_t)$ and $A_i(b_t, P_t)$, where the current estimate of the map probabilities P_t is used, instead of the true map probabilities P^* , in Eq. (14) to form $c_i(b_t, z_t, P_t)$, and in Eqs. (17)-(18) to form $A_i(b_t, P_t)$.

C. Hybrid Bayesian/ Frequentist SLAM

At this point, we formulate a hybrid methodology to generalize the frequentist mapping methodology to the simultaneous localization and mapping (SLAM) case. It is assumed that the system localizes itself with respect to a sparse, well separated set of features/ landmarks $\Theta = \{\theta_1, \dots, \theta_K\}$. Then, the belief (or probability distribution) over the pose-features pair is formed recursively using a Bayes filter (such as a Kalman filter in the Linear Gaussian case):

$$b_t(s, \Theta) = p(z_t^\theta / \Theta, s) \sum_{s'} p(s/s', u_{t-1}) b_{t-1}(s', \Theta), \quad (21)$$

where z_t^θ represents the noise corrupted observation of the landmarks Θ at time t , and u_{t-1} denotes the control acting on the system at time $t - 1$. The identification and recognition of these features and landmarks in an autonomous fashion is a

challenging problem in itself, but can be solved using suitable feature-based SLAM algorithms [4], [7], [8]. Given the joint distribution of the pose-landmark pair, the belief on the pose of the vehicle is formed by marginalizing the dependence on the landmarks, and is output to the frequentist part of the mapping algorithm. The frequentist part of the method, i.e., Estimator E2, is now used to map the rest of the (dense) environment using the belief output from the Bayesian part of the methodology. Thus, the hybrid methodology can be represented as the following algorithm:

Hybrid SLAM

Input $b_0(s)$, initial map occupancy probabilities $P_i(0)$, and reading of environment z_1 , $t = 0$.

Do till convergence of map probabilities

Bayesian: Extract the readings of the landmarks, z_t^θ , from the raw sensor readings z_t , and form the belief on the state of the robot using Eq. (21) and marginalizing over the landmarks.

Frequentist: Take the rest of the data z_t^Q and update the occupancy probabilities of those components of the (dense) map Q that are observed given z_t^Q (cf. Eq. 14), using the frequentist estimator E2 (cf Eq. 20).

end

Given a reading of the map z_t , and given that the Bayesian part of the hybrid problem formulation is solved using a particle filter, the computational complexity of the frequentist estimator in updating any map component is $O(N)$ where N is the number of particles used to represent the belief state. The Bayesian part of the formulation inherits the computational complexity of whatever method is used to solve that part of the problem. Also, note that each map component is updated completely independent of the others and hence, the method has complexity linear in the map components, i.e., $O(M)$, where M is the number of map components. Contrast this with the $O(M^2)$ complexity of the Kalman filter based approach. In order to make this clear, suppose that there are $M + N$ total components in a map. At the basic algorithmic level, in the Kalman filter based approach, the computational complexity is $O(N + M)^2$. In the Hybrid formulation, suppose that N is the number of features that is used to localize the robot and M is the rest of the map. Then, at the basic algorithm level, the computational complexity of the hybrid approach is $O(N^2) + O(M)$. Thus, if $N \ll M$, then the hybrid formulation possesses orders of magnitude better computational benefits compared to Bayesian methods such as the EKF. Moreover, due to the sparseness of the landmark/ features, the data association problem for the Bayesian sub-problem is significantly simpler. In conjunction with the robustness of the frequentist estimator to the data association problem, this leads to significantly improved robustness of the hybrid formulation to the data association problem.

Thus, given that the belief on the pose of the robot is consistent, the hybrid formulation presented above has the following properties:

1) Robust to the data association problem, 2) It is provably consistent, and 3) has complexity linear in the map components since the correlations between the map components need not be maintained for consistency. The complexity of the feature-based SLAM part, i.e., the Bayesian part of the method, though not linear in the number of features, can be kept tractable because of the sparsity of the feature set and by using suitable sophisticated feature-based EKF/ SEIF SLAM algorithms.

D. Why Frequentist?

At this point, a very pertinent question is: why not follow the “localize first-then map” philosophy, except instead of using the frequentist mapping methodology for uncertain poses, we use Bayesian techniques such as Elfes’ OG method [24]. Then, we should be able to keep track of the map components independent of each other, just as in the case of the frequentist technique. In the following, through a very simple counterexample, we show that the correlations just cannot be neglected because they are key to consistency of the Bayesian approach whereas the frequentist method remains consistent without having to maintain these correlations, by virtue of the problem formulation.

Consider the situation shown in Fig. 1.1. We assume that the robot can be at one of the two locations s_1 and s_2 with probabilities b_1 and b_2 . We have a perfect range sensor. There are two grids: q_1 and q_2 , q_1 is empty and q_2 is occupied. Thus, there are only two possible observations: z_1 , that made from s_1 and z_2 , that made from s_2 , as shown in the figure. Since, we do not know exactly where the robot is located, we might sometimes think that we have a reading from q_3 , which is actually occluded. Note that we either have no information about q_3 or think its occupied, its never seen empty.

First, we show the consistency of the frequentist approach. Let the current estimates of the map occupancy probabilities for q_1 , q_2 and q_3 be p_1 , p_2 and p_3 respectively (probability that the map components are occupied). Consider the probability of observing q_i occupied given the above estimates, the reading z and the belief state b , $c_i(b, z, P)$ as defined previously (cf. 14):

$$c_i(b, z, P) = \sum_{s=s_1, s_2} 1(\hat{q}_i = O/z, s) \frac{p(z/s)b(s)}{p(z/b)}, i = 1, 2, 3. \quad (22)$$




Fig. 2. Convergence of Map Probabilities

We see that $p(z_1/s_1) = p(z_1/q_1 = E, q_2 = O, s_1)p(q_1 = E)p(q_2 = O) = (1 - p_1)p_2$. Similarly, $p(z_1/s_2) = (1 - p_1)(1 - p_2)p_3$, $p(z_2/s_1) = p_1$, $p(z_2/s_2) = (1 - p_1)p_2$, and hence, $p(z_1/b) = (1 - p_1)p_2b_1 + (1 - p_1)(1 - p_2)p_3b_2$, and $p(z_2/b) = p_1b_1 + (1 - p_1)p_2b_2$. Substituting the above into the expression for $c_i(b, z, P)$ leads us to the following: $c_1(b, z_1, P) = 0$, $c_1(b, z_2, P) = \frac{b_1p_1}{p(z_2/b)}$, $c_2(b, z_1, P) = \frac{(1-p_1)p_2b_1}{p(z_1/b)}$, $c_2(b, z_2, P) = 1$. Since, we have a perfect sensor, the observation matrices are identity, and thus, the update equation simply becomes:

$$p_{i,t} = (1 - \gamma_t)p_{i,t-1} + \gamma_t c_i(b, z_t, P_{t-1}), \quad (23)$$

where $\gamma_t = \frac{1}{at}$. The sanity check that needs to be performed on the above algorithm is that the true map probabilities $p_1 = 0, p_2 = 1$ are fixed points of the above algorithm. We show this for p_2 , the case of p_1 is trivial. Given $p_1 = 0, p_2 = 1$, and z_1 is observed, p_2 will be updated by the algorithm to $c_2(b, z_1, P) = \frac{(1-p_1)p_2b_1}{p(z_1/b)} = \frac{1 \cdot 1 \cdot b_1}{b_1 + 1 \cdot 0 \cdot p_3 \cdot b_2} = 1$. The update of p_2 for z_2 is always 1. Hence, indeed the true map probabilities are fixed points of the algorithm.

Of course, this does not mean that the algorithm converges to the true map probabilities given arbitrary initial conditions on p_1 , p_2 and p_3 . We prove this analytically in section 3. In Fig. 1.2, we show simulations of the progress of the mapping algorithm for various initial conditions of p_1 and p_2 , and for $b_1 = 0.4$ and $b_2 = 0.6$ (the behaviour is similar for other choices of b_1 and b_2). The plots show the trajectories of the algorithm in the $p_1 - p_2$ plane for various different initial conditions. It is seen from the figure that the algorithm does indeed converge to the true map probabilities for all initial conditions.

Now, we show that the Bayesian approach fails unless the correlation between the components q_1, q_2 and the robot pose s is maintained. Consider first the case when we do not consider correlations of any map component to any other map component as well as the robotic pose. In such an approach, the update of the map component would be

$$p(q_i = O/\mathcal{F}^t) = \eta p(z_t/b, q_i = O)p(q_i = O/\mathcal{F}^{t-1}). \quad (24)$$

The sanity check has to be that the true map probabilities is a fixed point of the above update equation. Consider the update of $p(q_2 = O)$ given $p_1 = 0, p_2 = 1$. If observation z_1 turns up, the update of $p(q_2 = O)$ is $p(q_2 = O) = \frac{p(z_1/b, q_2=O)p(q_2=O)}{p(z_1/b, q_2=E)p(q_2=E) + p(z_1/b, q_2=O)p(q_2=O)} = \frac{(1-p_1)p_2b_1}{(1-p_1)p_2b_1 + p_3b_2} = \frac{b_1}{b_1 + p_3b_2} \neq 1$, since p_3 cannot be equal to 0 for its either occluded or mistakenly thought to be occupied. Hence, the true map probabilities are not a fixed point of the above update equation, and hence, the update is not consistent.

Next, we show the case where we maintain the correlations between the map component and the robot pose but not with other map components. In this case, the joint distribution is

$$p(s, q_1, q_2) = b(s)p(q_1/s)p(q_2/s). \quad (25)$$

Thus

$$p(q_i) = \sum_s b(s)p(q_i/s). \quad (26)$$

The update equation for $p(q_i/s)$ is then given by:

$$p(q_i/s, \mathcal{F}^t) = \eta p(z/s, q_i)p(q_i/s, \mathcal{F}^{t-1}). \quad (27)$$

Let $p(q_2/s_1) = p_{21}, p(q_2/s_2) = p_{22}$. Again, we assume the true map probabilities, $p_1 = 0, p_2 = 1$, and show that they are not a fixed point of the above iteration. Given $z_1, p(q_2 = O/s_1) = \frac{p(z_1/s_1, q_2=O)p(q_2=O/s_1)}{p(z_1/s_1, q_2=E)p(q_2=E/s_1) + p(z_1/s_1, q_2=O)p(q_2=O/s_1)} = \frac{(1-p_1)p_{21}}{(1-p_1)p_{21} + 0} = 1$, and similarly given $z_1, p(q_2 = O/s_2) = \eta p(z_1/s_2, q_2 = O)p(q_2 = O/s_2) = 0$. Then, $p(q_2) = p_{21}b_1 + p_{22}b_2 = b_1 \neq 1$. Hence, even in this case, the true map probabilities are not the fixed point of the Bayesian update and hence, not consistent.

Now, suppose that we maintain the correlations between the pose s and the components q_1, q_2 . the recursive update law is as follows:

$$p(s, q_1, q_2/\mathcal{F}^t) = \eta p(z_t/s, q_1, q_2)p(s, q_1, q_2/\mathcal{F}^{t-1}). \quad (28)$$

now, we may verify that the true map components $p_1 = 0, p_2 = 1$ are indeed a fixed point of the above recursion. In order to show this, note that given the observation z_1 and the true map components, $p(z_1/s, q_1, q_2) = 0$ for all possible combinations of (s, q_1, q_2) except $s = s_1, q_1 = E, q_2 = O$. Hence, it follows that given the true map probabilities and the observation z_1 , the map probabilities are updated to $p(s_1, q_1 = E, q_2 = O) = 1$ and hence the true map probabilities $p_1 = 0, p_2 = 1$ do not change and thus, are fixed points of the above iteration. Similarly, it can be shown that given z_2 , the true map probabilities do not change. Thus, we see that the only way that a consistent estimate of the map can be formed in the Bayesian approach, is by considering the joint distribution $p(s, q_1, q_2)$, i.e., the map components have to be correlated to each other as well as the robot pose. Note that the above does not prove that the method is consistent just that the true map probabilities are indeed fixed points of the recursion, which is more than can be said of the Bayesian recursions if the correlations are not considered.

Thus, in summary, Bayesian methods can indeed be used in the “-then map” part of the “localize first - then map” philosophy adopted in this paper but it does not buy much since the rest of the map has still got to be correlated to each other and the robot pose, to be consistent. On the other hand, the frequentist method, by virtue of its formulation, avoids the pitfall and can maintain the estimate of the map components independent of each other while remaining consistent.

III. EXPERIMENTS

In this section the methodology developed so far is applied to large maps with multiple cycles in them.. We have chosen to test our methods on such maps because of the well-known difficulties such maps pose to SLAM algorithms. The mobile robot is a differential drive vehicle. The equations of motion of such a robot is given by the following [35]:

$$\dot{x} = \frac{R}{2}(u_l + u_r)\cos\theta, \quad (29)$$

$$\dot{y} = \frac{R}{2}(u_l + u_r)\sin\theta, \quad (30)$$

$$\dot{\theta} = \frac{R}{L}(u_r - u_l), \quad (31)$$

where (x, y, θ) specify the pose of the robot, (u_l, u_r) are the left and right wheel angular velocities, and (R, L) are the radius of the wheel, and width of the robot respectively. The dimensions of the robot were as follows: a wheel radius of 25 cm, and a width of 50 cm. Experiments are performed for two different kinds of sensors: a) a noisy 2-D sensor with both range and bearing errors (such as a sonar) with $\sigma_r = 0.2$ m and $\sigma_\theta = 0.6$ deg, and b) an accurate 1-D range sensor such as a SICK laser range sensor with $\sigma_r = 0.01$ m and $\sigma_\theta = 0.05$ deg. We note here that the actual sensor noise given above is pumped up by a factor of approximately 10 in the EKF in order to maintain consistency (stabilizing noise), as otherwise the filter rapidly becomes overconfident, thereby leading to catastrophic failure on part of the hybrid scheme. The maximum range of the sensors was assumed to be 40 m. We assumed that there were feature points in the map that could be identified from the raw sensory data using suitable signal processing techniques. Mostly, these feature points were assumed to be either corners, or mid-points of corridors/ straight line segments. In the third example, we assume the availability of point landmarks. However, in our simulations, we do not identify these feature points from the raw sensory data as our intent is to show the efficacy of our methodology, and not the autonomous identification of such feature points. Any suitable feature-based SLAM algorithm can be used to achieve this goal. The observation of the feature-points are used in an EKF to form the belief on the robot pose, i.e., solve the Bayesian sub-problem of the hybrid formulation. Next, the belief state is used to map the rest of the dense map using the frequentist estimator E2. The process noise in the wheel encoders is $\sigma_u = 0.5$ rad/s. The average robot wheel speed is 5 rad/sec and the integration time step for the EKF is 0.5 sec. In the following, we show how the observation models $A(s)$ required in the frequentist estimator E2 are formed from the knowledge of the sensor noise specifications.

In the following, we show how the observation models $A(s)$ that are to required in the frequentist estimators E1/ E2 are extracted from the sensor noise models. We show the case of a range sensor corrupted by Gaussian noise. The method can be extended to sensors with both range and bearing errors in a relatively straightforward fashion.

Suppose in a gridded map the sensor is at some point A which is in grid q_A and the actual range reading which it should make be r^* . Let the range reading r^* correspond to some point R^* which is say in grid q_{r^*} . Let grid q_k be any grid that is being intersected by the ray AR^* . Then the sensor model in the discretized sense is given as the probability that the observation \hat{q}_k for grid q_k is occupied given that grid q_r is occupied.

$$P(\hat{q}_k = O/q_r = O) = \int_{q_k} f(r)dr = \int_{q_k} \frac{1}{\sigma\sqrt{2\pi}} \exp\left(-\frac{(r-r^*)^2}{2\sigma^2}\right)dr, \quad (32)$$

where $f(r)$ is the probability density function for normal distribution function with mean r^* and variance $\sigma(r^*)$. The integration is done over the interval for which the range reading r corresponds to the grid q_k . It should be noted that the sensor can get a reading from q_{r^*} only when all the grids between q_o and q_{r^*} are empty and this assumption is made implicitly. Thus we get the discretized sensor model.

While making observations, the assumption that all the grids between the vantage point and the observed grid is empty is not true, and the observation may very well be due to one of the intermediate grids too. We need $P(\hat{q}_k = O/q_k = O)$ and $P(\hat{q}_k = O/q_k = E)$ independent of this assumption where E means that the grid is empty. We label the grids that are being intersected by the ray AP as follows, where A is the vantage point, P is the point at which the observation falls, and q_k is the grid cell in the map in which the point P falls: the grids between q_A and q_K as $\{q_{l1}, q_{l2}, \dots, q_{ln}\}$ and the grids after grid q_k that are in the map as $\{q_{h1}, q_{h2}, \dots, q_{hm}\}$. Now we calculate the required probabilities using the following formulas which are a direct consequence of the law of total probability, and the assumption of the mutual independence of the different map components :

$$\begin{aligned} P(\hat{q}_k=O/q_k=O) &= P(\hat{q}_k=O/q_{l1}=O)P(q_{l1}=O) \\ &+ P(\hat{q}_k=O/q_{l1}=E, q_{l2}=O)P(q_{l2}=O)P(q_{l1}=E) + \dots \\ &+ P(\hat{q}_k=O/q_{l1}=E, \dots, q_{l(n-1)}=E, q_{ln}=O)P(q_{ln}=O) \\ &\quad \prod_{i=1}^{n-1} P(q_{li}=E) \\ &+ P(\hat{q}_k=O/q_{l1}=E, \dots, q_{ln}=E, C_k=O) \prod_{i=1}^n P(q_{li}=E) \end{aligned} \quad (33)$$

$$\begin{aligned} P(\hat{q}_k=O/q_k=E) &= P(\hat{q}_k=O/q_{l1}=O)P(q_{l1}=O) \\ &+ P(\hat{q}_k=O/q_{l1}=E, q_{l2}=O)P(q_{l2}=O)P(q_{l1}=E) + \dots \\ &+ P(\hat{q}_k=O/q_{l1}=E, \dots, q_{l(n-1)}=E, q_{ln}=O)P(q_{ln}=O) \\ &\quad \prod_{i=1}^{n-1} P(q_{li}=E) \\ &+ P(\hat{q}_k=O/q_{l1}=E, \dots, q_{ln}=E, q_{h1}=O)P(q_{h1}=O) \\ &\quad \prod_{i=1}^n P(q_{li}=E) \\ &+ P(\hat{q}_k=O/q_{l1}=E, \dots, q_{ln}=E, \hat{q}_k=E, q_{h1}=E, q_{h2}=O) \\ &\quad P(q_{h2}=O)P(q_{h1}=E) \prod_{i=1}^n P(q_{li}=E) + \dots \\ &+ P(\hat{q}_k=O/q_{l1}=E, \dots, q_{ln}=E, \hat{q}_k=E, q_{h1}=E, \dots, q_{hm}=O) \\ &\quad P(q_{hm}=O) \prod_{i=1}^{m-1} P(q_{hj}=E) \prod_{i=1}^n P(q_{li}=E) \end{aligned} \quad (34)$$

Now we can define the observation matrix A_k for a grid q_k when observing from point O as follows.

$$\Gamma_k = \begin{pmatrix} P(\hat{q}_k = O/q_k = O) & P(\hat{q}_k = O/q_k = E) \\ P(\hat{q}_k = E/q_k = O) & P(\hat{q}_k = E/q_k = E) \end{pmatrix}$$

$P(\hat{q}_k = E/q_k = O) = 1 - P(\hat{q}_k = O/q_k = O)$ and $P(\hat{q}_k = E/q_k = E) = 1 - P(\hat{q}_k = O/q_k = E)$. Note here that the observation matrix is dependent on the real probabilities of the map components $P(q_K = O)$ and $P(q_k = E)$. In fact this is what we are trying to estimate. So, to compute the observation matrix, as was mentioned in Section II, the current estimates of the map components are used.

Figs 3-9 show the results of our simulation experiments. Four different maps are considered in these experiments. The first map (Map 1) consists of a large cyclic corridor of side 100m, Map 2 is a long hallway (100m x 40m) with 4 cyclic corridors, Map 3 is a long hallway (220m x 40m) with 10 cyclic corridors, Map 4 is a large map with 20 cycles and is a kilometer long and 200 m wide (and could be thought of as several city blocks). A total of 2 laps of each map is made. The total length of the runs was approximately 1 km for Maps 1 and 2, approximately 2.5 Kms for Map3, and 8 km for Map 4. Each of these maps, except map 4, are sensed using both the noisy 2-D sensor, as well as the much more accurate Laser range sensor. Figs

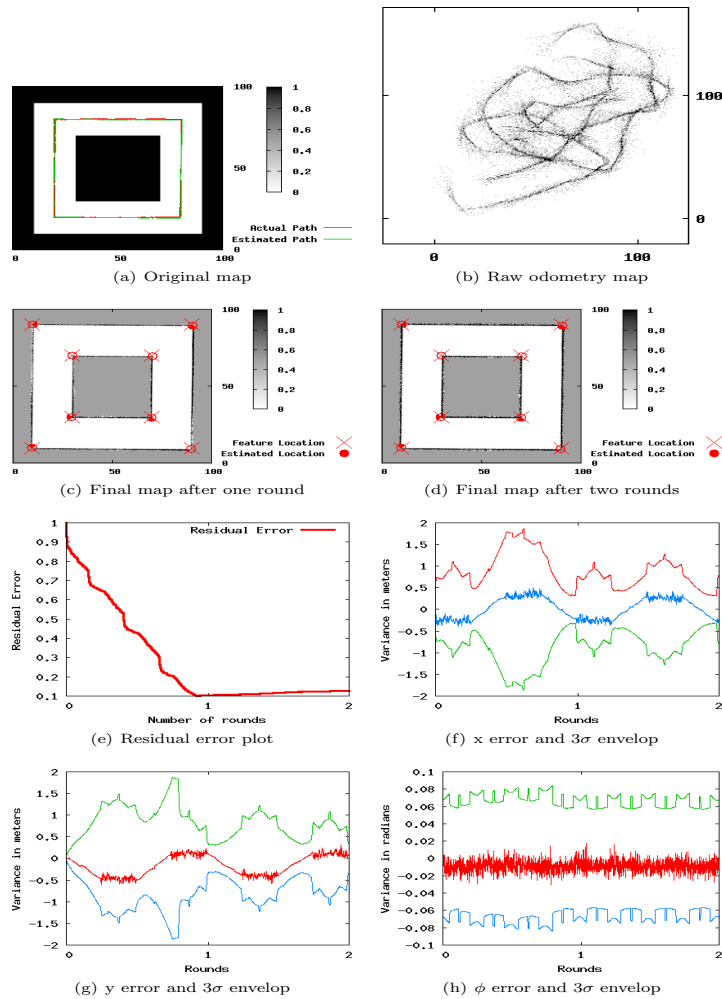


Fig. 3. Experimental results for Map 1 with noisy 2-D sensor

3-4 represent the results for Map 1, Figs 5-6 represent the results for Map 2, Figs. 7-8 represent the results for Map 3 and Fig. 9 for Map 4. In each of these figures, Subfigure (a) shows the original map, Subfigure (b) shows the raw odometry data, Subfigures (c) and (d) shows the estimated map, along with the features and their estimates, after the completion of one and three rounds respectively, Subfigure (e) shows the total error in the map as a function of the number of rounds the robot makes, Subfigures (f)-(h) show the error in the x , y , and θ co-ordinates of the robot along with their associated 3σ bounds. The features used in these maps were the corners of the corridors and were assumed to be reliably identified.

These figures give us an idea as to how well the algorithm is performing and also give us valuable practical insight into the algorithm. One of the reasons we chose these examples is because of the well-known challenge maps with multiple cycles pose to SLAM algorithms which is evidenced from the raw odometry plots (Subfigure (b) in the plots). The results show that the algorithm is able to map a large area with multiple cycles without much of a problem even though the sensors are quite noisy. Moreover, qualitatively (Subfigures (c), (d) in the plots) as well as quantitatively (Subfigure (e) in the plots), the results improve manifold with the more accurate Laser range sensor when compared to the noisier 2-D sensor. Thus, this reinforces the intuitive idea that much larger maps can be mapped more accurately with an accurate sensor such as the range sensor when compared to a noisier sensor such as the 2-D sensor. The algorithm had no problems in closing large loops as the ones shown here and we did not have to make any heuristic corrections when such a loop was closed. In fact, the size of the map, or the number of cycles in it, is really not a problem for this algorithm as long as the EKF remains consistent. However, if the EKF loses consistency then the guarantees of consistency for the frequentist part of the hybrid formulation are no longer valid.

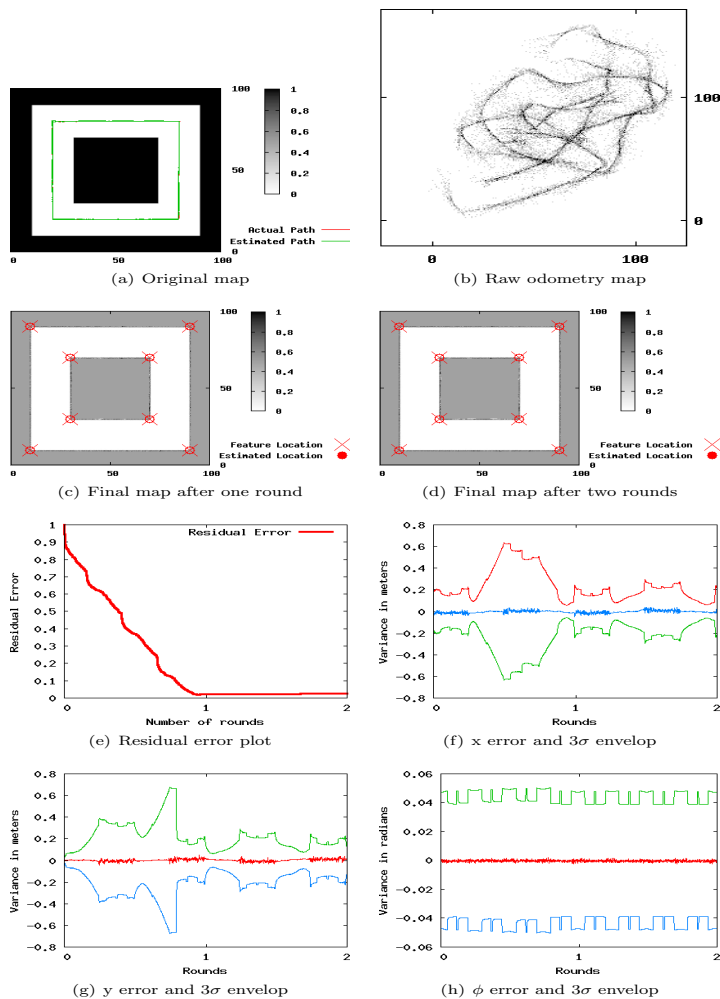


Fig. 4. Experimental results for Map 1 with accurate Laser range sensor

Thus, sophisticated Kalman/ Information filter based methods for feature-based SLAM can play an important role in ensuring the consistency of the hybrid formulation for maps on a much larger scale such as cities, planetary terrains etc., where the order of the distances is in hundreds or thousands of kilometers, and consequently, where the number of features/ landmarks would increase by several orders of magnitude when compared to the maps shown here. It can also be seen from the total map error plots (Subfigure (e) in the plots) that the mapping algorithm converges exponentially. In fact, it is our conjecture that this can be rigorously established, and is supported by the experimental evidence which suggests that this indeed seems to be the case. Hence, we may conclude that the results show sufficient evidence of the efficacy/ applicability of the methodology proposed in this paper.

IV. CONSISTENCY

In this section, the consistency of the frequentist mapping algorithm under uncertain robot poses is established. In our context, consistency implies that the estimated map probabilities converge to the true map probabilities with probability one, or almost surely. We shall prove the result using the powerful ODE method from the stochastic approximation literature [36], [37]. To begin, we present a short introduction of the method to clarify the basic idea behind the method. Consider the mapping algorithm as presented previously:

$$P_{i,t+1} = P_{i,t} + \gamma[c_i(X_t, P_t) - A_i(X_t, P_t)P_{i,t}], \quad (35)$$

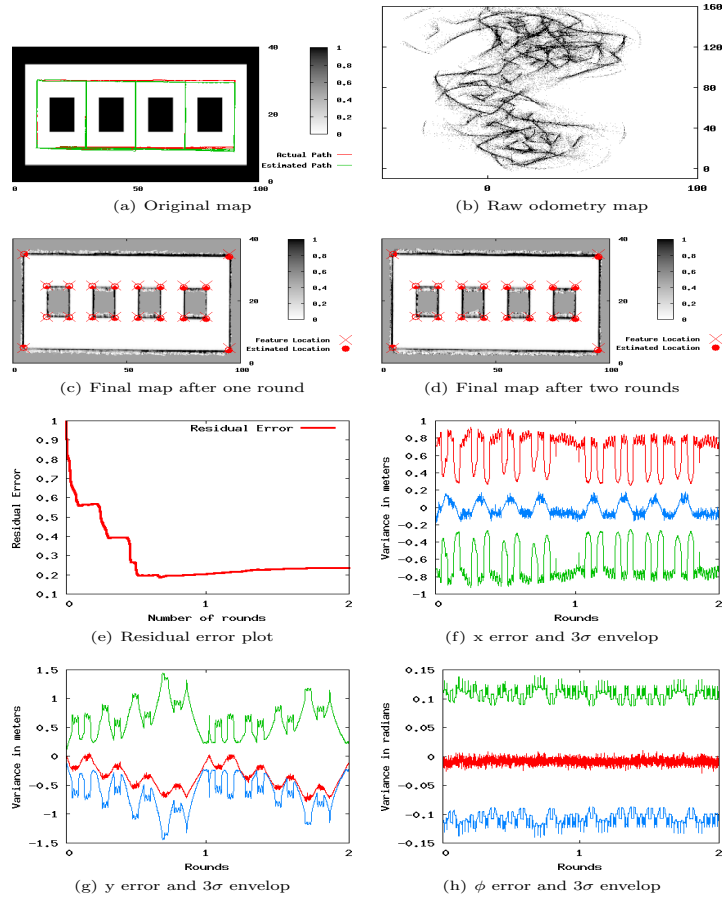


Fig. 5. Experimental results for Map 2 with noisy 2-D sensor

where $X_t = (b_t, z_t)$, the 2-tuple consisting of the belief state and the observation at any instant. If the learning rate parameter γ is small, then the value of the map probabilities does not change quickly, and can be assumed to be essentially equilibrated over N steps, and then

$$P_{i,t+N} \approx P_{i,t} + \gamma N \left(\frac{1}{N} \sum_{k=1}^N [c_i(X_{t+k}, P_t) - A_i(X_{t+k}, P_t) P_{i,t}] \right), \quad (36)$$

where X_{t+k} is the state of a Markov chain that depends on the current estimate of the map probabilities $P_{i,t}$. It can be seen that the belief process is a Markov process that would depend on the map probabilities in closed loop, or otherwise would be independent of the map probabilities. The method works either ways and possibly might be better in open loop in terms of convergence rates. Then, using the law of large numbers, it follows that if N is large enough:

$$\frac{1}{N} \sum_{k=1}^N [c_i(X_{t+k}, P_t) - A_i(X_{t+k}, P_t) P_{i,t}] \approx \bar{h}_i^*(P_t) - \bar{A}_i(P_t) P_{i,t}, \quad (37)$$

where $\bar{h}_i^*(\cdot)$ and $\bar{A}_i(\cdot)$ are the averaged values of $c_i(\cdot)$ and $A_i(\cdot)$ respectively. Then, it follows that

$$P_{i,t+N} - P_{i,t} \approx N\gamma [\bar{h}_i^*(P_t) - \bar{A}_i(P_t) P_{i,t}], \quad (38)$$

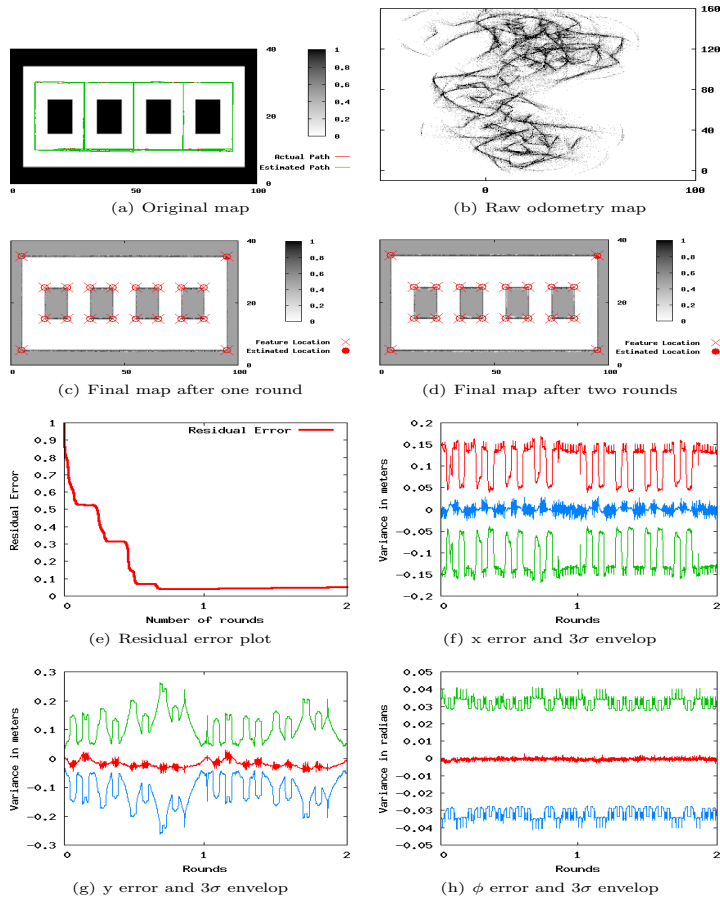


Fig. 6. Experimental results for Map 2 with accurate Laser range sensor

where $\bar{h}_i(P_t) = \bar{A}_i(P_t)P_{i,t}$, which happens to be the forward Euler approximation (with step size $N\gamma$) of the differential equation:

$$\dot{P}_i = \bar{h}_i^*(P) - \bar{h}_i(P). \quad (39)$$

The idea behind the method is that the asymptotic performance of the estimation/ mapping algorithm can be analyzed by analyzing the behaviour of the “mean/ average” ODE above. This method is very popular in analyzing the behaviour of algorithms in many different fields including reinforcement learning [38], neural networks [39], system identification and stochastic adaptive control [36], [37]. In the following, we analyze the mapping algorithm rigorously using the ODE method.

Consider the mapping algorithm as presented previously:

$$P_{i,t+1} = \Pi_{\mathcal{P}}\{P_{i,t} + \gamma_t[c_i(X_t, P_t) - A_i(X_t, P_t)P_{i,t}]\}, \quad (40)$$

where $X_t = (b_t, z_t)$ is the 2-tuple consisting of the belief state and the observation at any instant. The mean true observation probabilities of the i^{th} map component and the mean “current” predicted value are defined as:

$$h_i^*(b, p) \equiv E_z^*[c_i(b, z, P)] = \sum_z p^*(z/b)c_i(b, z, P), \quad (41)$$

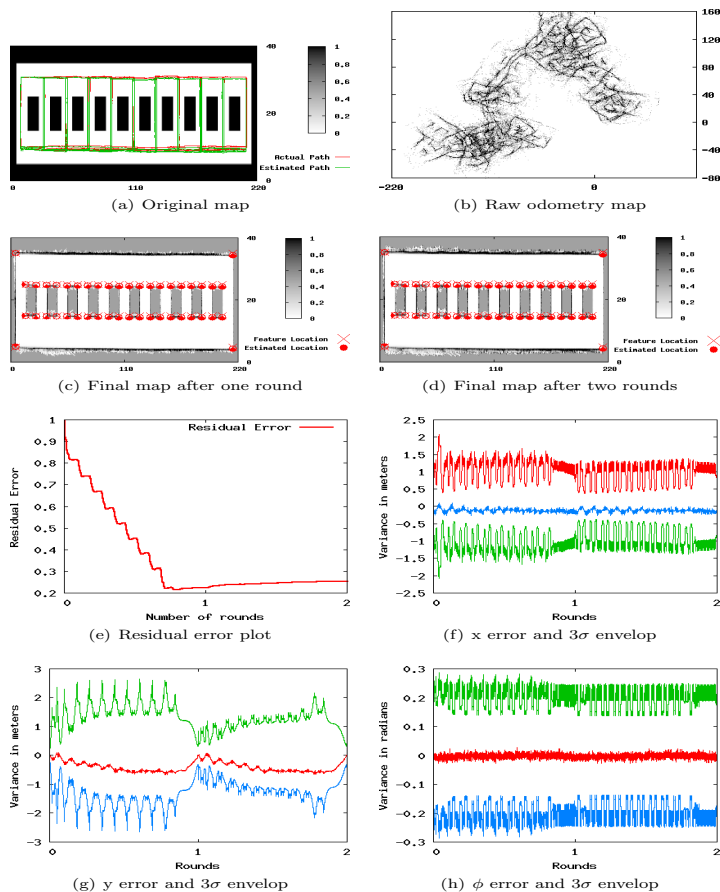


Fig. 7. Experimental results for Map 3 with noisy 2-D sensor

$$h_i(b, P) \equiv E_z[c_i(b, z, P)] = A_i(b, z, P)P_i = \sum_z p(z/b)c_i(b, z, P), \quad (42)$$

where $p^*(z/b)$ is the probability of an observation z given the true map probabilities P^* , and $p(z/b)$ are the probabilities given the estimate of the map probabilities P , and given that the belief state is b . Recall that $c_i(b, z, P)$ is the vector containing the observation probabilities of the i^{th} map component in its various different states, given that the belief on the robot pose is b , the reading from the sensor is z and the estimate of the map probabilities is P . Note that $c_i(b, z, P)$ is the approximation of the true observation probability vector $c_i^*(b, z, P)$ where the vector of true map probabilities, P^* , is replaced by the approximate map probabilities P (see the end of Section II B). If the map probabilities were truly P , then if we averaged $C_i(b, z, P)$ over all observations z , we would obtain the quantity $h_i(b, P)$. However, since the observation z is generated by the true map probabilities P^* , not P , we would obtain the quantity $h_i^*(b, P)$ which is in general, different from $h_i(b, P)$. In fact, only at $P = P^*$ are the two quantities equal and the algorithm uses this fact to guide the map estimates towards the true map probabilities. It can be seen that

$$p^*(z/b) = \sum_{s, q_1, q_2, \dots, q_M} p(z/s, q_1, \dots, q_M)p^*(q_1) \cdots p^*(q_M)b(s), \quad (43)$$

$$p(z/b) = \sum_{s, q_1, q_2, \dots, q_M} p(z/s, q_1, \dots, q_M)p(q_1) \cdots p(q_M)b(s). \quad (44)$$

In the following, for the sake of simplicity (the extension is reasonably straightforward), we deal exclusively with the

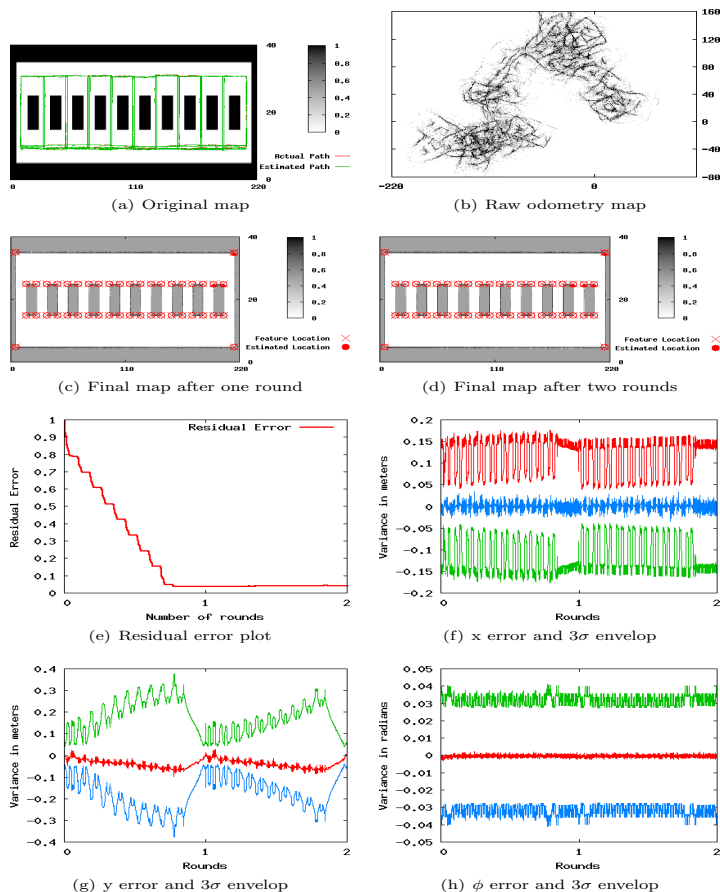


Fig. 8. Experimental results for Map 3 with accurate Laser range sensor

2-dimensional case, i.e., when each map component q_i can take one of two values 0/1, the extension to $D > 2$ is similar with a few modifications. The map probabilities can now be denoted by the vector $P = [P_1, P_2, \dots, P_M]$, where P_i denotes the probability that the i^{th} map component takes value 1 (i.e., is occupied).

First, we make the following assumption.

A 1. *Corresponding to every map probability P , let the belief process b_t have a stationary distribution $\pi_\infty(b, P)$. Moreover, let the belief process be geometrically ergodic. Let*

$$\bar{h}_i^*(P) = \int_{b \in B_i} h_i^*(b, P) \frac{d\pi_\infty(b, P)}{\pi_\infty(B_i, P)}, \bar{h}_i(P) = \int_{b \in B_i} h_i(b, P) \frac{d\pi_\infty(b, P)}{\pi_\infty(B_i, P)}, \quad (45)$$

where B_i are all the belief states that map component i is observed from. In particular, the above assumption implies that there exist $K < \infty$, $\rho < 1$ such that

$$E[C_i(b_t, z_t, P) - \bar{h}_i^*(P)] \leq K\rho^t, E[A_i(b_t, P)P_i - \bar{h}_i(P)] \leq K\rho^t, \quad (46)$$

i.e., the quantities $C_i(\cdot)$ and $A_i(\cdot)$ converge to their average values exponentially fast.

Define $\bar{H}^*(P) = [\bar{h}_1^*(P), \dots, \bar{h}_M^*(P)]^t$, and $\bar{H}(P) = [\bar{h}_1(P), \dots, \bar{h}_M(P)]^t$. Then, under assumption A 1, it can be shown that the asymptotic behaviour of the mapping algorithm is characterized by the solution of the following ODE ([36], pg. 187,

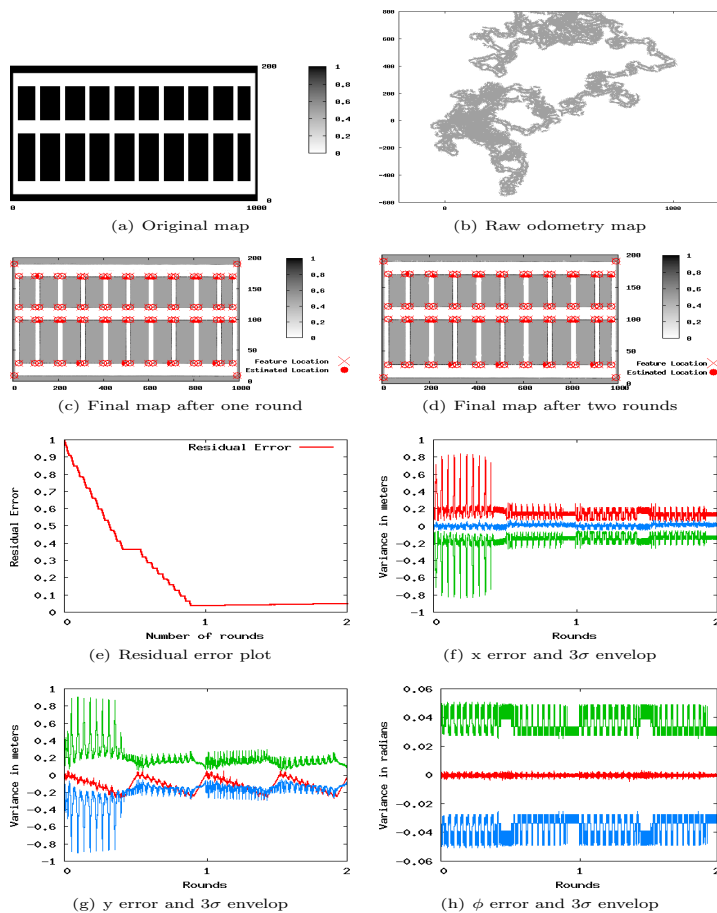


Fig. 9. Experimental results for Map 4 with accurate Laser range sensor

Ch. 6, Theorem 6.1):

$$\dot{P} = \bar{H}^*(P) - \bar{H}(P). \quad (47)$$

In particular, the following result holds.

Proposition 1. *Let the point $P = P^*$ be an asymptotically stable equilibrium of the ODE (47) with domain of attraction D^* . Let $C \subseteq D^*$ be some compact subset of D^* . Let the learning rate parameters $\{\gamma_t\}$ be such that $\sum_t \gamma_t = \infty$, and $\gamma_t \rightarrow 0$ as $t \rightarrow \infty$. If the trajectories of the mapping algorithm (40) enter the subset C infinitely often, then the estimates $P_t \rightarrow P^*$ almost surely.*

Hence, it is left for us to show that the set of true map probabilities $P^* = [P_1^*, \dots, P_M^*]$ is an asymptotically stable equilibrium of ODE (47). In order to show this, we will show that the linearization of the mean ODE (47) about P^* is asymptotically stable and hence, so is the nonlinear ODE ([40], Chapter 3, Theorem 3.7). We limit our treatment to the case when $D = 2$, i.e., the map components can take one of two values. The gradient of the vector field $\bar{H}^*(P) - \bar{H}(P)$ is defined by the matrix:

$$\nabla(\bar{H}^*(P) - \bar{H}(P)) = [\partial_i(\bar{h}_j^*(P) - \bar{h}_j(P))]. \quad (48)$$

We make the following assumption.

A 2. We assume that

$$\partial_i(\bar{h}_i^*(P) - \bar{h}_i(P)) < -\epsilon, \forall i, \quad (49)$$

$$\left| \sum_{j \neq i} [\partial_j(\bar{h}_i^*(P) - \bar{h}_i(P))] \right| \leq |\partial_i(\bar{h}_i^*(P) - \bar{h}_i(P))|, \quad (50)$$

where all the partial derivatives above are evaluated at $P = P^*$, and $\epsilon > 0$ is a positive constant.

We will provide the justification of the above assumptions after the following proposition.

Proposition 2. Let assumption A 2 be satisfied. Then the true map probability vector P^* is an asymptotically stable equilibrium of ODE 47 with a non-empty region of attraction D^* . Thus, if the mapping algorithm estimates P_t visit a compact subset $C \subseteq D^*$ infinitely often, then $P_t \rightarrow P^*$ almost surely, due to Proposition 1.

Proof. Under assumption A 2, the linearization of the mean ODE (47) about P^* is row-dominant, and hence, all its eigenvalues lie in the open left half plane and their real parts are bounded atleast ϵ away from the imaginary axis [41]. Therefore, it follows that P^* is an asymptotically stable equilibrium point of the mean ODE and hence, all the other results follow. \square

Finally, we furnish the justification for assumption A 2.

In order to do this, note that

$$\partial_j(\bar{h}_i^*(P) - \bar{h}_i(P)) = \int_{b \in B_i} \partial_j(h_i^*(b, P) - h_i(b, P)) \frac{d\pi_\infty(b, P)}{\pi_\infty(B_i, P)}. \quad (51)$$

Consider the term $\partial_j(h_i^*(b, P) - h_i(b, P))$ for some belief state b . In the following to simplify notation, all the partial derivatives are assumed to be evaluated at $P = P^*$. It may be shown that:

$$\partial_j(h_i^*(b, P) - h_i(b, P)) = - \sum_z \partial_j p(z/b) c_i(b, z, P), \quad (52)$$

where the partial derivative above is evaluated at $P = P^*$, and that:

$$\partial_j p(z/b) = p^*(z/b, q_j = 1) - p^*(z/b, q_j = 0), \quad (53)$$

i.e., the difference in the probabilities of observing z given belief state b , and whether q_j is in state 1 or state 0. Then, it follows that

$$\partial_j(h_i^*(b, P) - h_i(b, P)) = - \sum_z (p^*(z/b, q_j = 1) - p^*(z/b, q_j = 0)) p^*(\hat{q}_i = 1/z, b), \quad (54)$$

where recall that the variable $p^*(\hat{q}_i = 1/z, b) = c_i^*(b, z)$ is the ‘‘true’’ probability that the observation of the i^{th} map component is 1 given the belief state b and the observation z (see Section 2.2). Hence, it follows that

$$\begin{aligned} \partial_i(h_i^*(b, P) - h_i(b, P)) &= - \sum_z [p^*(z/b, q_i = 1) - p^*(z/b, q_i = 0)] p^*(\hat{q}_i = 1/z, b), \\ &= - [p^*(\hat{q}_i = 1/q_i = 1, b) - p^*(\hat{q}_i = 1/q_i = 0, b)]. \end{aligned} \quad (55)$$

Hence, $\partial_i(h_i^*(b, P) - h_i(b, P)) < -\epsilon$ if $p^*(\hat{q}_i = 1/q_i = 1, b) - p^*(\hat{q}_i = 1/q_i = 0, b) > \epsilon$. The above equation thus implies that the probability of observing the map component occupied when it is actually occupied should be more than the probability of seeing it occupied when it is actually unoccupied (i.e., a spurious observation due to some other map component). This in turn is a ‘‘good sensor’’ assumption, i.e, we see the right observation more times than the wrong one. This corresponds to the heuristic that we neglect or disregard the observations of the map components that are too far from our current belief state. Thus, the set B_i in Eq. 51 above should consist of only those belief states from which the observation of map component i can be reliable. Ensuring that the set B_i is chosen in the above fashion implies that $\partial_i(h_i^*(b, P) - h_i(b, P)) < -\epsilon$ for all $b \in B_i$, and hence, it follows that Eq. (49) is automatically satisfied.

Recall the definitions of $h_i(b, P)$ and $h_i^*(b, P)$ (cf. Eqs. (41)-(42)). The difference between the two signifies the average observation prediction error of the i^{th} map component given that the map probability estimates are P . Recall that it is zero for $P = P^*$. Thus, $\partial_j(h_i^*(b, P) - h_i(b, P))$ represents the sensitivity of this error to the j^{th} component of the map probabilities. Now, if we require that

$$\left| \sum_{j \neq i} [\partial_j(\bar{h}_i^*(b, P) - \bar{h}_i(b, P))] \right| \leq |\partial_i(\bar{h}_i^*(b, P) - \bar{h}_i(b, P))|, \quad (56)$$

for all $b \in B_i$, then Eq. (50) in assumption 2 is automatically satisfied. The equation above implies that the sensitivity of the observation error of the i^{th} map component to its own map probabilities should dominate the cumulative sensitivity of the error to all other map components. This is a structural assumption that is required regardless of whether the robot pose is uncertain or perfectly known. In fact, it may be reasonably expected that it is satisfied if the map components are updated

only from “good” belief states, i.e., belief states from which the sensors can be assumed to be reliable. Experimental evidence seems to suggest the same as well. The development above can then be summarized in the following proposition.

Proposition 3. *Given any map component q_i let there exist a set of belief states G_i s.t. for all beliefs $b \in G_i$, the following hold:*

$$p^*(\hat{q}_i = 1/q_i = 1, b) - p^*(\hat{q}_i = 1/q_i = 0, b) > \epsilon,$$

$$|\partial_i(h_i^*(b, P) - h_i(b, P))| \geq \left| \sum_{j \neq i} \partial_j(h_i^*(b, P) - h_i(b, P)) \right|,$$

where all the partial derivatives above are evaluated at $P = P^*$. If the sets B_i are chosen such that $B_i \subseteq G_i$, then assumption 2 is automatically satisfied and hence, Proposition 2 holds.

This completes our proof of the consistency of the mapping algorithm given an uncertain time-varying belief state. In practice, the sets B_i have to be chosen heuristically, and some trial and error might be required before we arrive at the right choice for a particular type of sensor. In the examples in the following section, the sets B_i were chosen such that we discard observations of map components that are further than 40 m from the mean of the robot pose belief state. We conclude this section by noting that the consistency of the Bayesian sub-problem is a necessary condition for the consistency of the hybrid method since the above proof assumes that the true belief on the pose of the robot is output to the frequentist mapping method.

V. CONCLUSION

In conclusion, we have presented a hybrid methodology for the SLAM problem. We have proved the consistency of the algorithm and shown that it is robust to the data association problem while having complexity linear in the map components. The methodology was tested on large maps with multiple cycles. The experimental results seem to confirm the applicability of the mapping methodology outlined in the development above. However, much still needs to be done including the testing of the algorithms on larger and more complicated maps than the ones presented in the paper, though we do not see that as a big bottleneck, as well as experimental evaluation on a robotic testbed in the real world. Also, the extension of the methodology to dynamic maps and multiple robot problems has to be considered. Moreover, the methodology is very general and applies to any mobile sensor system, and as such, future research can also focus on the use of such systems to map general environmental characteristics such as chemical concentrations in the atmosphere, ocean temperature and soil salinity.

ACKNOWLEDGEMENTS

This work was supported in part by Army Research Office (ARO) grant number W911NF-07-1-0429. The views and conclusions contained herein are those of the authors and should not be interpreted as necessarily representing the official policies or endorsements, either expressed or implied, of the ARO or the U.S. Government.

REFERENCES

- [1] M. W. M. Dissanayake, P. Newman, S. Clark, and H. F. Durrant-Whyte, “A solution to the simultaneous localization and map building (slam) problem,” *IEEE Transactions on Robotics and Automation*, vol. 17, pp. 229–241, 2001.
- [2] R. Smith, M. Self, and P. Cheeseman, “Estimating uncertain spatial relationships in robot vehicles,” in *Autonomous Robotic Vehicles*, I. J. Cox and G. T. Wilfon, eds., 1990.
- [3] P. Moutarlier and R. Chatila, “Stochastic multisensor data fusion for mobile robot localization and environment modeling,” in *Proceedings Int. Symp. Rob. Res.*, 1989.
- [4] J. Leonard and H. Durrant-Whyte, *Directed Sonar Sensing for Mobile Robot Navigation*. Boston, MA: Kluwer academic, 1992.
- [5] T. Bailey and H. Durrant-Whyte, “Simultaneous localization and mapping (slam): Part ii,” *IEEE Robotics and Automation Magazine*, pp. 108–117, Sep. 2006.
- [6] H. Durrant-Whyte and T. Bailey, “Simultaneous localization and mapping (slam): Part i,” *IEEE Robotics and Automation Magazine*, pp. 99–108, June 2006.
- [7] S. Thrun, Y. Liu, D. Koeller, A. Y. Ng, S. Ghahramani, and H. Durrant-Whyte, “Simultaneous localization and mapping with sparse extended information filters,” *International Journal of Robotics Research*, vol. 23, pp. 693–716, 2004.
- [8] R. Eustice, H. Singh, and J. Leonard, “Exactly sparse delayed state filters,” in *Proc. IEEE Int. Conf. Rob. Aut.*, 2005.
- [9] M. Walter, R. Eustice, and J. Leonard, “A provably consistent method for imposing sparsity in feature-based slam information filters,” in S. Thrun, R. Brooks, H. Durrant-Whyte (Eds.), *Robotics Research, STAR 28*, 2007, pp. 214–234.
- [10] M. Montemerlo, S. Thrun, and W. Whittaker, “Conditional particle filters for simultaneous mobile robot localization and people tracking,” in *IEEE ICRA*, 2002.
- [11] M. M. et. al., “Fastslam 2.0: An improved particle filtering technique for simultaneous localization and mapping that provably converges,” in *Proc. of IJCAI*, 2003.
- [12] K. Murphy, “Bayesian map learning in dynamic environments,” in *Advances in Neural Information Processing Systems (NIPS)*, 1999.
- [13] A. Eliazar and R. Parr, “Dp-slam: Fast, robust simultaneous localization and mapping without predetermined landmarks,” in *Proceedings of the Eighth International Joint Conference on Artificial Intelligence (IJCAI 03)*, 2003.
- [14] G. Grisetti, C. Stachnis, and W. Burgard, “Improved techniques for grid mapping with rao-blackwellized particle filters,” *IEEE Trans. Robotics*, vol. 23, pp. 34–46, 2007.
- [15] G. G. et. al., “Fast and accurate slam with rao-blackwellized particle filters,” *Rob. Aut. Syst.*, vol. 55, pp. 30–38, 2007.
- [16] C. Stachnis, *Exploration and Mapping with Mobile Robots, PhD Thesis*. Freiburg, Germany: University di Freiburg, 2006.
- [17] F. Lu and E. Milios, “Globally consistent range scan alignment for environment mapping,” *Autonomous Robots*, vol. 18, pp. 249–275, 1997.

- [18] J. S. Gutmann and K. Konolige, "Incremental mapping of large cyclic environments," in *Proc. IEEE Int. Symp. on Comp. Intell. in Rob. and Aut. (CIRA)*, 2000.
- [19] S. Thrun and M. Montemerlo, "The GraphSLAM algorithm with applications to large-scale mapping of urban structures," *International Journal on Robotics Research*, vol. 25, no. 5/6, pp. 403–430, 2005.
- [20] T. Bailey, J. Nieto, J. Guivant, M. Stevens, and E. Nebot, "Consistency of the ekf-slam algorithms," in *Proceedings of the IEEE Int. Conf. Rob. automat.*, 2006.
- [21] T. Bailey, J. Nieto, and E. Nebot, "Consistency of the fastslam algorithm," in *Proceedings of the IEEE Int. Conf. Rob. automat.*, 2006.
- [22] M. Klaas, N. DeFreitas, and A. Doucet, "Towards practical n^2 monte carlo: the marginal filter," in *Uncertainty in Artificial Intelligence*, 2005.
- [23] R. Martínez-Cantin, N. deFreitas, and J. A. Castellanos, "Analysis of particle methods for simultaneous robot localization and mapping and a new algorithm: Marginal-slam," in *Proceedings of the IEEE Int. Conf. Rob. automat.*, 2007.
- [24] A. Elfes, "Using occupancy grids for mobile robot perception and navigation," *IEEE Computer*, vol. 22, pp. 46–57, 1989.
- [25] J. Nieto, J. Guivant, and E. Nebot, "Denselam: Simultaneous localization and dense mapping," *The International J. Rob. Res.*, vol. 25, pp. 711–744, 2006.
- [26] J. N. et. al., "The hybrid metric maps (hymm)," in *IEEE Int. Conf. Rob. Automat.*, 2004.
- [27] S. Thrun, D. Fox, and W. Burgard, "A probabilistic approach to concurrent mapping and localization for mobile robots," *Mach. Learn. Aut. Robots*, vol. 31, pp. 29–53, 1998.
- [28] H. Shatkay and L. P. Kaelbling, "Learning topological maps with local odometric information," in *Proc. IJCAI-97*, 1997.
- [29] W. Burgard, D. Fox, H. Jans, C. Matenar, and S. Thrun, "Sonar-based mapping of large scale mobile robot environments using em," in *Proc. Int. Conf. Mach. Learn.*, 1999.
- [30] L. R. Rabiner, "A tutorial on hidden markov models and selected applications in speech recognition," *Proceedings of the IEEE*, vol. 77, pp. 257–286, 1988.
- [31] Y. Ephraim and N. Merhav, "Hidden markov processes," *IEEE Transactions Inf. Theory*, vol. 48, pp. 1518–1569, 2002.
- [32] S. Chakravorty and J. L. Junkins, "Motion planning in uncertain environments with vision like sensors," *Automatica*, vol. 43, pp. 2104–2111, 2007.
- [33] J. Davis and S. Chakravorty, "Motion planning under uncertainty: Application to an unmanned helicopter," *Journal of Guid., Contrl. and Dyn.*, vol. 30, pp. 1268–1277, 2007.
- [34] S. Chakravorty, "A hybrid approach to the simultaneous planning localization and mapping problem," in *Proceedings of the American Control Conference*, 2008.
- [35] S. M. LaValle, *Planning Algorithms*. Cambridge, UK: Cambridge University Press, 2005.
- [36] H. J. Kushner and G. G. Yin, *Stochastic Approximation and Recursive Algorithms and Applications*. New York, NY: Springer, 2003.
- [37] A. Benveniste, M. Metivier, and P. Priouret, *Adaptive Algorithms and Stochastic Approximations*. Berlin Heidelberg: Springer-Verlag, 1990.
- [38] D. P. Bertsekas and J. N. Tsitsiklis, *Neuro-Dynamic Programming*. Boston, MA: Athena Scientific, 1996.
- [39] S. Haykin, *Neural Networks*. Upper Saddle River, NJ: Prentice-Hall, 1994.
- [40] H. K. Khalil, *Nonlinear Systems*. Upper Saddle River, NJ: Prentice Hall, 1996.
- [41] R. A. Horn and C. R. Johnson, *Matrix Analysis*. Cambridge, UK: Cambridge University Press, 1993.