A Hierarchical Bayesian Approach to the Revisiting Problem in Mobile Robot Map Building

Dieter Fox[1], Jonathan Ko[1], Kurt Konolige[2], and Benjamin Stewart[1]

[1] University of Washington, Dept. of Computer Science & Engineering, Seattle, WA
[2] SRI International, Artificial Intelligence Center, Menlo Park, CA

Abstract. We present an application of hierarchical Bayesian estimation to robot map building. The *revisiting problem* occurs when a robot has to decide whether it is seeing a previously-built portion of a map, or is exploring new territory. This is a difficult decision problem, requiring the probability of being *outside* of the current known map. To estimate this probability, we model the structure of a "typical" environment as a hidden Markov model that generates sequences of views observed by a robot navigating through the environment. A Dirichlet prior over structural models is learned from previously explored environments. Whenever a robot explores a new environment, the posterior over the model is estimated using Dirichlet hyperparameters. Our approach is implemented and tested in the context of multi-robot map merging, a particularly difficult instance of the revisiting problem. Experiments with robot data show that the technique yields strong improvements over alternative methods.

1 Introduction

Building maps of unknown environments is one of the fundamental problems in mobile robotics. As a robot explores an unknown environment, it incrementally builds a map consisting of the locations of objects or landmarks. Typically, as it explores larger areas, its uncertainty relative to older portions of the map increases; for example, in closing a large loop. Thus, a key problem is determining whether the current position of the robot is in an unexplored area or in the already-constructed map (the *revisiting problem*). The revisiting problem for single robots is illustrated in Fig. 1(a). Shown there is a map built by a robot during exploration. The robot started in the lower right hallway and moved clockwise around the large loop. At the end, it moves down the right hallway, but due to the accumulated uncertainty in its own position, it can not determine whether it is in the same hallway as in the beginning or whether it is in a parallel hallway.

Multiple robots exploring an environment from unknown start locations face a particularly difficult instance of the revisiting problem. For coordinated exploration, the robots have to merge their maps so as to build a shared world model. Map merging requires the determination of the robots' relative location. Consider the situation shown in Fig. 1(b). Here, two robots have explored parts of the large environment shown below. In order to merge the partial maps, they have to determine whether they visited the same locations in the environment and if so, they have to determine the offset between their maps. The main difficulty of this problem lies in the first step, *i.e.* in deciding whether there is an overlap between the two maps

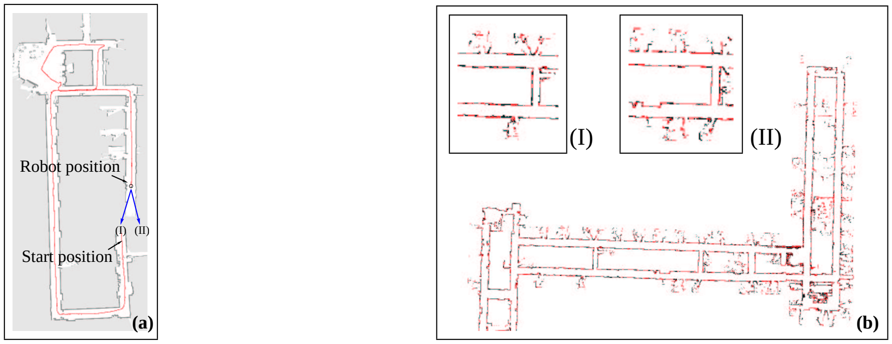

Fig. 1. (a) Loop closing: A robot explores an environment and has to decide whether it returned to the hallway it started in (I) or whether it is in a parallel hallway (II). (b) Multi-robot map merging: Two robots built the partial maps (I) and (II) and have to decide whether they explored an overlapping part of the environment, *i.e.* whether they can merge their maps.

or not. To avoid this decision problem, existing approaches to multi-robot mapping assume knowledge about the robots' relative start locations. At the minimum, these techniques require either that one robot is known to start in the map already built by the other robot [4,15,14] or that there exists an overlap between the maps [3].

If we consider the revisiting problem in a Bayesian context, then to make an informed decision, we require probabilities for two different hypotheses, one for the robot moving through area that has already been mapped, and one for the robot moving through unexplored area. To compute these probabilities, it is necessary to determine the likelihood of sensor measurements under the two hypotheses. While it is well-understood how to compute the measurement likelihood in areas already mapped by a robot, it is not clear how to compute the likelihood of sensor measurements in areas the robot has not yet explored. Existing approaches to map building implicitly determine the likelihood for "out of map" measurements under the assumption that objects are distributed uniformly, *i.e.* they assign fixed, identical likelihoods to all observations in unexplored areas [12,11,7,14]. Obviously, such approaches ignore valuable information since most environments are structured rather than randomly patched together.

The key contribution of this paper is a method for estimating the probability of the out-of-map hypothesis. In a nutshell, we construct a structural model of a typical environment; when the robot is outside the partial map, we use the model to predict what a typical view would look like, given the robot's history of observations. The current observation is then compared against the generated view to compute a likelihood. More specifically, we introduce a hierarchical Bayesian approach that captures the structure of an environment by a hidden Markov process that represents transitions between views of the environment. An offline learning process takes a set of maps and generates a Dirichlet prior over map structures. The prior is the "typical" generative map used by the robot at the start of exploration. An adaptation process refines the model distribution as the robot encounters new views of its environment.

To prove the validity of the approach, we have constructed an efficient implementation, using a particle filter that derives the likelihoods of the out-of-map hypothesis under the structural model. Views are discrete features extracted from laser range-finder scans. Experiments using a multi-robot exploration scenario show that our technique clearly outperforms alternative approaches to map merging.

This paper is organized as follows. In the next section, we will describe the Bayesian approach to learning and estimating the structure of environments. Section 3 outlines the generative model for map merging. Experiments are described in Section 4, followed by a discussion.

2 Bayesian Estimation of Map Structures

Our model of map structures is based on the idea that indoor environments consist of collections of local patches. These patches, especially the way they are connected, generate *sequences of views* observed by a robot as it moves through an environment. For example, many indoor environments consist of straight hallways, hallway crossings, and rooms. These local pieces are not patched together by pure chance, but rather according to the global structure of the environment. We use discrete, multinomial distributions to describe the connectivity between the views (patches) observed by a robot. Before we give the details of our approach, we will review some properties of Dirichlet distributions, which form the basis for estimating multinomial distributions (see also [10,9,6] for details).

2.1 Dirichlet Hyperparameters

Assume we want to estimate the parameters of a multinomial distribution $\boldsymbol{q} = \langle q_1, q_2, \ldots, q_n \rangle$ with n bins. The value of each q_i gives the probability of bin i and the parameters are constrained to sum up to 1. In Bayesian estimation, these parameters are treated as random variables and we estimate distributions over their values. Because of its convenient mathematical properties, the *Dirichlet* distribution is a standard choice for estimating multinomials. The Dirichlet distribution over multinomials with n bins is parameterized by a vector $\boldsymbol{\alpha} = \langle \alpha_1, \alpha_2, \ldots, \alpha_n \rangle$, with all components $\alpha_i > 0$. These parameters are also called *hyperparameters*, since they represent distributions over distributions (multinomials in our case). Under a Dirichlet with parameter $\boldsymbol{\alpha}$, the probability of a multinomial $\boldsymbol{q}$ is given by

$$p(\boldsymbol{q} \mid \boldsymbol{\alpha}) \;\sim\; \text{Dirichlet}(\boldsymbol{q} \mid \boldsymbol{\alpha}) \;=\; \frac{\Gamma(\sum_i \alpha_i)}{\prod_i \Gamma(\alpha_i)} \prod_i q_i^{\alpha_i - 1}, \tag{1}$$

where Γ is the gamma distribution. An important inference task in our context is to estimate posteriors over multinomials $\boldsymbol{q}$ given frequency counts $\boldsymbol{f} = \langle f_1, f_2, \ldots, f_n \rangle$. Each count f_i describes how often the i-th bin of the multinomial $\boldsymbol{q}$ was observed. Let the prior distribution over multinomials $\boldsymbol{q}$ be given by a Dirichlet with parameter $\boldsymbol{\alpha}$. Then it can be shown that the posterior is Dirichlet with the following parameters:

$$p(\boldsymbol{q} \mid \boldsymbol{f}, \boldsymbol{\alpha}) \;=\; \frac{p(\boldsymbol{f} \mid \boldsymbol{q}, \boldsymbol{\alpha})\; p(\boldsymbol{q} \mid \boldsymbol{\alpha})}{p(\boldsymbol{f} \mid \boldsymbol{\alpha})} \;=\; \text{Dirichlet}(\boldsymbol{q} \mid \boldsymbol{f} + \boldsymbol{\alpha}) \tag{2}$$

The property that a Dirichlet prior $\boldsymbol{\alpha}$ along with frequency counts $\boldsymbol{f}$ results in a Dirichlet posterior over $\boldsymbol{q}$ is also called conjugacy and is one of the key advantages of Dirichlets for estimating multinomials. Furthermore, (2) shows that the posterior is given by simply adding the observed frequency counts to the prior. Hence, the prior can be seen as initial counts in the different bins observed *before* the data $\boldsymbol{f}$. The higher the values of the α_i's, the stronger the prior, *i.e.* the more data $\boldsymbol{f}$ is needed to dominate over the prior.

Another important task is to determine the expected probability of the different bins i of the multinomial given the prior and the observed data. This value, also called the posterior predictive distribution, is computed by integrating over all possible multinomials weighted by their probability.

$$p(i \mid \boldsymbol{f}, \boldsymbol{\alpha}) \;=\; \int \text{Dirichlet}(\boldsymbol{q} \mid \boldsymbol{f} + \boldsymbol{\alpha})\; q_i \; d\boldsymbol{q} \;=\; \frac{f_i + \alpha_i}{\sum_{i'} f_{i'} + \alpha_{i'}}. \tag{3}$$

The rightmost equation follows from the properties of the Dirichlet. (3) shows that the predictive probability of bin i has an extremely convenient form, since it is proportional to the number of times i was observed plus the initial counts provided by the prior. The denominator is simply a normalizer to ensure that the probabilities sum up to one. This finalizes the review of Dirichlet distributions and in the next two sections we will describe how to estimate and learn map structures.

2.2 Inferring Map Structure

As noted above, we represent the structure of an environment by the way the local patches, or views, are observed by a robot moving through the environment. We assume that a robot can observe a finite number ν of distinctive views. The structure of an environment is captured by $\nu \times \nu$ parameters $q_{i|j}$, which describe the probability of observing view i given that the robot previously saw view j. Let $\mathbf{q}_{|j} = \langle q_{1|j}, q_{2|j}, \ldots, q_{\nu|j} \rangle$ denote the multinomial distribution over views following view j. The complete structure of an environment is thus represented by a collection of ν multinomial distributions $\mathbf{q}_{|j}$, one for each view. Since the structure of an environment is not directly observable, it has to be estimated from data collected by a robot. As a robot moves through the environment, it observes a sequence of views, which results in frequency counts $\mathbf{f}_{|j} = \langle f_{1|j}, f_{2|j}, \ldots, f_{\nu|j} \rangle$, where each $f_{i|j}$ describes how often the robot observed view i after observing view j [1]. We can use these counts to estimate the parameters of the multinomials $\mathbf{q}_{|j}$. To do so, let us first assume that the Dirichlet priors $\boldsymbol{\alpha}_j = \langle \alpha_{1_j}, \alpha_{2_j}, \ldots, \alpha_{\nu_j} \rangle$ for these distributions are known. Then, given the prior $\boldsymbol{\alpha}_j$ and the counts $\mathbf{f}_{|j}$, the posterior distribution over $\mathbf{q}_{|j}$ is given by (4), which follows directly from (2).

$$p(\mathbf{q}_{|j} \,|\, \mathbf{f}_{|j}, \boldsymbol{\alpha}_j) = \text{Dirichlet}(\mathbf{q}_{|j} \mid \mathbf{f}_{|j} + \boldsymbol{\alpha}_j) \tag{4}$$

[1] The robot actually does not observe discrete views, but rather continuous, noisy versions thereof. We determine the frequency counts $f_{i|j}$ using the views that are most likely to have generated the observations. See [1] for an approach for partially observable views.

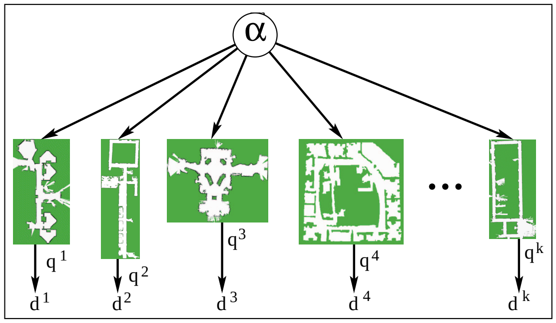

Fig. 2. Hierarchical Bayesian model: The hyperparameter $\boldsymbol{\alpha}$ represents the prior distribution over maps. The structure of each map is captured by a Dirichlet process q_l which describes how map patches are connected. The Dirichlet processes generate data d^l, sequences of views observed by a robot during exploration.

The key inference task for addressing the revisiting problem is to use the structural parameters to predict the next view observed by the robot. Let v_t denote the random variable over views at time t. Following (3), the predictive distribution for v_t given the previous view v_{t-1}, the frequency counts $\mathbf{f}_{|j}$ observed so far, and the prior $\boldsymbol{\alpha}_j$, is given by

$$p(v_t = i \mid v_{t-1} = j, \boldsymbol{\alpha}_j, \mathbf{f}_{|j}) = \frac{f_{i|j} + \alpha_{i_j}}{\sum_{i'} f_{i'|j} + \alpha_{i'_j}}. \quad (5)$$

As can be seen, the prior and the frequency counts are sufficient statistics for the predictive distribution. Furthermore, whenever a robot makes an observation in a new environment, (5) can be updated by simply incrementing the frequency count $f_{i|j}$ of the most recently observed view transition. It remains to be shown how to determine the Dirichlet prior $\boldsymbol{\alpha}$ over map structures.

2.3 Learning Prior Map Structure

Estimating a Dirichlet prior over map structures is an important component of our approach since this prior will be used by a robot when it enters an unknown environment. We estimate the prior using a hierarchical Bayesian approach [6] based on data collected in previously encountered environments. The hierarchical Bayesian model is illustrated in Fig. 2. Shown there are maps of typical indoor environments. As a robot moves through one of these environments, it observes a sequence of views d^l distributed according to the transition parameters q^l of the map structure. We assume that indoor environments are similar in the way their local patches (views) are connected. This similarity is captured by the common hyperparameter $\boldsymbol{\alpha}$, which serves as the prior distribution from which all map structures are drawn. While a full Bayesian treatment would require to use the data so as to learn a *distribution* over hyperparameters $\boldsymbol{\alpha}$, we restrict our model to the MAP estimate $\boldsymbol{\alpha}^*$:

$$\boldsymbol{\alpha}^* \;=\; \operatorname*{argmax}_{\boldsymbol{\alpha}} p(\boldsymbol{\alpha} \mid d) \;=\; \operatorname*{argmax}_{\boldsymbol{\alpha}} \frac{p(d \mid \boldsymbol{\alpha})\, p(\boldsymbol{\alpha})}{p(d)} \;\approx\; \operatorname*{argmax}_{\boldsymbol{\alpha}} p(d \mid \boldsymbol{\alpha}) \quad (6)$$

Here the rightmost term follows from a non-informative prior over the hyperparameter $\boldsymbol{\alpha}$ and the fact that $p(d)$ has no impact on the MAP estimate. The data $d = \langle d^1, \ldots, d^k \rangle$ consists of frequency counts observed in the k previously explored maps. More specifically, each d^l contains $\nu \times \nu$ counts $f^l_{i|j}$ specifying how often the robot observed a transition from view j to view i in environment l. Assuming independence between the different maps and between the Dirichlet priors for the different views, we can maximize (6) over the individual priors $\boldsymbol{\alpha}_j$ as follows:

$$p(d \mid \boldsymbol{\alpha}_j) = \prod_{l=1,\ldots,k} p(d^l \mid \boldsymbol{\alpha}_j) = \prod_{l=1,\ldots,k} \frac{\prod_i \Gamma(f^l_{i|j} + \alpha_{i_j})}{\Gamma(\sum_i f_{i|j} + \alpha_{i_j})} \frac{\Gamma(\sum_i \alpha_{i_j})}{\prod_i \Gamma(\alpha_{i_j})} \tag{7}$$

The rightmost equation follows by algebraic manipulation using properties of the Dirichlet distribution [9]. We determine the MAP $\boldsymbol{\alpha}^*$ by maximizing the log of (7) using a conjugate gradients method for each component $\boldsymbol{\alpha}_j$ (see also [9,10]).

To summarize, the structure of an environment is captured by a collection of multinomial distributions $\mathbf{q}_{|j}$ describing the sequence of views observed by a robot as it navigates through the environment. A Dirichlet prior $\boldsymbol{\alpha}$ over these structural parameters is learned from data collected in previously explored environments. As the robot moves through a new environment, it estimates the posterior over the structure of this environment. Sufficient statistics for the posterior over multinomials are given by the Dirichlet prior and the frequency counts of view transitions observed in the new environment. In the next section we will outline how this predictive model can be used in the context of multi-robot map merging (our implementation of map merging is not the main focus of this paper and details can be found in [13,8]).

3 Application to Multi-robot Map Merging

As described in Section 1, the multi-robot map merging problem is a particularly difficult instance of the revisiting problem. Imagine two robots exploring an environment from different, unknown start locations. As soon as they can communicate via wireless connection, the robots try to *co-locate*, that is they try to determine the relative offset between their maps (the robots can not see each other). To do so, one robot transmits the sensor data it collected so far and the other robot estimates the location of this robot relative to its own, partial map. The main difficulty of this task is to determine whether the paths of the robots overlap at all. If the relative offset between the maps can be established, map merging can be performed by a mapping algorithm such as [14,7,4,15].

Our approach addresses the co-location problem by estimating the location of a robot both inside and outside the partial map of the other robot. We do this by using a particle filter similar to robot localization in complete maps [5]. Particle filters represent posteriors over a robot's continuous position by sets $S_t = \{\langle x_t^{(i)}, w_t^{(i)} \rangle \mid i = 1, \ldots, N\}$ of N weighted samples distributed according to the posterior. Here each $x_t^{(i)}$ is a robot position (or state), and the $w_t^{(i)}$ are non-negative numerical factors called *importance weights*, which sum up to one. Sets at time t are generated from previous sets S_{t-1} by a sampling procedure often referred to as SISR, sequential importance sampling with re-sampling. SISR implements the recursive Bayes filter in a three stage process: First, draw states $x_{t-1}^{(i)}$ from the previous sample set with probability given by the importance weights $w_{t-1}^{(i)}$. Then draw for each such state a new state from the motion model $p(x_t \mid x_{t-1}^{(i)}, u_{t-1})$, where u_{t-1} typically describes an odometry measurement. Finally, weight these new states/samples proportional to the observation likelihood $p(z_t \mid x_t)$, which describes the likelihood of observing the sensor measurement z_t given the robot's location x_t.

As noted above, we estimate the posterior over robot locations x_t both inside and outside the partial map (see [8,13] for details). If a particle is inside the partial map, the likelihood of a measurement z_t can be computed by comparing the measurement with the measurement expected at the location in the map (identical to regular robot localization). If, on the other hand, the particle is outside the partial map, then we compute the expected observation using the structural model discussed in Section 2.2. To do so, we extract discrete views v_t from laser range-scans. These views roughly correspond to map patches such as hallways, openings, rooms, *etc.* [13]. At each iteration, the next view is predicted using the estimate of the previous view v_{t-1} and the view transition given by (5). The transition model is based on the frequency counts $\mathbf{f}_{|j}$ already observed in this environment and the priors $\boldsymbol{\alpha}_j$, which are computed from previously explored maps by maximization of (7). The likelihood weight for particles outside the map is then determined by comparing the predicted view to the view extracted from the current laser measurement.

To summarize, our approach to map merging sequentially estimates a robot's location both inside and outside the partial map built by the other robot. Particles inside the map are updated using expected measurements extracted from the map, and locations outside the map are updated using measurements predicted by the structural model. At each iteration, the parameters of the model are updated using the frequency counts of view transitions extracted from the observations.

4 Experiments

Our technique to map merging under global uncertainty was tested using data collected in the five environments shown in Fig. 2. We generated 15 partial maps from different environments and estimated the location of a robot relative to these maps [13]. The relative position was estimated from sensor logs collected in the same environments. The sensor logs were chosen randomly and some of them had no overlap with the corresponding partial map at all. For each map-trajectory pair we proceeded as follows. At each iteration of the particle filter, we determined the most likely hypothesis for the robot's location. If the probability of this hypothesis exceeded a certain threshold θ then this hypothesis was considered valid. For each threshold θ, *precision* measures the fraction of correct valid hypotheses, *i.e.* hypotheses above the threshold. Correctness is tested by comparing the position of the hypothesis to a ground truth estimate computed offline. To determine recall, we first checked at what times the robot was in the partial map. *Recall*, then, measures

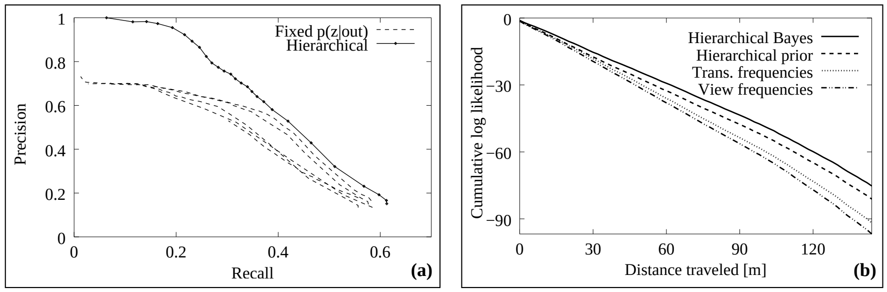

Fig. 3. (a) Precision *vs*. recall: Each point represents an average over 375 pairs of partial maps and trajectories. Each curve shows the trade-off for different thresholds θ (0.05-0.99). The dashed lines indicate results obtained with different fixed values for $p(z \mid \text{outside})$ and the solid line represents the results for our approach. (b) Predictive likelihood of different approaches averaged over 45 data sequences in three environments.

the fraction of this time for which the approach generated a correct hypothesis, *i.e.* at the correct position and with probability above the threshold θ. We compared our approach to an alternative method that uses a *fixed* likelihood $p(z|\text{outside})$ for locations outside the partial map (this corresponds to existing mapping techniques).

The precision-recall trade-offs for different thresholds θ are shown in Fig. 3(a) (each point represents a different threshold). The solid line represents the results obtained with our approach and the dashed lines are results for the fixed approach using different likelihoods $p(z|\text{outside})$ for measurements outside the maps (data points are omitted for clarity). The graphs clearly show the superior performance of our approach. It achieves 26% higher precision than the best likelihood value for the alternative method. Note that high precision values are more important than high recalls since low precision results in wrong map merging while low recall only delays the map merging decision. Note also that one cannot expect very high recall values since a robot has to be in the partial map for a certain duration before a valid hypothesis can be generated.

To evaluate the predictive quality of our approach, we used sequences of data collected in three environments. At each iteration (after approximately 2m of robot motion), we computed the likelihood of the next view in the data log given the view prediction of the structural model. The predictive quality is then determined by accumulating the logarithm of the measurement likelihoods over time. Fig. 3(b) shows the results for alternative techniques, averaged over 45 data sequences. The solid line represents the results for our approach, *i.e.* using (5) to predict the next view. The dashed line gives the results using our approach, but *without* updating the structural model, *i.e.* only the Dirichlet prior learned from other maps is used. Even though the (logarithmic) difference between these top two graphs seems small, the average likelihood of a complete sensor sequence using our adaptive approach is approximately 360 times as high as with the prior only approach. This indicates that it is important to update the structural model during exploration. The dotted line shows the result if we predict views using the frequency counts of view transitions observed in the other maps. These predictions are clearly inferior to those of the

Dirichlet prior learned with the hierarchical Bayesian approach (dashed line), which shows that our learning method improves the performance over straightforward transition frequency counting. Finally, the dashed-dotted line gives the result based on frequency counts of individual views, *i.e.* without considering transitions between views. This graph demonstrates that considering the *connectivity* of environments is superior to predicting views simply based on their frequency. These graphs are averages over different environments. We found our approach to yield much stronger improvements in more predictable environments such as the rightmost one in Fig. 2.

5 Conclusions and Future Work

In this paper, we introduced a novel approach to addressing the revisiting problem in mobile robot map building. Multi-robot map merging, a particularly difficult instance of this problem, requires the localization of one robot relative to a partial map built by another robot. The key problem in map merging without knowledge about the robots' relative locations is to get accurate estimates for the likelihoods of observations *outside* the partial map. To solve this problem, we introduce a structural model of an environment that can be used to predict the observations made by the robot. The structural model is a hidden Markov model that generates sequences of views observed by a robot when navigating through the environment. The parameters of the model are updated during exploration via Dirichlet hyperparameters. A Dirichlet prior is learned from previously encountered environments.

The structural model is integrated into a particle filter that uses samples to represent a robot's location and that updates the structural parameters as more data becomes available. Extensive experiments show that our approach clearly outperforms alternative techniques. In [8], we additionally show how to integrate this approach into a multi-robot exploration strategy. The approach enables teams of robots to efficiently explore environments from different, completely unknown start locations. The system has been shown to efficiently generate maps that consistently combine the information collected by multiple robots.

Our Bayesian approach can be readily applied to the loop closing problem in single robot mapping (see Fig. 1(a)). Here, a robot has to decide whether it came back to a previously explored location, or whether it moves through a similar, unexplored area. Especially mapping approaches based on Rao-Blackwellised particle filters [11] and topological SLAM [2] can easily incorporate our structural model.

Despite these encouraging results, this is only the first step towards using structural models of environments. For example, our current approach uses maximum likelihood estimates to update the parameters of the model. More sophisticated EM-based techniques such as [1] might yield further improvements. Other areas for improvement are better algorithms for extracting views from sensor data. Another application of our method is to improve robot exploration strategies by *predicting* partial maps into unexplored areas. Thereby, for example, a robot can actively try to close loops so as to improve map quality.

We consider hierarchical Bayesian techniques such as the one used in this paper to be an extremely promising tool for achieving more robust estimation and reasoning processes in robotics. Most existing approaches to state estimation in robotics are fixed in that they do not adapt to the environment. For example, if a map building approach is based on the assumption that the environment is rectilinear, then it will fail in environments that violate this assumption. On the other hand, not making use of the fact that most environments are rectilinear obviously discards valuable information. Using a hyperparameter that models the type of environment, a mapping approach can work reliably in different types of environments while still being able to make use of the structure underlying a specific environment.

Acknowledgments
This work has partly been supported by the NSF under CAREER grant number IIS-0093406 and by DARPA's SDR program (contract number NBCHC020073).

References

1. M.J. Beal, Z. Ghahramani, and C.E. Rasmussen. The infinite hidden Markov model. In *Advances in Neural Information Processing Systems (NIPS)*, 2002.
2. H. Choset and K. Nagatani. Topological simultaneous localization and mapping (SLAM): toward exact localization without explicit localization. *IEEE Transactions on Robotics and Automation*, 17(2), 2001.
3. G. Dedeoglu and G.S. Sukhatme. Landmark-based matching algorithm for cooperative mapping by autonomous robots. In *Proc. of the 5th International Symposium on Distributed Autonomous Robotic Systems (DARS)*, 2000.
4. J.W. Fenwick, P.M. Newman, and J.J. Leonard. Cooperative concurrent mapping and localization. In *Proc. of the IEEE International Conference on Robotics & Automation (ICRA)*, 2002.
5. D. Fox. Adapting the sample size in particle filters through KLD-sampling. *International Journal of Robotics Research (IJRR)*, 22(12), 2003.
6. A. Gelman, J. B. Carlin, H. S. Stern, and D. B. Rubin. *Bayesian Data Analysis*. Chapman and Hall/CRC, 1997.
7. J.S. Gutmann and K. Konolige. Incremental mapping of large cyclic environments. In *Proc. of the IEEE International Symposium on Computational Intelligence in Robotics and Automation (CIRA)*, 2000.
8. J. Ko, B. Stewart, D. Fox, K. Konolige, and B. Limketkai. A practical, decision-theoretic approach to multi-robot mapping and exploration. In *Proc. of the IEEE/RSJ International Conference on Intelligent Robots and Systems (IROS)*, 2003.
9. D.J.C. MacKay and L.C. Bauman Peto. A hierarchical Dirichlet language model. *Natural language engineering*, 1(3), 1995.
10. T. Minka. Estimating a Dirichlet distribution. Technical report, MIT, 2000.
11. M. Montemerlo and S. Thrun. Simultaneous localization and mapping with unknown data association using FastSLAM. In *Proc. of the IEEE International Conference on Robotics & Automation (ICRA)*, 2003.
12. J. Neira and J.D. Tardós. Data association in stochastic mapping using the joint compatibility test. *IEEE Transactions on Robotics and Automation*, 17(6), 2001.
13. B. Stewart, J. Ko, D. Fox, and K. Konolige. The revisiting problem in mobile robot map building: A hierarchical Bayesian approach. In *Proc. of the Conference on Uncertainty in Artificial Intelligence (UAI)*, 2003.
14. S. Thrun. A probabilistic online mapping algorithm for teams of mobile robots. *International Journal of Robotics Research*, 20(5), 2001.
15. S.B. Williams, G. Dissanayake, and H. Durrant-Whyte. Towards multi-vehicle simultaneous localisation and mapping. In *Proc. of the IEEE International Conference on Robotics & Automation (ICRA)*, 2002.