2008 American Control Conference
Westin Seattle Hotel, Seattle, Washington, USA
June 11-13, 2008

Simultaneous Planning Localization and Mapping: A Hybrid Bayesian/ Frequentist Approach

S. Chakravorty * and R. Saha

Abstract

In this paper, the problem of mapping and planning in an uncertain environment is studied. A hybrid Bayesian/ frequentist formulation of the simultaneous planning, localization and mapping (SPLAM) problem is presented wherein the environment is modeled as a stationary, spatially uncorrelated random process whose stationary probabilities are fixed but unknown, and have to be estimated as the autonomous system moves through the environment and makes observations using its sensors. The environmental random process is estimated using stochastic approximation algorithms. Under a certain "reliable sensor assumption", it is shown that the mapping algorithms converge with probability one, and that the convergence of the mapping algorithms is independent of the planning policy, as long as it is non-anticipative, akin to the celebrated "Separation Principle" in Classical Linear Control theory. Further, the computational burden of the mapping algorithms is significantly reduced when compared to Bayesian SPLAM techniques.

1. Introduction

In this paper, the problem of mapping and planning in an uncertain environment is considered. The uncertainty of the autonomous system is modeled as a completely known Markov decision process while the environment is modeled as a stationary, spatially uncorrelated (independent) random process whose parameters are unknown and have to be estimated while the system navigates through the environment. It is assumed that there is error in observing the environment. The observation model is assumed to be known, given the location of the autonomous system. A frequentist (stochastic approximation based) approach is presented for the estimation of the environment process, i.e., the mapping problem. It is shown that under a certain "Reliable Sensor Assumption" the estimation and planning algorithms are guaranteed to converge with probability one. Moreover, it is shown that the convergence of the estimation algorithms are completely independent of the planning algorithms and thus, satisfy a "Separation Principle" as in classical linear control theory. It is further shown that the computation burden of the algorithms are significantly less when compared to Bayesian SPLAM techniques.

The problem of simultaneous localization and mapping (SLAM) or additionally simultaneous planning, localization and mapping (SPLAM) has received considerable attention in the Robotics community in the past several years . The generic SLAM problem consists of an autonomous system navigating in an unknown environment, which it is trying to map while simultaneously localizing itself with respect to the map that it is building. This creates a philosophical "chicken and egg" problem which leads to a very high dimensional computational problem [1–4]. In most SLAM techniques, the localization and mapping problem is posed as a Bayesian filtering problem wherein the environment is considered to be a fixed but unknown parameter. There are two basic approaches to solving the Bayes filtering problem. the first alternative is to use the Kalman filtering technique which is applicable to linear-Gaussian systems [4–7]. However, this method cannot accommodate cases where the distributions are non-Gaussian and cannot provide a solution to the so called "data association" problem [2, 8]. The second method consists of solving the Bayesian filtering problem using particle filtering techniques [10]. These methods can accommodate the multi-modal nature of the probability distributions and the data association problem gracefully [2, 12]. The basic drawback with the Bayesian formulation of the SLAM problem is that the estimates of the various environmental components (features) become correlated even though their measurements are mutually independent. This is a basic structural property of the Bayesian formulation of the SLAM problem [1–4, 12].

*Suman Chakravorty is an Assistant Professor, Department of Aerospace Engineering, Texas A&M University, College Station, schakrav@aero.tamu.edu

The correlation between the environmental components can be eliminated by using Rao-Blackwellized particle filters that keep track of the whole trajectory of the robot instead of just the current pose [3, 12]. However, in this case the "curse of dimensionality" in the environment is traded off for the "curse of history" and can be a big limitation in large environments. Thus, the filter has to estimate a very high dimensional probability distribution on the environment which becomes increasingly computationally intractable as the size of the environment increases [2,3,12,13]. The addition of planning to the SLAM problem, resulting in the SPLAM problem, adds further to the complexity of the problem [14–17]. In fact, the planning problem on its own is computationally quite intractable under uncertainty. In this paper, an alternative to the Bayesian formulation of the SPLAM problem is proposed. In this formulation, the environment is modeled as a stationary but unknown random process that has to be estimated as the autonomous system moves through it and makes observations of the environment. It is assumed that the robot localizes in the environment based on a few landmarks located throughout the environment (which is the Bayesian SLAM problem but with very few parameters compared to case when the whole environment is considered to be an unknown parameter), and then maps the rest of the environment based on this estimate using frequentist methods (stochastic approximation methods). This problem formulation ensures that the estimates of the landmarks and the environmental components never get correlated and hence, each individual environmental component can be estimated completely independently of the other components. This results in the computational complexity of the formulation being linear in the environmental components as opposed to exponential for the Bayesian formulation. Further, it is shown that the mapping algorithms converge regardless of the planning policy, as long as the planning is not anticipative, i.e., does not depend on the future. the closest in philosophy to our methodology is the DenseSLAM method which too localizes the robot with respect to a sparse set of landmarks and then, maps the rest of the environemnt based on this estimate [18]. However, the mapping in this methodology is done in a Bayesian fashion (as opposed to frequentist in our method) and hence, the convergence of the map estimatescannot be guaranteed.

The rest of the paper is organized as follows. In section II, the hybrid methodology for the simultaneous mapping, localization and planning problem is proposed and the strong consistency of the mapping algorithms, and their separation from the planning algorithms, established. In section III, a computer simulation of the implementation of the methodology to a simple SLAM problem is presented in order to verify the theoretical results of the paper.

2. Mapping and Planning: Imperfect State Sensing

Let the environment consist of a set of landmarks $\Theta = \{\theta_1, \cdots, \theta_K\}$ and the stationary environmental process composed of independent components, $Q = \{q_1, q_2, \cdots, q_M\}$. Note that the set of landmarks in the environment in general would be much smaller when compared to the full environment, i.e., $K << M$. Let the state/ pose of the autonomous system be denoted by s. In the following, the hybrid SPLAM methodology is outlined and the convergence of these algorithms analyzed. In the following we shall use state and pose of the autonomous system in an interchangeable fashion, and the reader is advised that they are synonymous in the context of this paper. **We would like to stress here that we shall not consider any specific planning algorithm in this paper, only show that the mapping and planning problems can be solved independent of each other without affecting convergence.**

2.1. Mapping and Localization

The environment as mentioned above is partitioned into a set of discrete-valued deterministic landmarks, and a stationary environmental process which is characterized by the probabilities:

$$P(Q/Q') = P^*(Q) = \prod_{i=1}^{M} p^*(q_i). \tag{1}$$

The probabilities $p^*(q_i)$ are fixed but unknown and have to be estimated during the course of the algorithm. The environment observation model is as follows. Let $p(\hat{Q}/Q, s)$ denote the sensor observation model, i.e., the model represents the probability that an observation, $\hat{Q}$, is made when the environment is actually at the state Q given that the observation is made from state s. Further, it is assumed that the environmental observation model can be factored as follows:

$$p(\hat{Q}/Q, s) = \prod_{k=1}^{M} p(\hat{q}_k/q_k, s), \tag{2}$$

where $\hat{q}_k$ is the noise corrupted observation of the k^{th} environmental component, i.e., the observation of the individual environmental components are independent of each other. Thus, the observation model for the k^{th}

environmental component is represented by the stochastic matrix (since the rows of the matrix add to one), $A^{(k)}(s) = [p^k_{ij}(s)] = p(\hat{q}_k = i/q_k = j, s)$.
At any instant of time, an estimate of the location of the robotic pose and the environment needs to be formed. Let $\mathscr{F}^t = \{(\hat{\Theta}_0, \hat{Q}_0), \cdots, (\hat{\Theta}_t, \hat{Q}_t)\}$ denote the history of the algorithm till time t, where $\hat{\Theta}_t$ denotes the observation of the landmarks at instant t. The belief on the state/ pose-landmark-environment triple is then made in the following fashion:
First, a belief $b_t(s, \Theta)$ is formed on the state-landmark pair using only the observations of the landmark according to the standard Bayes filter,

$$b_t(s,\Theta) = \eta P(\hat{\Theta}_t/\Theta, s) \sum_{s'} p(s/s', u_{t-1}) b_{t-1}(s', \Theta), \tag{3}$$

where η is a suitable normalization constant. Then, a belief of the environment given the state s is formed using the stationary environment process for Q as follows:

$$b_t(Q/s) = \eta p(\hat{Q}_t/Q, s) P^*(Q), \tag{4}$$

which can be shown to be factored as

$$b_t(Q/s) = \prod_{k=1}^{M} b_t(q_k/s), \tag{5}$$

$$b_t(q_k/s) = \eta_k p(\hat{q}_{k,t}/q_k, s) p^*(q_k), \tag{6}$$

where η_k is a suitable normalizer for the belief on the k^{th} environmental component. Then the joint belief on the (s, Θ, Q) triple is formed as:

$$b_t(s, \Theta, Q) = b_t(s, \Theta) b_t(Q/s). \tag{7}$$

Due to the problem formulation, the burden of estimating the belief on the environment is on the stationary environmental process and by definition, can be maintained independent of each other. **The above comprises the Bayesian part of the hybrid methodology presented here**. Note that the current state of the environment is inferred using a Bayesian scheme given the stationary environmental probabilities $p^*(q_k)$. However, these probabilities need to be estimated since they are unknown a priori. **These are estimated using frequentist methods and comprise the frequentist part of the methodology,** and are outlined below.
Let $X_{k,t} = (b_t(s, \Theta), \hat{q}_{k,t})$ and let

$$A^{(k)}(X_{k,t}) = [\sum_{s,\Theta} p^k_{ij}(s) b_t(s, \Theta)], \tag{8}$$

$$b^{(k)}(X_{k,t}) = [\beta_i(X_{k,t})], \beta_i(X_{k,t}) = 1(\hat{q}_{k,t} = i). \tag{9}$$

There are two estimation algorithms for estimating the stationary probabilities P^*_k of the k^{th} environmental component:
Estimator E1:

$$P_k(t) = arg \min_{p \in \mathscr{P}} ||A_k(t)P - b_k(t)||, \tag{10}$$

$$A_k(t) = (1-\gamma_t)A_k(t-1) + \gamma_t A^{(k)}(X_{k,t}), \tag{11}$$

$$b_k(t) = (1-\gamma_t)b_k(t-1) + \gamma_t b^{(k)}(X_{k,t}). \tag{12}$$

Estimator E2:

$$P_k(t) = \Pi_{\mathscr{P}}\{(1-\gamma_t)P_k(t-1) + \gamma_t(b^{(k)}(X_{k,t}) - A^{(k)}(X_{k,t})P_k(t-1))\}, \tag{13}$$

where $\Pi_{\mathscr{P}}$ represents the projection into the subspace of probability vectors.
The learning rate parameters satisfy $\sum_t \gamma_t = \infty$ and $\sum_t \gamma_t^2 < \infty$, and $\mathscr{P}$ represnts the space of probability vectors in $\Re^D$. The algorithms can be heuristically understood as follows. Note that $A^{(k)}(X_{k,t})$ is the averaged observation model for environmental component q_k at time t, averaged upon the belief on the pose of the autonomous system, $b_t(s)$, at time t. Then, the matrix $A_k(t)$ in estimator E1 may be understood as the averaged observation model for the k^{th} environmental component till time t, i.e., formed by a time average of the instantaneous observation models $A^{(k)}(X_{k,t})$. The quantity $b_k(t)$ contains the relative frequencies of the observed values of the k^{th} environmental component in its various discrete states. Then, if the vector $b_k(t)$ is interpreted as a probability vector, it can be expected that in the limit as $t \to \infty$, $A_k(t)P^*_k = b_k(t)$, since P^*_k is the true environmental probability vector for the k^{th} component. This can be seen through a simple application of the law of total probability $\bar{p}(\hat{q}_k) = \sum_{q_k} \bar{p}(\hat{q}_k/q_k)) p^*(q_k)$ where $\bar{p}(\hat{q}_k = i)$ is given asymptotically by the i^{th} component of the vector $b_k(t)$ and $\bar{p}(\hat{q}_k = i/q_k = j)$ is given asymptotically by the $(i,j)^{th}$ element of the matrix $A_k(t)$. Since in general $P_k(t) = A_k(t)^{-1}b_k(t)$ need not be a probability vector for finite time t, the closest approximation of the solution, in the mean square sense, in the space of probability vectors $\mathscr{P}$ is used instead, in estimator E1. The estimator E2 may be understood as a purely incremental version of the algorithm in which we only keep track of the current estimate of the environmental components P_t instead of inferring it indirectly from $A_k(t)$ and $b_k(t)$. The asymptotic behaviour of the algorithms is identical.

2.2. Convergence Analysis

In this section, the convergence of the mapping algorithms presented in the previous subsection is established. The proofs are left out of this document because of the paucity of space. However, the interseted reader may see the whole paper at the first Author's website Suman Papers. The main results needed to prove the convergence are stated while the proofs are left out for lack of space.

A 1. *It is assumed that the estimate of the k^{th} environemntal component is updated according to estimators E1/ E2 at time t, if and only if all the diagonal elements of the current averaged observation model, $A^k(X_{k,t})$ are greater than $0.5+\varepsilon$.*

The above assumption implies that the observation (based on the current belief over the pose $b_t(s)$) is realiable in the sense that it is correct more than 50% of the time. Recall that $X_{k,t} = (b_t(s,\Theta), \hat{q}_{k,t})$, i.e., the ordered pair of the state-landmark belief and the noise corrupted observation of the k^{th} environmental component. The key here is that the sequence $X_{k,t}$ is a Markov process and it can be shown that its transition probabilities are given by

$$P(X_{k,t+1}/X_{k,t}) = p(b_{t+1}/b_t, u_t)p(\hat{q}_{k,t+1}/b_{t+1}), \quad (14)$$

where

$$p(\hat{q}_k/b) = \sum_{(s,\Theta,q_k)} p(\hat{q}_k/q_k, s)p^*(q_k)b(s,\Theta), \quad (15)$$

and the transition probabilities for the belief state $b_t(s,\Theta)$, $p(b/b',u)$, can be obtained similar to the belief state Markov Decision Process (MDP) for any Partially Observed Markov Decision Process (POMDP) [19]. The transition probabilities for the belief $b(s,\Theta)$ however are not required by the algorithms. The structure of the transition probabilities of the Markov Process above is a direct consequence of the problem formulation.

Let $\mathscr{F}^t = \{(\hat{\Theta}_0, \hat{Q}_0), \cdots, (\hat{\Theta}_t, \hat{Q}_t)\}$ denote the history of the process till the time instant t. The following result is a structural property of the algorithm:

Lemma 1. *Let $A^{(k)}(X_{k,t})$ and $b^{(k)}(X_{k,t})$ be as defined by equations (8) and (9). Then, for every environmental component, q_k, the following result holds:*

$$E[A^{(k)}(X_{k,t+1})P_k^* - b(X_{k,t+1})/\mathscr{F}^t] = 0. \quad (16)$$

Also, recall the "reliable sensor" assumption 1. This assumption is key in proving the strong consistency of the mapping/ estimation algorithms.

Proposition 1. *The estimates of the environmental probabilities $P_k(t) \to P_k^*$ for all environmental components k, with probability 1, under either of the estimation (E1/ E2) schemes, as long as the control policy is non-anticipative and the environmental components are observed infinitely often.*

The result above is valid for any non-anticipative control policy, i.e., a policy such that the control at the current instant is not dependent on the future of the algorithm. Thus, the above result also establishes an important "Separation Principle" that the mapping algorithms can be designed totally independent of the planning algorithms without affecting convergence.

3. A Simple SLAM Example

In the simulation, we have a land vehicle fitted with a radar and we want to map the unknown environment based on the observations from the radar. The terrain consists of point landmarks which reflect the radar waves and we want to map these landmarks. The position of the vehicle is unknown so we also need to localize the position of the vehicle relative to the observed landmarks. A cartesian coordinate system is selected coincident with the initial pose of the vehicle. All further calculations are done on this coordinate system. The vehicle starts at $(1,5)$ and keeps moving in circles of radius 4 units centered at $(5,5)$. We are interested in the rectangular domain $(0,0)$ to $(10,10)$. Thirteen landmarks are located in this domain (see Fig. 1(a)) out of which we use the central one to localize the vehicle and the other twelve are mapped using the stochastic approximation scheme E1.

The vehicle state is defined as $\mathbf{x}_p = [x,y,\phi]^T$, where x and y are the coordinates of the center of the rear axle with respect to some global coordinate frame and ϕ is the orientation of the vehicle in the same global frame. The global frame is fixed relative to the starting pose of the vehicle.

The dynamics of the vehicle is governed by the following equations

$$\begin{aligned} \dot{x} &= V\cos\phi \\ \dot{y} &= V\sin\phi \\ \dot{\phi} &= \frac{V}{r}, \end{aligned} \quad (17)$$

where r is the radius of the path we want the vehicle to traverse. $V = u + \omega_v, u$ is the velocity input and ω_v is gaussian noise with zero mean. The landmarks are assumed to be stationary point objects. Only one landmark was used for localization of the vehicle while the other landmarks are mapped separately using the stochastic approximation schemes outlined in the previous section.

Each observation consists of the range $r^i(k)$ and bearing $\theta^i(k)$ to the landmark being observed. It follows the model

$$r^i(k) = \sqrt{(x^i - x(k))^2 + (y^i - y(k))^2} + \omega_r(k)$$

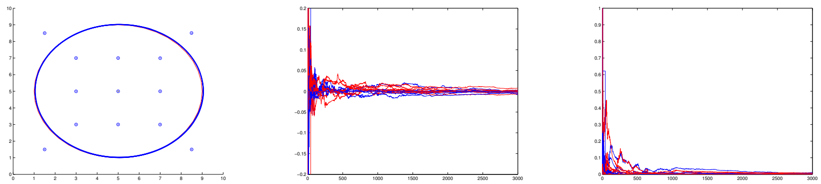

(a) True and estimated vehicle path and landmark locations (b) Mean of the error in the Landmark estimates (c) Variance of the error in the Landmark estimates

Figure 1. Performance without the Reliable Sensor Correction

$$\theta^i(k) = \arctan\frac{y^i - y(k)}{x^i - x(k)} - \phi(k) + \omega_\theta(k) \qquad (18)$$

where ω_r and ω_θ are the noise sequences associated with the range and bearing measurements, and $[x(k), y(k)]$ is the location of the radar.

The vehicle dynamics and observation model described here are both nonlinear. Hence, an EKF is employed to generate the estimates. The process model $f(.)$ and observation model $h(.)$ are obtained as given in the previous subsections.

The belief over the pose of the vehicle, $b_t(s)$, is needed for use in the mapping algorithms E1 and E2 (for calculation of $A^{(k)}(X_{k,t})$). In this case, this amounts to calculating the belief over the (x, y) co-ordinates of the autonomous vehicle. We know that $\mathbf{x}$ has Gaussian distribution with mean $\mathbf{x}(t)$ and covariance matrix $P(t)$. Any linear transformation of a gaussian vector is also Gaussian. Thus $x(1:2)$ has a guassian distribution with mean $[x(t), y(t)]$ and the covariance matrix being the top left 2x2 sub-matrix of P. Let us call this P_0. Now we can find the probability that the vehicle is in some grid, in the x, y domain, simply by integrating the multivariate normal distribution over the required domain. This gives us the belief states. The landmarks are mapped using estimator E1 from the previous section. In order to apply the inherently dicrete mapping methodology, the location of the vehicle is discretized into a 20 x 20 grid on the X-Y plane while the location of any landmark is assumed to lie in a subset of these grids which is identified based on the observation of the landmark assuming perfect data association. The reliable sensor criterion is implemented by imposing the condition that only those readings are counted for which the diagonal terms of the averaged observation model at time t, $A^{(k)}(X_{k,t})$, are greater that 50%. This helps us to obtain better convergence properties for the landmark estimates. The simulation was run for two cases. In both cases there were thirteen total landmarks and only one (fixed apriori) landmark was use to localize the vehicle and then the stochastic approximation algorithms were used to map the remaining twelve landmarks. In the first case, the mapping algorithms were implemented without imposing the reliable sensor condition (Fig.1). In the second case, the reliable sensor assumption was imposed (Fig. 2). It can be seen from Figs. 1(b), (c) and 2(b), (c) that the estimates of the landmarks indeed converge as predicted by the theoretical results. Also note that the estimates converge even without imposing the reliable sensor assumption (Figs. 1(b), (c)) though the convergence is slower in this case than when the assumption is enforced.

In summary, the fundamental difference between the methodology presented in this paper and existing SPLAM methods is in the hybrid problem formulation which has the following advantages over its Bayesian counterpart:

- Computationally much more tractable than the Bayesian formulation since the computational complexity is linear in the environmental components as compared to exponential complexity for the Bayesian problem.

- Provable strong convergence of mapping algorithms.

- Separation of Mapping and Planning Problems, in the sense that the convergence of the mapping algorithms is completely unaffected by the planning policy as long as it is non-anticipative.

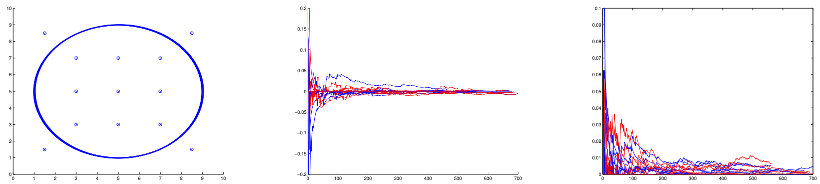

(a) True and estimated vehicle path and landmark locations (b) Mean of the error in the Landmark estimates (c) Variance of the error in the Landmark estimates

Figure 2. Performance after imposing the Reliable Sensor Correction

It has to be noted here that the implementation of the methodology as presented in this paper is far from being efficient and implementable in a real robotic SPLAM system. However, the purpose of this paper to show the basic soundness of the methodology and should be understood in this light. Our current research is focused on the efficient implementation of the hybrid methodology presented here. Also, the current methodology cannot handle the issue of data association which is also another of our current research directions.

References

[1] S. Thrun, "A probabilistic on-line mapping algorithm for teams of mobile robots," *International Journal of Robotics Research*, vol. 20, pp. 335–363, 2001.

[2] S. Thrun, D. Fox, and W.Burgard, "A probabilisitic approach to concurrent mapping and localization for mobile robots," *Mach. Learn. Aut. Robots*, vol. 31, pp. 29–53, 1998.

[3] K. Murphy, "Bayesian map learning in dynamic environments," in *Advances in Neural Information Processing Systems (NIPS)*, 1999.

[4] 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.

[5] 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.

[6] P. Moutarlier and R. Chatila, "Stochastic multisensor data fusion for mobile robot localization and environment modeling," in *Proceedings Int. Symp. Rob. Res.*, 1989.

[7] J. Leonard and H. Durrant-Whyte, *Directed Sonar Sensing for Mobile Robot Navigation*. Boston, MA: Kluwer academic, 1992.

[8] Y. Bar-Shalom and T. Fortmann, *Tracking and Data Association*. New York: Academic Press, 1988.

[9] A. Doucet, J. F. G. deFrietas, N. J. Gordon, and eds., *Sequential Monte Carlo Methods in Practice*. New York: Springer-Verlag, 2001.

[10] J. Liu and R. Chen, "Sequential monte carlo methods for dynamical systems," *Journal of the American Statistical Association*, vol. 93, pp. 1032–1044, 1998.

[11] J. Nieto, J. Guivant, E. Nebot, and S. Thrun, "Real time data association for fastslam," in *Proceedings IEEE Int. Conf. Rob. Aut.*, 2003.

[12] J. M. O'Kane, B. Tovar, P. Cheng, and S. M. LaValle, *Autonomous mobile Robots: Sensing, Control, Decison-Making and Applications*. New York, NY: Marcel-Dekker, 2006.

[13] R. A. Brooks, "Robust layered control system for a mobile robot," *IEEE J. Rob. Automat.*, vol. 2, pp. 14–23, 1986.

[14] B. J. Kuipers and Y. T. Bun, "A robot exploration and mapping startegy based on semantic hierarchy of spatial representations," *J. Robot. Autonom. Syst.*, vol. 8, pp. 47–63, 1991.

[15] T. S. Levitt and D. T. Lawton, "Qualitative navigation for mobile robots," *Artif. Intell. J.*, vol. 44, pp. 305–360, 1990.

[16] W. Burgard, D. Fox, M. Moors, R. Simmons, and S. Thrun, "Collaborative multi-robot exploration," in *Proceedings of IEEE Int. Conf. Rob. Automat.*, 2000.

[17] J. Nieto, J. Guivant, and E. Nebot, "Denseslam: Simultaneous localization and dense mapping," *The International J. Rob. Res.*, vol. 25, pp. 711–744, 2006.

[18] D. P. Bertsekas, *Dynamic Programming and Optimal Control, vols I and II*. Cambridge: Athena Scientific, 2000.