Bayesian Inference in the Space of Topological Maps

Ananth Ranganathan, *Member, IEEE*, Emanuele Menegatti, *Member, IEEE*, and Frank Dellaert, *Member, IEEE*

***Abstract*—While probabilistic techniques have previously been investigated extensively for performing inference over the space of metric maps, no corresponding general-purpose methods exist for topological maps. We present the concept of probabilistic topological maps (PTMs), a sample-based representation that approximates the posterior distribution over topologies, given available sensor measurements. We show that the space of topologies is equivalent to the intractably large space of set partitions on the set of available measurements. The combinatorial nature of the problem is overcome by computing an approximate, sample-based representation of the posterior. The PTM is obtained by performing Bayesian inference over the space of all possible topologies, and provides a systematic solution to the problem of perceptual aliasing in the domain of topological mapping. In this paper, we describe a general framework for modeling measurements, and the use of a Markov-chain Monte Carlo algorithm that uses specific instances of these models for odometry and appearance measurements to estimate the posterior distribution. We present experimental results that validate our technique and generate good maps when using odometry and appearance, derived from panoramic images, as sensor measurements.**

***Index Terms*—Bayesian inference, Markov-chain Monte Carlo (MCMC), mobile robots, perceptual aliasing, probability distributions, sample-based representations, topological maps.**

I. Introduction

MAPPING an unknown and uninstrumented environment is one of the foremost problems in robotics. For this purpose, both metric maps [11], [36], [37], and topological maps [4], [25], [31], [43] have been explored in depth as viable representations of the environment. In both cases, probabilistic approaches have had great success in dealing with the inherent uncertainties associated with robot sensorimotor control and perception that would otherwise make map building a very brittle process. Lately, the vast majority of probabilistic solutions to the mapping problem also solve the localization problem simultaneously, since these two problems are intimately connected. A solution to this simultaneous localization and mapping (SLAM) problem demands that the algorithm maintains beliefs over the pose of the robot as well as the map of the environment. The pose and the map are then updated, either recursively, assuming the belief about the other quantity to be fixed [36], or simultaneously [52].

The majority of the work in robot mapping deals with the construction of metric maps. Metric maps provide a fine-grained representation of the actual geometric structure of the environment. This makes navigation easy, but also introduces significant problems during their construction. Due to systematic errors in odometry, the map tends to accumulate errors over time, which makes global consistency difficult to achieve in large environments.

Topological representations, on the other hand, offer a different set of advantages that are useful in many scenarios. Topological maps attempt to capture spatial connectivity of the environment by representing it as a graph, with arcs connecting the nodes that designate significant places in the environment, such as corridor junctions and room entrances [27]. The arcs are usually annotated with navigation information.

Arguably the hardest problem in robotic mapping is the perceptual aliasing problem, which is an instance of the data-association problem, also variously known as "closing the loop" [17] or "the revisiting problem" [49]. It is the problem of determining whether sensor measurements taken at different points in time correspond to the same physical location. When a robot receives a new measurement, it has to decide whether to assign this measurement to one of the locations it has visited previously, or to a completely new location. The aliasing problem is hard, as the number of possible choices grows combinatorially. Indeed, we demonstrate below that the number of choices is the same as the number of possible partitions of a set, which grows hyperexponentially with the cardinality of the set. Previous solutions to the aliasing problem [24], [52] commit to a specific correspondence between measurements and locations at each step, so that once a wrong decision has been made, the algorithm cannot recover. Other solutions [9], [25], [26] perform well in most situations, but fail silently in environments where the problem is particularly hard, thus returning a wrong map of the environment.

The major contribution of this paper is the idea of defining a probability distribution over the space of topological maps and the use of sampling in order to obtain this distribution. The key realization here is that a distribution over this combinatorially large space can be succinctly approximated by a sample set drawn from the distribution. In this paper, we describe probabilistic topological maps (PTMs), a sample-based representation that captures the posterior distribution over all possible topological maps, given the available sensor measurements.

The intuitive reason for computing the posterior is to solve the aliasing problem for topologies in a systematic manner. The set of all possible correspondences between measurements and the physical locations from which the measurements are taken is exactly the set of all possible topologies. By inferring the posterior

Manuscript received October 25, 2004; revised April 7, 2005. This paper was recommended for publication by Associate Editor K. Yoshida and Editor S. Hutchinson upon evaluation of the reviewers' comments. This paper was presented in part at the IEEE/RSJ International Conference on Intelligent Robots and Systems, Sendai, Japan, September 2004.

A. Ranganathan and F. Dellaert are with the GVU Center, Georgia Institute of Technology, Atlanta, GA 30332-0760 USA (e-mail: ananth@cc.gatech.edu; dellaert@cc.gatech.edu).

E. Menegatti is with the Intelligent Autonomous Systems Laboratory (IAS-Lab), Department of Information Engineering, University of Padua, I-35131 Padova, Italy (e-mail: emg@dei.unipd.it).

Digital Object Identifier 10.1109/TRO.2005.861457

1552-3098/$20.00 © 2006 IEEE

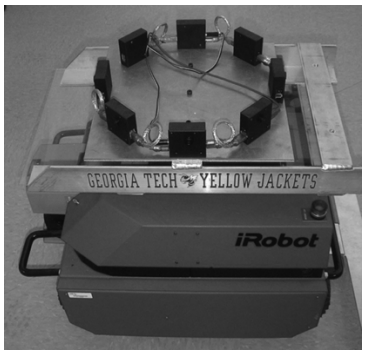

Fig. 1. Camera rig mounted on the robot used to obtain panoramic images.

on this set, whereby each topology is assigned a probability, it is possible to locate the more probable topologies without committing to a specific correspondence at any point in time, thus providing the most general solution to the aliasing problem. Even in pathological environments, where almost all current algorithms fail, our technique provides a quantification of uncertainty by pegging a probability of correctness to each topology. While sample-based estimation of the posterior over data associations has previously been discussed in computer vision [5], its use in robotic mapping to find a distribution over all possible maps is completely novel, to the best of our knowledge.

As a second major contribution, we show how to perform inference in the space of topologies, given uncertain sensor data from the robot. We provide a general theory for incorporating odometry and appearance measurements in the inference process. More specifically, we describe an algorithm that uses Markov-chain Monte Carlo (MCMC) sampling [13] to extend the highly successful Bayesian probabilistic framework to the space of topologies. To enable sampling over topologies using MCMC, each topology is encoded as a set partition over the set of landmark measurements. We then sample over the space of set partitions, using as target distribution the posterior probability of the topology given the measurements.

Another important aspect of this paper is the definition of a simple, but effective, prior distribution on the density of landmarks in the environment. We demonstrate that given this prior, the additional sensor information needed can be very scant indeed. In fact, our method is general, and can deal with any type of sensor measurement or prior knowledge. Moreover, any newly available information from the sensors can be used to augment the previous available information to improve the quality of the PTM. In our experiments, we show that even when the system produces nice maps of the environment using only odometry measurements, we can get better quality maps by using appearance data in addition to the odometry.

We provide a general-purpose appearance model in this paper, and illustrate its application using Fourier signatures [20], [32] of panoramic images. The panoramic images are obtained from a camera rig mounted on a robot, as shown in Fig. 1. An example of such a panoramic image is shown in Fig. 2. Fourier signatures, which have previously been used in the context of memory-based navigation [32] and localization using omnidirectional vision [33], are a low-dimensional representation of images using Fourier coefficients. They allow easy matching of images to determine correspondence. Further, due to the periodicity of panoramic images, Fourier signatures are rotation-invariant. This property is of prime importance when determining correspondence, since the robot may be moving in different directions when the images are obtained.

Fig. 2. Panoramic image obtained from the robot camera rig.

In subsequent sections, we provide related work in probabilistic mapping, in general, and topological mapping, in particular. Then, we define PTMs formally, and provide a theory for estimating the posterior over the space of topologies. Subsequently, we describe an implementation of the theory using MCMC sampling in topological space, followed by a section with details about the specific odometry and appearance models we used. A prior over the space of topologies is also described. Finally, in Section VII, we provide experimental validation for our technique.

II. Related Work

Our work relates to the area of probabilistic mapping and, more specifically, to topological mapping. Below we review relevant prior research in these areas.

A. Probabilistic Mapping and SLAM

Early approaches to the mapping problem (usually obtained by solving the SLAM problem) used Kalman filters and extended Kalman filters (EKFs) [3], [6], [10], [28], [48], [47] on landmarks and robot pose. Kalman-filter approaches assume that the motion model, the perceptual model (or the measurement model), and the initial state distribution are all Gaussian distributions. EKFs relax these assumptions by linearizing the motion model using a Taylor series expansion. Under these assumptions, the Kalman-filter approach can estimate the complete posterior over maps efficiently. However, the Gaussian assumptions mean that it cannot maintain multimodal distributions induced by the measurement of a landmark that looks similar to another landmark in the environment. In other words, the Kalman-filter approach is unable to cope with the correspondence problem.

A well-known algorithm that incorporates smoothing, instead of Kalman filtering, is the Lu–Milios algorithm [30], a laser-specific algorithm that performs maximum-likelihood (ML) estimation of the correspondence. It iterates over a map estimation and a data-association phase that enable it to recover from wrong correspondences in the presence of small errors. However, the algorithm encounters limitations when faced with large pose errors, and fails in large environments.

Rao–Blackwellized particle filters (RBPFs) [38], of which the FastSLAM algorithm [36] is a specific implementation, are also theoretically capable of maintaining the complete posterior distribution over maps under assumption of Gaussian measurements. This is possible since each sample in the RBPF can represent a different data-association decision [35]. However, in practice, the dimensionality of the trajectory space is too large

to be adequately represented in this approach, and often the ground-truth trajectory along with the correct data association will be missed altogether. This problem is a fundamental shortcoming of the importance sampling scheme used in the RBPF, and cannot be dealt with satisfactorily except by an exponential increase in the number of samples, which is intractable. One solution to this problem, which involves keeping all possible data associations in a tree and searching through them at each time step, is given in [18]. Another problem with RBPFs is their sensitivity to odometry drift over time. Recent work by Haehnel *et al.* [17] tries to overcome the odometry drift by correcting for it through scan matching, but this only alleviates the problem without solving it completely.

Yet another approach to SLAM that has been successful is the use of the expectation-maximization (EM) algorithm to solve the correspondence problem in mapping [2], [52]. The algorithm iterates between finding the most likely robot pose and the most likely map. EM-based algorithms do not compute the complete posterior over maps, but instead perform hill climbing to find the most likely map. Such algorithms make multiple passes over sensor data, which makes them extremely slow and unfit for online, incremental computation. In addition, EM cannot overcome local minima, resulting in incorrect data associations. Other approaches exist that report loop closures and redistribute the error over the trajectory [16], [39], [49], [51], but these decisions are again irrevocable, and hence, mistakes cannot be corrected.

Recent work by Duckett [7] on the SLAM problem is similar to our own, in the sense that he too deals with the space of possible maps. The SLAM problem is presented as a global optimization problem, and metric maps are coded as chromosomes for use in a genetic algorithm. The genetic algorithm searches over the space of maps (or chromosomes) and finds the most likely map using a fitness function. This approach differs from ours by computing the ML solution, as opposed to the complete posterior, as we do. In consequence, it suffers from brittleness similar to other techniques described previously.

B. Topological Maps

Maintaining the posterior distribution over the space of topologies results in a systematic and robust solution to the aliasing problem that plagues robot mapping. Though probabilistic methods have been used in conjunction with topological maps before, none exist that are capable of dealing with the inference of the posterior distribution over the space of topologies. A recent approach gives an algorithm to build a tree of all possible topological maps that conform to the measurements, but in a nonprobabilistic manner [43], [44]. Dudek *et al.* [8] have also given a technique that maintains multiple hypotheses regarding the topological structure of the environment in the form of an exploration tree. Most instances of previous work extant in the literature that incorporate uncertainty in topological map representations do not deal with general topological maps, but with the use of Markov decision processes to learn a policy that the robot follows to navigate the environment.

Simmons and Koenig [46] model the environment using a partially observable Markov decision process (POMDP) in which observations are used to update belief states. Another approach, that is closer to the one presented here, in the sense of maintaining a multihypothesis space over correspondences, is given by Tomatis *et al.* [55] and also uses POMDPs to solve the correspondence problem. However, while in their case, a multihypothesis space is maintained, it is used only to detect the points where the probability mass splits into two. Also, like a lot of others, this work uses specific qualities of the indoor environment, such as doors and corridor junctions, and hence, is not generally applicable to any environment. Shatkay and Kaelbling [45] use the Baum–Welch algorithm, a variant of the EM algorithm used in the context of hidden Markov models (HMMs), to solve the aliasing problem for topological mapping. Other examples of HMM-based work include [23], [15], and [1], where a second-order HMM is used to model the environment.

Lisien *et al.* [29] have provided a method that combines locally estimated feature-based maps with a global topological map. Data association for the local maps is performed using a simple heuristic, wherein each measurement is associated with the existing landmark having the minimum distance to the measured location. A new landmark is created if this distance is above a threshold. The set of local maps is then combined using an "edge-map" association, i.e., the individual landmarks are aligned and the edges compared. This technique, while suitable for mapping environments where the landmark locations are sufficiently dissimilar, is not robust in environments with large or multiple loops.

Many topological approaches to mapping, related to our work only in the sense that they form a significant part of the topological mapping literature, include robot *control* to help solve the correspondence problem. This is achieved by maneuvering the robot to the exact spot it was in when visiting the location previously, so that correspondence becomes easier to compute. Examples of this approach include Choset's generalized Voronoi graphs [4] and Kuipers' spatial semantic hierarchy [27]. Other approaches that involve behavior-based control for exploration-based topological mapping are also fairly common. Mataric [31] uses boundary-following and goal-directed navigation behaviors in combination with qualitative landmark identification to find a topological map of the environment. A complete behavior-based learning system based on the spatial semantic hierarchy that learns at many levels, starting from low-level sensorimotor control to topological and metric maps, is described in [41]. Yamauchi *et al.* [56], [57] use a reactive controller in conjunction with an adaptive place network that detects and identifies special places in the environment. These locations are subsequently placed in a network denoting spatial adjacency.

Some other approaches use a nonprobabilistic approach to the correspondence problem by applying a clustering algorithm to the measurements to identify distinctive places, an instance being [26]. Finally, SLAM algorithms used to generate metric maps have also been applied to generating integrated metric and topological maps with some success. For instance, Thrun *et al.* [53] use the EM algorithm to solve the correspondence problem while building a topological map. The computed correspondence is subsequently used in constructing a metric map. By contrast, Thrun [50] first computes a metric map using value

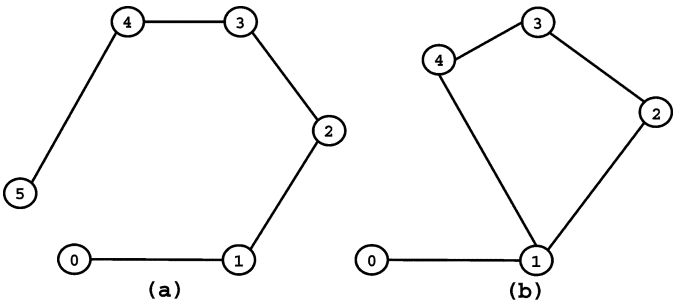

Fig. 3. Two topologies with six observations, each corresponding to set partitions (a) with six landmarks (0, 1, 2, 3, 4, 5), and (b) with five landmarks (0, 1, 5, 2, 3, 4), where the second and sixth measurement are from the same landmark.

iteration, and uses thresholding and Voronoi diagrams to extract the topology from this map.

III. Probabilistic Topological Maps

A PTM is a sample-based representation that approximates the posterior distribution $P(T|Z)$ over topologies T, given observations Z. While the space of possible maps is combinatorial, a probability density over this space can be approximated by drawing a sample of possible maps from the distribution. Using the samples, it is possible to construct a histogram on the support of this sample set.

We do not consider the issue of landmark detection in this paper. Instead, we assume the availability of a "landmark detector" that simply detects a landmark when the robot is near (or on) a landmark. Subsequently, odometry and appearance measurements from the landmark location are stored, the appearance measurements being in the form of images. The odometry can be said to measure the landmark location, while the images measure the landmark appearance. No knowledge of the correspondence between landmark measurements and the actual landmarks is given to the robot; indeed, that is exactly the topology that we seek. The problem then is to compute the discrete posterior probability distribution $P(T|Z)$ over the space of topologies.

Our technique exploits the equivalence between topologies of an environment and set partitions of landmark measurements, which group the measurements into a set of equivalence classes. When all the measurements of the same landmark are grouped together, this naturally defines a partition on the set of measurements. It can be seen that a topology is nothing but the assignment of measurements to sets in the partition, resulting in the above-mentioned isomorphism between topologies and set partitions. An example of the encoding of topologies as set partitions is shown in Fig. 3.

We begin our consideration by assuming that the robot observes N "special places" or landmarks during a run, not all of them necessarily distinct. The number of distinct landmarks in the environment, which is unknown, is denoted by M. Formally, for the N element measurement set $Z = \{Z_i | 1 \leq i \leq N\}$, a partition T can be represented as $T = \{S_j \mid j \in [1, M]\}$, where each S_j is a set of measurements such that $S_{j1} \cap S_{j2} = \phi \; \forall j1, j2 \in [1, M]$, $j1 \neq j2$, $\bigcup_{j=1}^{M} S_j = Z$, and $M \leq N$ is the number of sets in the partition. In the context of topological mapping, all members of the set S_j represent landmark observations of the jth landmark. The cardinality of the set of all possible topologies is identical to the number of set partitions of the observation N set. This number is called the Bell number b_N [40], defined as $b_N = 1/e \sum_{k=0}^{\infty} k^N/k!$, and grows hyperexponentially with N, for example $b_2 = 2$, $b_3 = 5$, but $b_{15} = 1\,382\,958\,545$. The combinatorial nature of this space makes exhaustive evaluation impossible for all but trivial environments.

IV. A General Framework for Inferring PTMs

The aim of inference in the space of topologies is to obtain the posterior probability distribution on topologies $P(T|Z)$. All inference procedures that compute sample-based representations of distributions require that evaluation of the sampled distribution be possible. In this section, we describe the general theory for evaluating the posterior at any given topology.

Using Bayes' Law on the posterior $P(T|Z)$, we obtain

$$P(T|Z) \propto P(Z|T)P(T) \tag{1}$$

where $P(T)$ is a prior on topologies and $P(Z|T)$ is the observation likelihood.

In this paper, we assume that the only observations we possess are odometry and appearance. Note that this is not a limitation of the framework, and other sensor measurements, such as laser range scans, can easily be taken into consideration. We factor the set Z as $Z = \{O, A\}$, where O and A correspond to the set of odometry and appearance measurements, respectively. This allows us to rewrite (1) as

$$\begin{aligned} P(T \mid O, A) &= kP(O, A|T)P(T) \\ &= kP(O|T)P(A|T)P(T) \end{aligned} \tag{2}$$

where k is the normalization constant, and we have assumed that the appearance and odometry are conditionally independent, given the topology. We discuss evaluation of the appearance likelihood $P(A|T)$, odometry likelihood $P(O|T)$, and the prior on topologies $P(T)$ in the following sections.

A. Evaluating the Odometry Likelihood

It is not possible to evaluate the odometry likelihood $P(O|T)$ without knowledge of the landmark locations. However, since we do not require the landmark locations when inferring topologies, we integrate over the set of landmark locations X and calculate the marginal distribution $P(O|T)$

$$P(O|T) = \int_X P(O|X, T)P(X|T) \tag{3}$$

where $P(O|X, T)$ is the measurement model, a probability density on O given X and T, and $P(X|T)$ is a prior over landmark locations. Note that (3) makes no assumptions about the actual form of X, and hence, is completely general. Evaluation of the odometry likelihood using (3) requires the specification of a prior distribution $P(X|T)$ over landmark locations in the environment, and a measurement model $P(O|X, T)$ for the odometry, given the landmark locations.

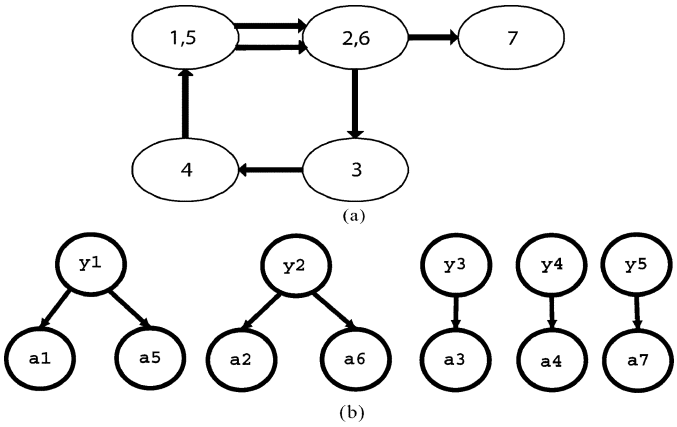

Fig. 4. The Bayesian network (b) that encodes the independence assumptions for the appearance measurements in the topology (a), given the true appearance $Y = \{y_1, \ldots, y_5\}$ at all the landmark locations. The measurements corresponding to different landmarks are independent.

B. Evaluating the Appearance Likelihood

Similar to the case of the odometry likelihood, estimation of the appearance likelihood $P(A \mid T)$, where $A = \{a_i | 1 \leq i \leq N\}$ is the set of appearance measurements, is performed by introducing the hidden parameter $Y = \{y_j | 1 \leq j \leq M\}$. This hidden parameter denotes the "*true appearance*" corresponding to each landmark in the topology. As we do not need to compute Y when inferring topologies, we marginalize over it so that

$$P(A \mid T) = \int_Y P(A \mid Y, T) P(Y \mid T) \tag{4}$$

where $P(A|Y,T)$ is the measurement model and $P(Y \mid T)$ is the prior on the appearance. We assume that the appearance of a landmark is independent of all other landmarks, so that each y_j is independent of all other $y_{j'}$. The prior $P(Y \mid T)$ can thus be factored into a product of priors on the individual y_j

$$P(Y \mid T) = \prod_{j=1}^{M} P(y_j). \tag{5}$$

The topology T introduces a partition on the set of appearance measurements by determining which "true appearance" y_j each measurement a_i actually measures, i.e., the partition encodes the correspondence between the set A and the set Y. Also, given Y, the likelihood of the appearance can be factored into a product of likelihoods of the individual appearance instances. This is illustrated using an example topology in Fig. 4, where the Bayesian network encodes the independence assumptions in the appearance measurements. Hence, denoting the jth set in the partition as S_j, we rewrite $P(A \mid Y, T)$ as

$$P(A \mid Y, T) = \prod_{j=1}^{M} \prod_{a_i \in S_j} P(a_i \mid y_j) \tag{6}$$

where the dependence on T is subsumed in the partition. Combining (4), (5), and (6), we get the expression for the appearance likelihood as

$$P(A \mid T) = \prod_{j=1}^{M} \int_{y_j} P(y_j) \prod_{a_i \in S_j} P(a_i \mid y_j). \tag{7}$$

In the above equation, $P(y_j)$ is a prior on appearance in the environment, and $P(a_i \mid y_j)$ is the appearance measurement model. Evaluation of the appearance likelihood requires the specification of these two quantities.

C. Prior on Topologies

The prior on topologies $P(T)$, required to evaluate (2), assigns a probability to topology T based on the number of distinct landmarks in T and the total number of measurements. The prior is obtained through the use of the classical occupancy distribution [21]. In the interest of continuity, the derivation of the prior is deferred to Appendix A. We simply state the expression for the prior, given the total number of landmarks in the environment L (including those not visited by the robot)

$$P(T|L) = k \frac{L^{-N} \times L!}{(L-M)!} \tag{8}$$

where N is the number of measurements, M is the number of distinct landmarks in the topology T, and k is a normalization constant. This prior distribution assigns equal probability to all topologies containing the same number of landmarks.

Note that the total number of landmarks in the environment, L, is not known. Hence, we assume a Poisson prior on L, giving $P(L|\lambda) = \lambda^L e^{-\lambda}/L!$, and marginalize over L to get the actual prior on topologies

$$\begin{aligned} P(T) &= \sum_L P(T|L) P(L|\lambda) \\ &\propto e^{-\lambda} \sum_{L=M}^{\infty} \frac{L^{-N} \times \lambda^L}{(L-M)!} \end{aligned} \tag{9}$$

where λ is the Poisson parameter, and the summation replaces the integral as the Poisson distribution is discrete. In practice, the prior on L is a truncated Poisson distribution, since the summation in (9) is only evaluated for a finite number of terms.

V. Inferring Probabilistic Topological Maps Using MCMC

The previous section provided a general theory for inferring the posterior over topologies using odometry and appearance information. We now present a concrete implementation of the theory that uses the Metropolis–Hastings algorithm [19], a very general MCMC method, for performing the inference. Fig. 5 depicts an example of the discrete posterior over topologies obtained using our MCMC-based technique. All MCMC methods work by running a Markov chain over the state space, with the property that the chain ultimately converges to the target distribution of our interest. Once the chain has converged, subsequent states visited by the chain are considered to be samples from the target distribution. The Markov chain itself is generated using a proposal distribution that is used to propose the next state in the chain, a move in state space, possibly by conditioning on the current state. The Metropolis–Hastings algorithm provides a technique whereby the Markov chain can converge to the target distribution using any arbitrary proposal distribution, the only important restriction being that the chain be capable of reaching all the states in the state space.

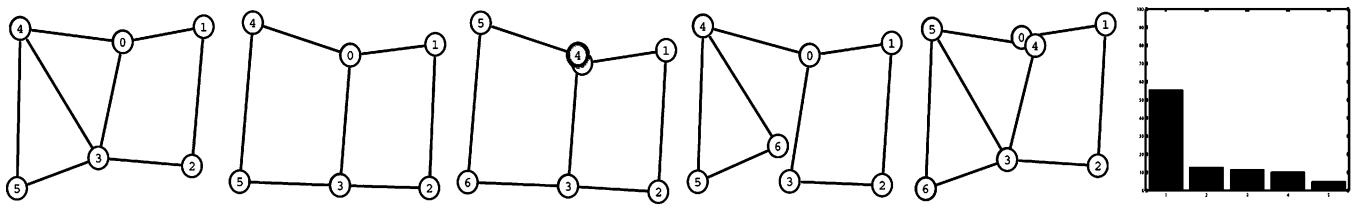

Fig. 5. Example of a PTM giving the most probable topologies in the posterior distribution obtained using MCMC sampling. The histogram gives the probability of each topology.

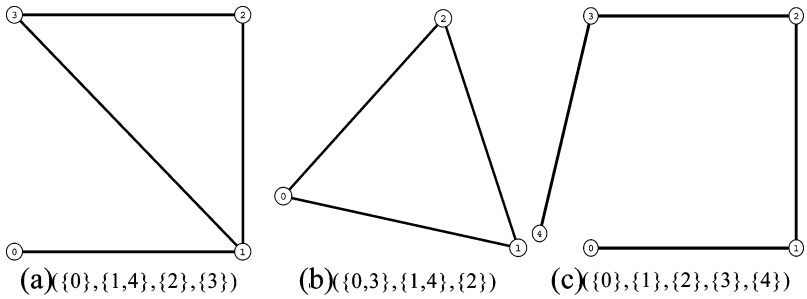

Fig. 6. Illustration of the proposal. Given a topology (a) corresponding to the set partition with $N = 5$, $M = 4$, the proposal distribution can (b) perform a merge step to propose a topology with a smaller number of landmarks corresponding to a set partition with $N = 5$, $M = 3$, or (c) perform a split step to propose a topology with a greater number of landmarks corresponding to a set partition with $N = M = 5$, or repropose the same topology.

Algorithm 1: The Metropolis–Hastings algorithm

1) Start with a valid initial topology T_t, then iterate once for each desired sample.
2) Propose a new topology T'_t using the *proposal distribution* $Q(T'_t; T_t)$.
3) Calculate the *acceptance ratio*

$$a = \frac{P(T'_t|Z^t)}{P(T_t|Z^t)} \frac{Q(T_t; T'_t)}{Q(T'_t; T_t)} \tag{10}$$

where Z^t is the set of measurements observed up to and including time t.
4) With probability $p = \min(1, a)$, accept T'_t and set $T_t \leftarrow T'_t$. If rejected, we keep the state unchanged (i.e., return T_t as a sample).

The pseudocode to generate a sequence of samples from the posterior distribution $P(T|Z)$ over topologies T using the Metropolis–Hastings algorithm is shown in **Algorithm 1** (adapted from [13]). In this case, the state space is the space of all set partitions, where each set partition represents a different topology of the environment. Intuitively, the algorithm samples from the desired probability distribution $P(T|Z)$ by rejecting a fraction of the moves generated by a proposal distribution $Q(T'_t; T_t)$, where T_t is the current state and T'_t is the proposed state. The fraction of moves rejected is governed by the acceptance ratio a given by (10), which is where most of the computation takes place. Computing the acceptance ratio, and hence, sampling using MCMC, requires the design of a proposal density and evaluation of the target density, the details of which are discussed below.

We use a simple split-merge proposal distribution that operates by proposing one of two moves, a split or a merge, with equal probability at each step. Given that the current sample topology has M distinct landmarks, the next sample is obtained by splitting a set, to obtain a topology with $M+1$ landmarks, or merging two sets, to obtain a topology with $M-1$ landmarks. The proposal is illustrated in Fig. 6 for a trivial environment. If the chosen move is not possible, the current topology is reproposed. An example of an impossible move is a merge move on a topology containing only one landmark.

The *merge move* merges two randomly selected sets in the partition to produce a new partition with one less set than before. The probability of a merge is simply $1/N_M$, where N_M is the number of possible merges and is equal to the binomial coefficient $\binom{M}{2,(M>1)}$.

The *split move* splits a randomly selected set in the partition to produce a new partition with one more set than before. To calculate the probability of a split move, let N_S be the number of nonsingleton sets in the partition. Clearly, N_S is the number of sets in the partition that can be split. Out of these N_S sets, we pick a random set R to split. The number of possible ways to split R into two subsets is given by $\left\{ {|R| \atop 2} \right\}$, where $\left\{ {n \atop m} \right\}$ denotes the Stirling number of the second kind that gives the number of possible ways to split a set of size n into m subsets, and is defined recursively as

$$\left\{ {n \atop m} \right\} \triangleq \left\{ {n-1 \atop m-1} \right\} + m \left\{ {n-1 \atop m} \right\}$$

[40]. Combining the probability of selecting R and the probability of splitting it, we obtain the probability of the split move as

$$p_{\text{split}} = \left(N_S \left\{ {|R| \atop 2} \right\} \right)^{-1}.$$

Algorithm 2: The Proposal Distribution

1) Select a merge or a split with probability 0.5.
2) **Merge move:**
   a) if T contains only one set, repropose $T' = T$, hence $r = 1$
   b) otherwise, select two sets at random, say R and S.

2.1) $T' = (T - \{R\} - \{S\}) \cup \{R \cup S\}$ and $Q(T \to T') = 1/N_M$

2.2) $Q(T' \to T)$ is obtained from the reverse case 3.2), hence $r = N_M \left(N_S \left\{ {|R \bigcup S| \atop 2} \right\} \right)^{-1}$, where N_S is the number of possible splits in T'.

3) **Split move:**

a) if T contains only singleton sets, repropose $T' = T$, hence $r = 1$

b) otherwise, select a nonsingleton set U at random from T and split it into two sets R and S.

3.1) $T' = (T - \{U\}) \cup \{R, S\}$ and $Q(T \to T') = \left(N_S \left\{ {|U| \atop 2} \right\} \right)^{-1}$

3.2) $Q(T' \to T)$ is obtained from the reverse case 2.2), hence $r = N_M^{-1} N_S \left\{ {|U| \atop 2} \right\}$, where N_M is the number of possible merges in T'.

The proposal distribution is summarized in pseudocode format in **Algorithm 2**, where Q is the proposal distribution and $r = q(T' \to T)/q(T \to T')$ is the proposal ratio, a part of the acceptance ratio in **Algorithm 1**. Note that this proposal does not incorporate any domain knowledge, but uses only the combinatorial properties of set partitions to propose random moves.

In addition to proposing new moves in the space of topologies, we also need to evaluate the posterior probability $P(T|Z)$. This is done as described in Section IV. The specification of the measurement models and the details of evaluating the posterior probability using these models are given in the following section.

VI. Evaluating the Posterior Distribution

We evaluate the posterior distribution, which is also the MCMC target distribution, using the factored Bayes rule (2). It is important to note that we do not need to calculate the normalization constant in (2), since the Metropolis–Hastings algorithm requires only a ratio of the target distribution evaluated at two points, wherein the normalization constant cancels out. The odometry and appearance measurement models required to evaluate (2) are described below.

A. Evaluating the Odometry Likelihood

Evaluation of the odometry likelihood is performed using (3)

$$P(O|T) = \int_X P(O|X,T) P(X|T)$$

under the assumption, common in robotics literature, that landmark locations and odometry measurements have the 2-D form $X = \{l_i = (x_i, y_i) | 1 \le i \le N\}$ and $O = \{o_k = (x_k, y_k, \theta_k) | 1 \le k \le N-1\}$, respectively. This requires the definition of a prior on the distribution of the landmark locations X conditioned on the topology T, $P(X|T)$.

We use a simple prior on landmarks that encodes our assumption that landmarks do not exist close together in the environment. If the topology T places two distinct landmarks l_{i1} and l_{i2} within a distance d of each other, the negative log-likelihood corresponding to the two landmarks is given by the penalty function

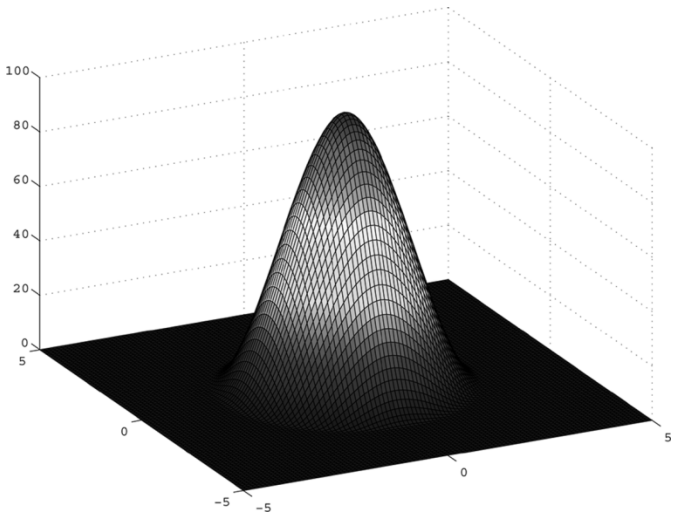

Fig. 7. Cubic penalty function (in this case, with a threshold distance of 3 m) used in the prior over landmark density.

$$L(l_{i1}, l_{i2}; T) = L(l_{i2}, l_{i1}; T) = \begin{cases} f(d), & d < D \\ 0, & d \ge D \end{cases} \tag{11}$$

where d is the Euclidean distance between l_{i1} and l_{i2}, D is a threshold value, called the "penalty radius," and we define $f(d)$ to be a cubic function, as shown in Fig. 7. The cubic function is defined using two parameters, the penalty radius D at which the function becomes zero, and the maximum value of the function at the origin. The total probability $P(X|T)$ of landmark locations X given topology T is then calculated as

$$-\log P(X|T) = \sum_{\substack{1 \le i1 < i2 \le N \\ l_{i1} \notin S(l_{i2})}} L(l_{i1}, l_{i2}) \tag{12}$$

where $S(l_{i2})$ denotes the set containing l_{i2}.

The odometry likelihood function $P(O|X,T)$ in (3) encodes the deviation between the measured odometry and the odometry predicted by the topology and the landmark locations. Intuitively, the topology T constrains some measurements as being from the same location, even though the odometry may put these locations far apart. The likelihood function accounts for the two types of errors, those from distorting the odometry and those from not conforming to the topology T. Hence, the log-likelihood for the odometry can be written as

$$-\log P(O|X,T) = \left(\frac{\|X - X_O\|}{\sigma_O} \right)^2 + \sum_{S \in T} \sum_{i1, i2 \in S} \left(\frac{l_{i1} - l_{i2}}{\sigma_T} \right)^2 \tag{13}$$

where S is a set in the partition corresponding to T, σ_O and σ_T are standard deviations explained below, and X_o is the set of landmark locations obtained from the odometry measurements. The first term on the right-hand side of (13) corresponds to the error from the odometry distortion, while the second term corresponds to the topology constraints. The standard deviations for the odometry and topology constraints, σ_O and σ_T, respectively, encode the amount of error that we are willing to tolerate in each of these quantities.

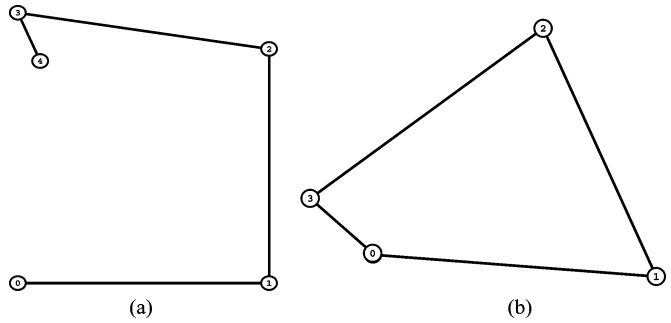

Fig. 8. Illustration of optimization of the odometry likelihood. The observed odometry in (a) is transformed to the one in (b) because the topology used in this case, ($\{0, 4\}$, $\{1\}$, $\{2\}$, $\{3\}$), tries to place the first and last landmarks at the same physical location.

A simple example illustrating the constraints is given in Fig. 8. In this example, the topology constrains X_o and X_4 (the first and last landmarks) to the same location, causing a distortion in the odometry. This results in the topology in Fig. 8(b).

1) Numerical Evaluation of the Odometry Likelihood: In some cases, it may be possible to evaluate the integral in (3) analytically, using the functional form of the log-likelihood given in (13) and (12). If closed-form evaluation is not possible, it may still be possible to use an analytical approximation technique, such as Laplace's method [54], to evaluate (3).

However, in general, it is not possible to use any form of analytical evaluation to compute (3). Instead, we employ a Monte Carlo approximation, using importance sampling [12] to approximate the integrand $P(O|X,T)P(X|T)$. Importance sampling works by generating samples from a proposal distribution that is easy to sample from. Each sample is then weighted by the ratio of the target distribution to the proposal distribution evaluated at the sample location. The Monte Carlo approximation is subsequently performed by summing the weighted samples. The primary condition on the proposal distribution is that it should be nonzero at all locations where the target distribution is nonzero. In addition, importance sampling is efficient if the proposal distribution is a close approximation to the target distribution.

In our case, the importance sampling proposal distribution is obtained from the odometry log-likelihood (13). This function is a lower bound on the log of the integrand, $\log(P(O|X,T)P(X|T))$, since the prior term given by (12) is never negative. Consequently, (13) can be used to obtain a valid importance sampling distribution. We employ Laplace's method to obtain a multivariate Gaussian distribution from $-\log P(O|X,T)$, which is used as the proposal distribution. This is achieved by computing the ML path $X^\star$ through a nonlinear optimization of $-\log P(O|X,T)$, and creating a local Gaussian approximation $Q(X|O,T)$ around $X^\star$

$$X^\star = \underset{X}{\operatorname{argmax}} \left(-\log P(O|X,T)\right)$$

$$Q(X \mid O,T) = \frac{1}{\sqrt{|2\pi\Sigma|}} e^{-(1/2)(X-X^\star)^T \Sigma^{-1}(X-X^\star)}$$

where Σ is the covariance matrix relating to the curvature of $\psi(X)$ around X^*. The distribution $Q(X|O,T)$ is then used as the proposal distribution for the importance sampler.

In practice, we use the Levenberg–Marquardt algorithm in conjunction with a sparse QR solver to perform the optimization described above. The Levenberg–Marquardt algorithm requires the derivative of the objective function that is being minimized, in this case, the function $-\log P(O|X,T)$ in (13). To compute the (sparse) Jacobian H given by $H = \partial\psi(X)/\partial X$, we use an automatic differentiation (AD) framework. AD is a technique for augmenting computer programs with derivative computations. It exploits the fact that by applying the chain rule of differential calculus repeatedly to elementary operations, derivatives of arbitrary order can be computed automatically and accurately to working precision. See [14] for more details.

The odometry likelihood given by (3) is now evaluated using the Monte Carlo approximation

$$\int_X P(O|X,T)P(X|T) \approx \frac{1}{N}\sum_{i=1}^{N} \frac{P(O|X^{(i)},T)P(X^{(i)}|T)}{Q(X^{(i)}|O,T)} \tag{14}$$

where the $X^{(i)}$ are samples obtained from the Gaussian proposal distribution $Q(X|O,T)$ and N is the number of samples.

B. Evaluating the Appearance Likelihood

Fourier signatures, which we use as appearance measurements, are computed by calculating the 1-D Fourier transform of each row of the panoramic image, and storing only the few coefficients corresponding to the lower spatial frequencies [33]. While more popular dimensionality reduction techniques, such as principal component analysis (PCA) [22] exist, the drawback of such systems is the need to further preprocess the measurement images in order to obtain rotational invariance. In contrast, the magnitudes of Fourier coefficients in a Fourier signature are rotation-invariant, since panoramic images are periodic. Hence, a Fourier signature yields a low-dimensional, rotation-invariant representation of the image. We use images obtained from an eight-camera rig mounted on a robot to produce panoramic images. The eight images obtained at each point in time are stitched together automatically to form a 360° view of the environment.

In our case, Fourier signatures are calculated using a modification of the procedure given in [32]. First, a single row image obtained by averaging the rows of the input image is calculated, and subsequently, the 1-D Fourier transform of this image is performed. This gives us the Fourier signature of the image. It is to be noted that Fourier signatures do not comprise a robust source of measurements, since the measurements contain many false positives, in the sense that images from distinct physical locations often yield similar Fourier signatures. This is due to perceptual aliasing and the extreme compression of the Fourier signature. However, they have the advantage of being simple to compute and model. Moreover, in conjunction with odometry, they still produce good results, as we demonstrate in Section VII.

Evaluation of the appearance likelihood is performed using (7). However, in this case, each appearance measurement a_i is a Fourier signature vector given as $a_i = \{a_{i1}, a_{i2}, \ldots, a_{iK}\}$, where a_{ik} is the kth Fourier component in the Fourier signature. Also, we assume a similar vector form for the hidden appearance

variables y_i, so that $y_i = \{y_{i1}, y_{i2}, \ldots, y_{iK}\}$. We can then write (7) as

$$P(A|T) = \prod_{j=1}^{M} \int_{y_j} P(y_{j1}, \ldots, y_{jK}) \times \prod_{a_i \in S_j} P(a_{i1}, \ldots, a_{iK} \mid y_{j1}, \ldots, y_{jK}). \quad (15)$$

Clearly, the various frequency components of the Fourier signature are independent, given the corresponding appearance variable, and hence, can be factored, as can be the prior over the hidden appearance variables. Consequently, we modify (15) to get the expression for the appearance likelihood as

$$P(A \mid T) = \prod_{j=1}^{M} \prod_{k=1}^{K} \int_{y_{jk}} P(y_{jk}) \prod_{a_i \in S_j} P(a_{ik} \mid y_{jk}). \quad (16)$$

We assume the measurement noise in the Fourier signatures to be Gaussian distributed, so that the model for appearance instance a_{ik}, belonging to the jth set S_j, is also a Gaussian centered around the "true appearance" y_{jk} with variance σ_{jk}^2. Since we do not know either of these parameters, we further model them hierarchically in a proper Bayesian manner. Hierarchical priors are placed on σ_{jk}^2 and y_{jk}: the prior on σ_{jk}^2 being an inverse gamma distribution, while the prior on y_{jk} is taken to be a Gaussian distribution with mean μ and variance σ_{jk}^2/κ. This particular choice of priors also allows the integration in (16) to be performed analytically. The appearance model can then be summarized as

$$\begin{aligned} a_{ik} &\sim \mathcal{N}(y_{jk}, \sigma_{jk}^2) \quad \text{where } a_i \in S_j \\ y_{jk} &\sim \mathcal{N}\left(\mu, \frac{\sigma_{jk}^2}{\kappa}\right) \\ \sigma_{jk}^2 &\sim IG(\alpha_k, \beta_k) \end{aligned} \quad (17)$$

where IG denotes the inverse gamma distribution. Note that while the value of κ is generally chosen so that the prior on y_{jk} is vague, we usually have some extra "world knowledge" that can be used to set the values of the hyperparameters α_k and β_k. For example, if we expect the value of the Fourier signature to vary by only a small amount in the neighborhood of a given location, the prior on σ_{jk}^2 should reflect this knowledge by being peaked about a specific value.

The generative model for Fourier signature measurements specified by (17) is now used to compute the appearance likelihood given by (16). In addition to integrating over y_{jk}, we also integrate over the variance σ_{jk}^2, as we are not interested in its value. It follows that

$$P(A \mid T) = \prod_{j=1}^{M} \prod_{k=1}^{K} \int_{\sigma_{jk}^2} IG(\alpha_k, \beta_k) \times \int_{y_{jk}} \mathcal{N}\left(\mu, \frac{\sigma_{jk}^2}{\kappa}\right) \left(\mathcal{N}(y_{jk}, \sigma_{jk}^2)\right)^{|S_j|}. \quad (18)$$

We prove in Appendix B that performing the integration over y_{jk} and σ_{jk}^2 gives the expression for the appearance likelihood as

$$P(A|T) \propto \prod_{j=1}^{M} C_j^K \prod_{k=1}^{K} \Gamma(\gamma_{jk}+1) \left(\beta + \frac{1}{2}\Phi_{jk}\right)^{-(\gamma_j+1)} \quad (19)$$

where

$$\begin{aligned} C_j &= (\kappa + |S_j|)^{-(1/2)} \\ \Phi_{jk} &= \kappa (\mu_k^\star - \mu)^2 + \sum_{a_i \in S_j} \left(a_{ik} - \mu_{jk}^\star\right)^2 \\ \mu_{jk}^\star &= \frac{\kappa\mu + \sum_{a_i \in S_j} a_{ik}}{\kappa + |S_j|} \\ \gamma_j &= \alpha + \frac{|S_j|}{2} + 1 \end{aligned}$$

and constants that do not affect the likelihood ratio have been omitted.

The appearance model presented above is not specific to Fourier signatures. Indeed, it is a general-purpose clustering model that assumes that the data to be clustered are distributed as a mixture of Gaussians with an unknown number of components. A topology labels each data instance as arising out of one of the mixture components, where the number of mixture components is determined by the topology.

C. Putting it Together

The odometry and appearance likelihoods and the prior on topologies required to compute the target distribution (2) are given by (14), (19), and (9), respectively. We use this target distribution to sample using **Algorithm 1**, as explained before.

VII. Experiments and Results

Three sets of experiments were performed to validate the PTM algorithm. All experiments were performed using an ATRV-Mini mounted with an eight-camera rig. The landmarks in the experiments were selected manually. In all cases, we initialized the sampler with the partition that assigned each measurement to its own set. We describe the experiments and results below.

The first experiment was conducted using a relatively short run of the robot. Nine landmark locations were observed during the run of approximately 15 m. The raw odometry obtained from the robot, labeled with the landmark locations, and the ground-truth topology are shown in Fig. 9. Only the odometry measurements were used in the experiment, no appearance information was provided to the algorithm. This was done by simply neglecting the appearance likelihood term in (2). The penalty radius was set to 2.5 meters for this experiment.

Fig. 10 shows the evolution of the MCMC sampler for different values of the maximum penalty parameter. In our algorithm, it is the penalty term that facilitates merging of nodes in the map that are the same. Without the penalty, the system has no incentive to move toward a topology with a lesser number of nodes, as this increases the odometry error. Fig. 10(a) illustrates this case. It can be seen that the topology that is closest to the raw odometry data and also having the maximum possible nodes gets the maximum probability mass. For the rest of the cases with maximum penalties equal to 50, 100, and 150, respectively, the most likely solution is the topology indicated by the raw odometry. The large error in odometry makes the ground-truth topology less likely in these cases. The ground-truth topology is the second-most likely topology for maximum

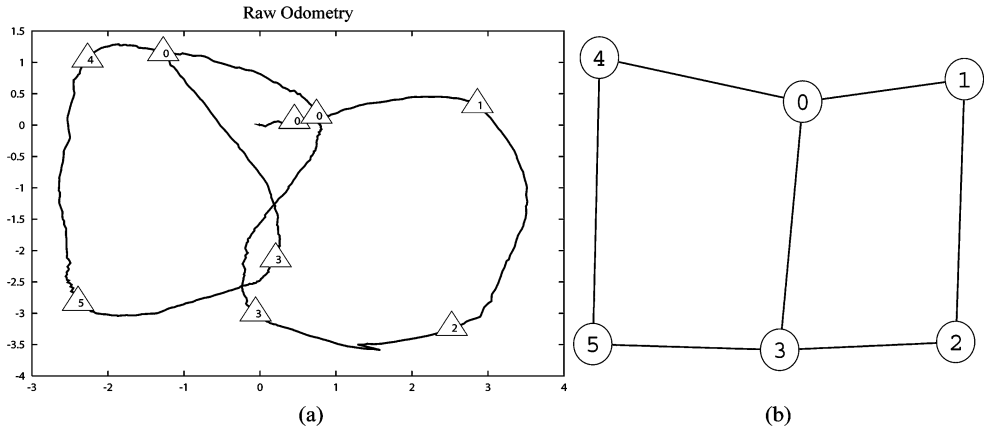

Fig. 9. (a) Raw odometry (in meters). (b) Ground truth topology from the first experiment involving nine observations.

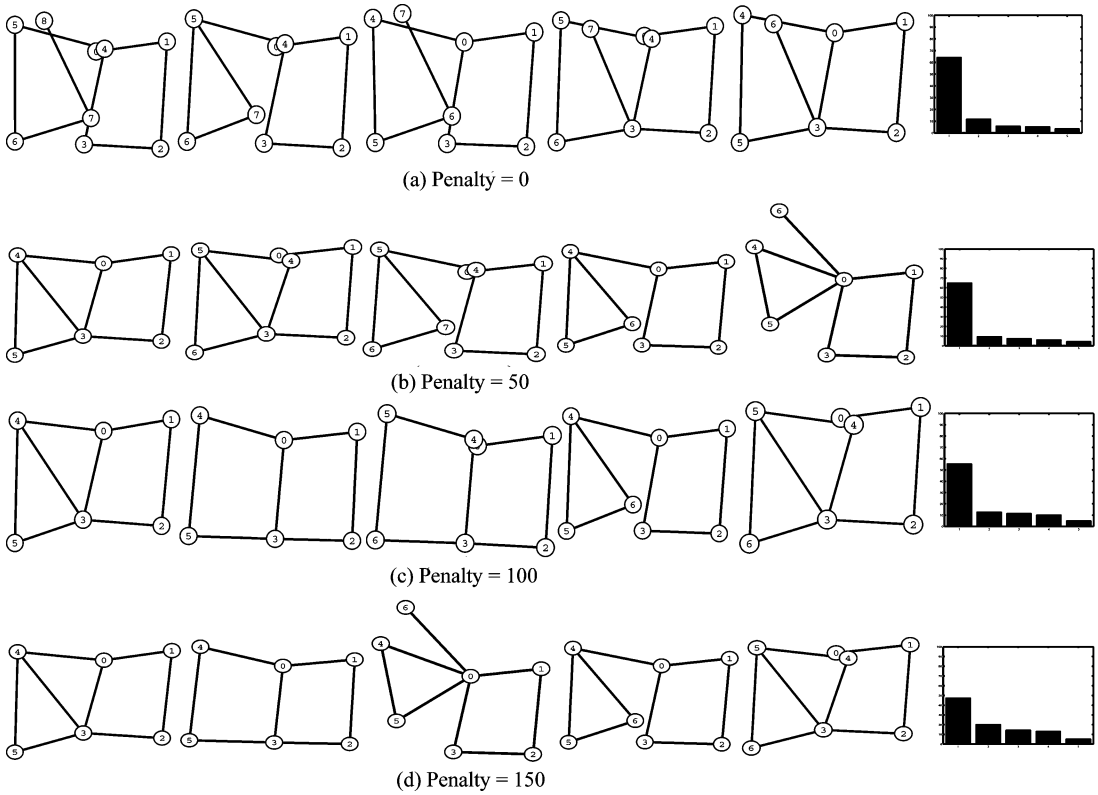

Fig. 10. Change in probability mass with maximum penalty of the five most probable topologies in the histogrammed posterior. The histogram at the end of each row gives the probability values for each topology in the row.

penalty values 100 and 150. This is because as the penalty is increased, the effect of odometry is diminished, and the ground-truth topology gains more of the probability mass. However, a very large penalty swamps odometry data and makes absurd topologies more likely.

The second experiment demonstrates the usefulness of appearance in disambiguating noisy odometry measurements. The experiment was conducted in an indoor office environment where the robot traveled along the corridors in a run of approximately 200 m and observed nine landmarks. A floorplan of the experimental area is shown in Fig. 11. The landmark locations obtained using odometry are shown in Fig. 12. As in the first experiment, the five most likely topologies from the target distribution were obtained using only odometry measurements. A penalty radius of 20 m and a maximum penalty of 100 were used to obtain the topologies, which are shown in Fig. 13. As before, the ground-truth topology receives only a small probability, due to noisy odometry.

We now repeat the experiment, but this time also using the appearance measurements, i.e., the Fourier signatures of the panoramic images obtained from the landmark locations, in addition to the odometry. The first five frequencies of the Fourier signatures were used for this purpose. The values of the variance hyperparameters in the appearance model were set so that the prior over the variance is centered at 500 with a variance of 50. The five most likely topologies in the resulting probability histogram are shown in Fig. 14. The ground-truth topology gets the majority of the probability mass. This experiment illustrates the

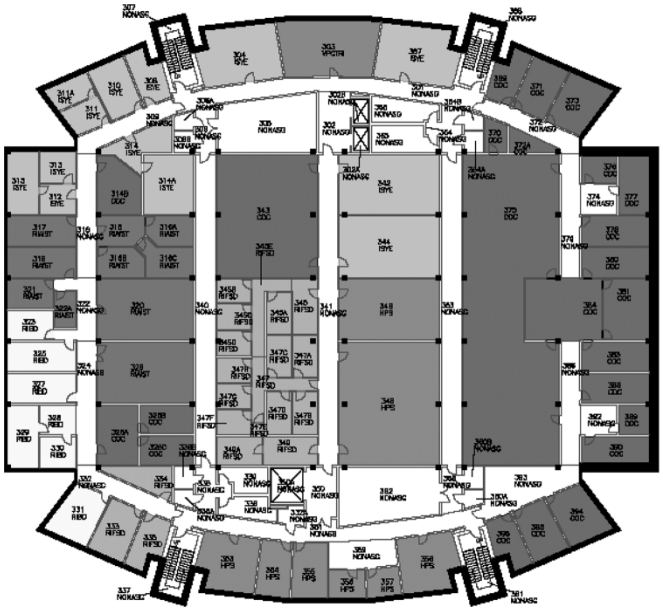

Fig. 11. Floorplan of experimental area for second experiment.

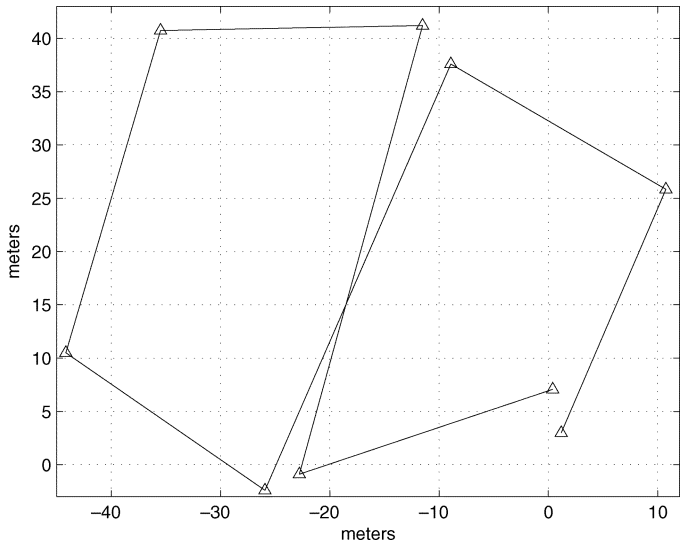

Fig. 12. Landmark locations (in meters) plotted using odometry for second experiment.

use of appearance measurements to disambiguate noisy odometry data. Additionally, it demonstrates that the Bayesian model used herein refines the posterior over topologies given more data.

The third experiment was conducted over an entire floor of a building, and was complex in the sense that the robot run contained two loops, a bigger loop enclosing a smaller loop. The odometry of the robot with the laser plotted on top is shown in Fig. 15. Twelve landmarks were observed by the robot during the run, shown overlaid on a floorplan of the experimental area in Fig. 16. A penalty radius of 3.5 m and a maximum penalty value of 100 were used in this experiment. Using only the odometry measurements, the ground-truth topology did appear in the five most probable topologies in the PTM, but received a low probability mass. These results are given in Fig. 17.

When appearance is also included (Fig. 18), the results shift dramatically, since there is little perceptual aliasing in this environment. Only two topologies appear in the PTM, with the ground truth receiving almost all the probability mass. This experiment illustrates the fact that when reliable measurements are available, the PTM computed by our approach is sharply peaked and concentrated on very few topologies.

VIII. Conclusions

In the first experiment, even though the environment is small, noisy odometry results in the ground-truth topology not receiving the highest probability mass. If an ML approach were used in this case, the result would just be an incorrect topology. However, using a Bayesian methodology to compute the posterior over topologies yields a robust result for the given data. Our technique yields the systematic, complete answer for the given data. Subsequently, the resulting posterior can be postprocessed, if necessary, using an application-specific technique to yield a single topology. On the other hand, computations such as planning, topological localization, and metric map creation can be performed using the full posterior without rejecting any possibility. For example, plans can be computed on multiple topological maps sampled from the posterior distribution; each plan being given a confidence rating proportional to the probability of the sampled topology. In contrast to ML or other truncated approaches, our technique allows such general-purpose, application-specific use of the output.

The first experiment also highlights the sensitivity of the inference to the penalty parameters. As pointed out by a referee, this problem is a fundamental one arising from the attempt to combine two incommensurable measures of goodness: one being the continuous probability mass derived from odometry, and the other being a discrete preference for a smaller number of distinct places. While we have used the prior on landmark locations to address this issue in this paper, this is by no means the optimal solution to the problem.

The second experiment illustrates the power of using a Bayesian approach in the sense that good results are obtained even with noisy data, when a large amount of data is available. Initially, due to odometry drift in the large environment, the information available to the algorithm is limited, and hence, an incorrect topology gets a large majority of the probability mass. However, the inclusion of appearance measurements, which are themselves noisy, in the inference results in a posterior in which the ground-truth topology is highly probable. Note, however, that this does not imply that the second posterior is better in any sense, since there exists only one posterior for a given set of measurements. The experiment merely affirms the fact that use of more data from varied sources improves Bayesian inference and yields more robust results. In particular, noisy odometry and aliased appearance may combine to prevent the ground-truth topology from becoming the maximum *a posteriori* topology, even after a long exploratory sequence.

These conclusions are confirmed by the third experiment, which, in addition, demonstrates that with reliable measurements, the posterior becomes sharply peaked, and the PTM ap-

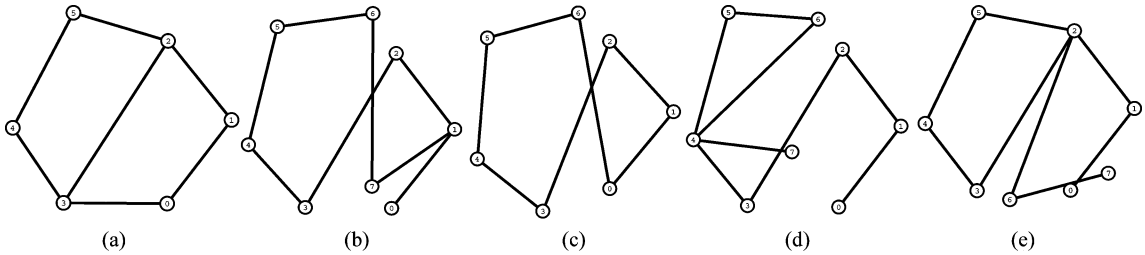

Fig. 13. Topologies with highest posterior probability mass for the second experiment using only odometry. (a) An incorrect topology receives 91% of the probability mass, while the ground-truth topology (b) receives 6%, and (c), (d), and (e) receive 0.9%, 0.8%, and 0.7%, respectively.

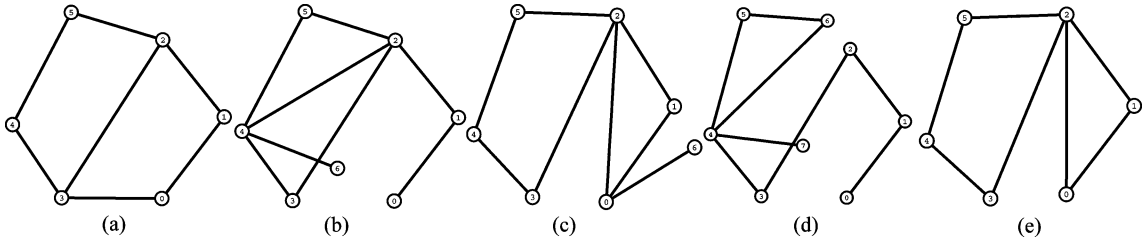

Fig. 14. Topologies with highest posterior probability mass for the second experiment using odometry *and* appearance. (a) The ground-truth topology receives 94% of the probability mass, while (b), (c), (d), and (e) receive 3.2%, 1.2%, 0.3%, and 0.3% of the probability mass, respectively.

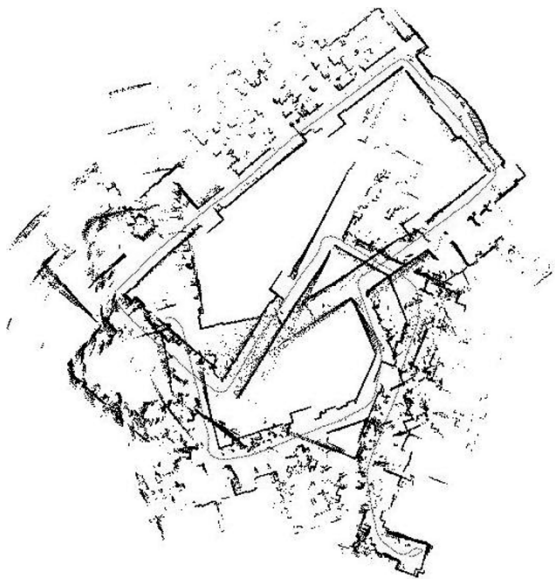

Fig. 15. Odometry of the robot plotted with the laser measurements for the third experiment.

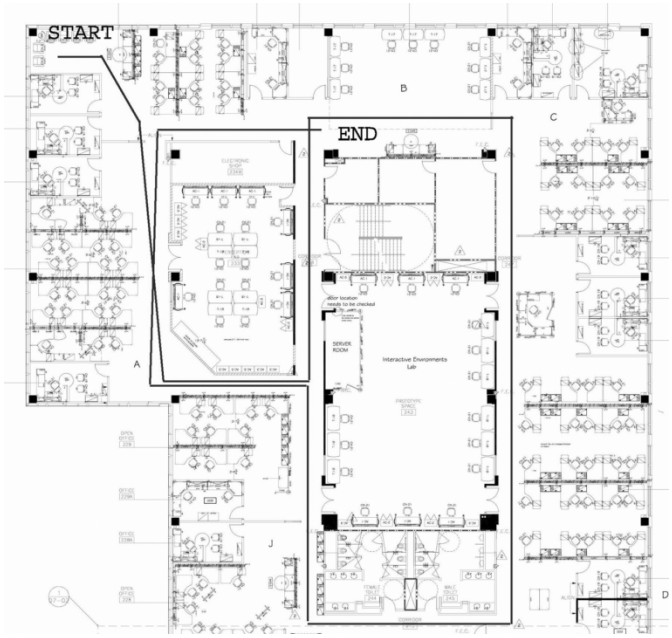

Fig. 16. Robot path overlaid on a floorplan of the environment for the third experiment.

proach defaults to a mapping method that finds a single, ML topology.

IX. Discussion

We presented the novel idea of computing discrete probability densities over the space of all possible topological maps. The PTMs are computed using MCMC sampling over set partitions that are used to encode the topologies. PTMs are a systematic solution to the perceptual aliasing problem in topological mapping, and provide an optimal estimation of the posterior distribution over topologies for the given measurements. We provide a general framework for estimating the posterior over the space of topologies, and two specific models for computing measurement likelihood, one using odometry and the other using appearance. The odometry likelihood computation uses a simple spatial prior on landmark distribution in the form of a cubic penalty function that disallows proximity among landmarks. The appearance model used in this paper deals with Fourier signatures of panoramic images. The model clusters similar appearance measurements as coming from the same spatial location. Experimental results on environments with varied sizes demonstrate the applicability of PTM.

One advantage of our approach is that an estimate of topology is possible, even if only a meager amount of information is available. It is not the purpose of this paper to find the best topological map, but to compute the posterior probability density over topological space, as per the Bayesian approach. We have shown this

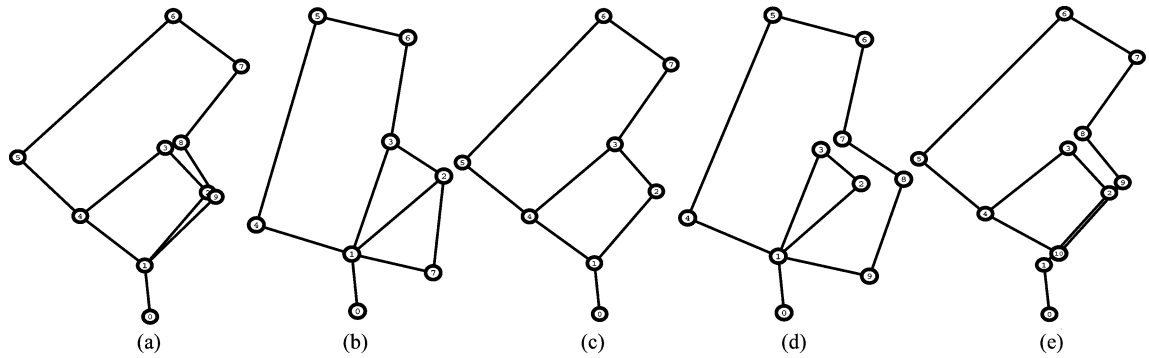

Fig. 17. Topologies with highest posterior probability mass for the third experiment using only odometry. (a) received 43% of the probability mass, while (b), (c), (d), and (e) receive 14%, 7.3%, 3.9%, and 2.8% of the probability mass, respectively. The ground-truth topology is (c).

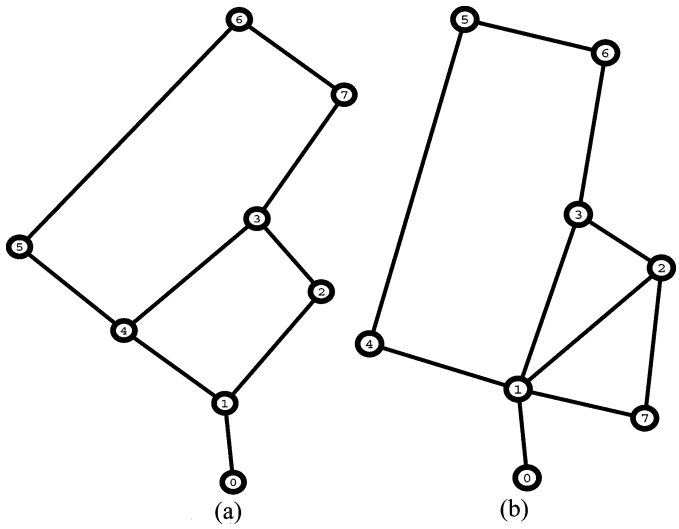

Fig. 18. Two topologies constituting the PTM when both odometry and appearance measurements are used. The ground-truth topology on the left received 99.5% of the probability mass.

capability in experiments that use only odometry from the robot to create distributions that can either correspond to the odometry or the prior (in this case, the spatial penalty function) as parameters are varied. Appearance modeling has largely been used in this work as a disambiguation mechanism for odometry, i.e., by either increasing or decreasing the evidence for the odometry. Of course, sophisticated appearance models that are more robust to perceptual aliasing could be used. However, our use of low-resolution Fourier signatures demonstrates the ability of our system to cope with environments containing barely distinguishable landmarks.

The PTM algorithm is a major step toward acknowledging the idea that the ideal mapping algorithm, capable of producing an accurate map in any environment using just the available measurements, may not exist. Instead, the mapping algorithm should be able to reason about and flag any uncertainty it might have about the maps it is generating. This is precisely what PTMs accomplish. PTMs could also be used as the basis to create a posterior over all possible metric maps using the approach given in [34].

We have shown that the inference space in which we operate is combinatorial. However, this does not cause problems in scalability, since in real environments, the measurements provide enough information so that the posterior is concentrated on a few topologies. Even if all the landmarks in the environment are perfectly aliased, inference based on only odometry still leads to a peaked posterior [42]. Only in the pathological case of very poor odometry *and* perfectly aliased appearance do we encounter problems of scalability. However, in this extreme case, the measurements do not provide sufficient information, and hence, the algorithm can hardly be blamed.

Currently, the PTM algorithm requires five parameters to be chosen by the user. These are the penalty radius and maximum penalty values for the odometry likelihood, and the α and β variance hyperparameters and the number of frequency components for the appearance likelihood. The penalty values depend on the size and scale of the environment being mapped and need to be empirically determined for each environment. This is also the case for the variance hyperpriors, which encode the variation in appearance values from the same location in the environment. Changes in lighting, camera distortion, and other measurement noise may make this variation large. It is our experience that there is rarely a need to use more than the first five frequency components in the appearance model. This is because the higher frequency components mainly contain noise, which we do not seek to model. It is also to be noted that while we use Fourier signatures in this paper, any other rotation-invariant dimensionality reduction technique can be used instead.

While we only provide likelihood models for odometry and appearance, a simple extension to laser data is also possible. If two lasers are used to gather 360° laser scans at the landmark locations, the likelihood of two scans being from the same location can be computed using scan matching. This likelihood, extended to multiple scan comparison, can be used to sample over partitions.

A problem with the current setup is the use of a single value for the penalty radius. This can cause poor performance if the distribution of landmarks varies across the environment, for example, if most of the landmarks occur in a closely spaced group, but the remainder are spread wide apart. Finding clusters at different scales is a well-researched problem in machine learning, and it is future work to apply those techniques to automate the process of setting the penalty radius.

APPENDIX A

To derive the expression for the prior over topologies given in (8), we note that the setup can be converted into an urn-ball model by considering landmarks to be urns and measurements to be balls, yielding L urns and N balls. We now show that the urn-ball model yields a prior over set partitions, which is also a prior over topologies, due to the isomorphism between topologies and set partitions.

A set partition on the measurements is created by randomly adding the balls to the urns, where it is assumed that a ball is equally likely to land in any urn (i.e., there is a uniform distribution on the urns). The distribution on the number of occupied urns, after adding all the N balls randomly to the urns, is given by the classical occupancy distribution [21] as

$$P(M) = \binom{L}{M} L^{-N} M! \left\{ \begin{matrix} N \\ M \end{matrix} \right\} \tag{20}$$

where $\left\{ \begin{matrix} N \\ M \end{matrix} \right\}$ is the Stirling number of the second kind.

The number of occupied urns after adding all the balls corresponds to the number of distinct landmarks in the topology, while the specific allocation of balls to urns (called an allocation vector) corresponds to the topology itself. Also, (20) assigns an equal probability to all ball allocations with the same number of occupied urns. Hence, we can interpret (20) as

$$\begin{aligned} P(M) \propto\ & P\ (\text{allocation vector with } M \text{ occupied urns}) \\ & \times \text{No. of allocation vectors with } M \text{ occupied urns}. \end{aligned} \tag{21}$$

The number of allocation vectors with M occupied urns is equal to the number of partitions of the set of balls into M subsets. This is precisely the Stirling number of the second kind $\left\{ \begin{matrix} N \\ M \end{matrix} \right\}$. Combining this observation with (20) and (21) yields

$$P\ (\text{allocation with } M \text{ occupied urns}) \ \propto \binom{L}{M} L^{-N} M!.$$

As mentioned previously, the probability of an allocation vector corresponds to the probability of a topology. Hence, the prior probability of a topology T with M landmarks is

$$P(T|L) = k \frac{L^{-N} \times L!}{(L-M)!}$$

which is the prior in (8). Specifying a different distribution on the allocation of balls to urns, rather than the uniform distribution assumed above, yields different priors on topologies.

APPENDIX B

To obtain the expression for the appearance likelihood given in (19), consider the integral in (7), which is the probability of a set in the topology taking into account only one frequency component

$$P(S) = \int_{\sigma_{jk}^2, y_{jk}} P(\sigma_{jk}^2) P(y_{jk}|\sigma_{jk}^2) \prod_{a_i \in S_j} P(a_{ik} \mid y_{jk}, \sigma_{jk}^2).$$

Plugging in the functional forms of the distributions defined in the model (17), we get

$$P(S) = K_j \int_{\sigma_{jk}^2} \left(\sigma_{jk}^2\right)^{-A_j} e^{-(\beta/\sigma_{jk}^2)} \int_{y_{jk}} e^{-(1/2\sigma_{jk}^2) B_{jk}}$$

where

$$\begin{aligned} K_j &= \frac{\beta^\alpha}{\Gamma(\alpha)} \frac{\kappa^{2/2}}{(2\pi)^{(|S_j|+1)/2}} \\ A_j &= \alpha + \frac{|S_j|}{2} + \frac{3}{2} \\ B_{jk} &= \kappa \left(y_{jk} - \mu\right)^2 + \sum_{a_i \in S_j} \left(a_{ik} - y_{jk}\right)^2. \end{aligned}$$

Performing the inner integration, we get

$$P(S) = K' \int_{\sigma_{jk}^2} \left(\sigma_{jk}^2\right)^{-\gamma_j} e^{-(1/\sigma_{jk}^2)(\beta + (1/2)\Phi_{jk})} \tag{22}$$

where

$$\begin{aligned} K' &= \frac{1}{(2\pi)^{|S_j|/2}} \frac{\beta^\alpha}{\Gamma(\alpha)} \left(\frac{\kappa}{\kappa + |S_j|}\right)^{1/2} \\ \Phi_{jk} &= \kappa \left(\mu^\star - \mu\right)^2 + \sum_{a_i \in S_j} \left(a_{ik} - \mu_{jk}^\star\right)^2 \\ \mu_{jk}^\star &= \frac{\kappa\mu + \sum_{a_i \in S_j} a_{ik}}{\kappa + |S_j|} \\ \gamma_j &= \alpha + \frac{|S_j|}{2} + 1. \end{aligned}$$

We now provide here a useful definition of the Gamma function

$$\int_0^\infty e^{-\alpha t} t^\gamma dt = \frac{\Gamma(\gamma+1)}{\alpha^{(\gamma+1)}}$$

using which, (22) can be integrated (note that t corresponds to σ_{jk}^{-2}) to yield

$$P(S) = K' \frac{\Gamma(\gamma_j + 1)}{\left\{\beta + \frac{1}{2}\Phi_{jk}\right\}^{(\gamma_j+1)}}$$

whence (19) follows.

ACKNOWLEDGMENT

The authors would like to thank S. C. Brubaker for his earlier work on this project, as well as R. Conroy-Dalton for many helpful discussions on spatial priors. They are also grateful to M. Kaess for providing the data used in the experiments.

REFERENCES

[1] O. Aycard, F. Charpillet, D. Fohr, and J. Mari, "Place learning and recognition using hidden Markov models," in *Proc. IEEE/RSJ Int. Conf. Intell. Robots Syst.*, 1997, pp. 1741–1746.

[2] 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. Learning*, Bled, Slovenia, 1999, pp. 67–76.

[3] J. A. Castellanos and J. D. Tardos, *Mobile Robot Localization and Map Building: A Multisensor Fusion Approach*. Boston, MA: Kluwer, 2000.

[4] H. Choset and K. Nagatani, "Topological simultaneous localization and mapping (SLAM): Toward exact localization without explicit localization," *IEEE Trans. Robot. Autom.*, vol. 17, no. 2, pp. 125–137, Apr. 2001.

[5] F. Dellaert, S. M. Seitz, C. E. Thorpe, and S. Thrun, "EM, MCMC, and chain flipping for structure from motion with unknown correspondence," *Mach. Learning*, vol. 50, no. 1–2, pp. 45–71, Jan.–Feb. 2003.

[6] G. Dissanayake, H. Durrant-Whyte, and T. Bailey, *A Computationally Efficient Solution to the Simultaneous Localization and Map Building (SLAM) Problem*. Sydney, Australia: Australian Centre for Field Robotics, Univ. Sydney, 2000, Working notes of ICRA 2000 Workshop W4: Mobile Robot Navigation and Mapping.

[7] T. Duckett, "A genetic algorithm for simultaneous localization and mapping," in *Proc. IEEE Int. Conf. Robot. Autom.*, 2003, pp. 434–439.

[8] G. Dudek, S. Hadjres, and P. Freedman, "Using local information in a nonlocal way for mapping graph-like worlds," in *Proc. Int. Joint Conf. Artif. Intell.*, 1993, pp. 1639–1645.

[9] G. Dudek and D. Jugessur, "Robust place recognition using local appearance-based methods," in *Proc. IEEE Int. Conf. Robot. Autom.*, 2000, pp. 1030–1035.

[10] H. F. Durrant-Whyte, S. Majunder, S. Thrun, M. de Battista, and S. Scheding, "A Bayesian algorithm for simultaneous localization and map building," in *Proc. 10th Int. Symp. Robot. Res.*, 2001, [CD-ROM].

[11] A. Elfes, "Occupancy grids: A probabilistic framework for robot perception and navigation," *J. Robot. Autom.*, vol. RA-3, no. 3, pp. 249–265, Jun. 1987.

[12] A. Gelman, J. B. Carlin, H. S. Stern, and D. B. Rubin, *Bayesian Data Analysis*. London, U.K.: Chapman & Hall, 1995.

[13] W. R. Gilks, S. Richardson, and D. J. Spiegelhalter, Eds., *Markov Chain Monte Carlo in Practice*. London, U.K.: Chapman & Hall, 1996.

[14] A. Griewank, "On automatic differentiation," in *Mathematical Programming: Recent Developments and Applications*, M. Iri and K. Tanabe, Eds. Norwell, MA: Kluwer, 1989, pp. 83–108.

[15] R. Gutierrez-Osuna and R. C. Luo, "Lola: Probabilistic navigation for topological maps," *AI Mag.*, vol. 17, no. 1, pp. 55–62, 1996.

[16] J.-S. Gutmann and K. Konolige, "Incremental mapping of large cyclic environments," in *Proc. IEEE Int. Symp. Comput. Intell. Robot. Autom.*, Nov. 2000, pp. 318–325.

[17] D. Haehnel, W. Burgard, D. Fox, and S. Thrun, "A highly efficient FastSLAM algorithm for generating cyclic maps of large-scale environments from raw laser range measurements," in *Proc. IEEE/RSJ Int. Conf. Intell. Robots Syst.*, 2003, pp. 206–211.

[18] D. Hähnel, W. Burgard, B. Wegbreit, and S. Thrun, "Toward lazy data association in SLAM," in *Proc. 11th Int. Symp. Robot. Res.*, Sienna, Italy, 2003, [CD-ROM].

[19] W. K. Hastings, "Monte Carlo sampling methods using Markov chains and their applications," *Biometrika*, vol. 57, pp. 97–109, 1970.

[20] H. Ishiguro, K. C. Ng, R. Capella, and M. M. Trivedi, "Omnidirectional image-based modeling: Three approaches to approximated plenoptic representations," *Mach. Vis. Applic.*, vol. 14, no. 2, pp. 94–102, 2003.

[21] N. L. Johnson and S. Kotz, *Urn Models and their Applications*. New York: Wiley, 1977.

[22] I. T. Jolliffe, *Principal Component Analysis*. New York: Springer, 1986.

[23] L. P. Kaelbling, A. R. Cassandra, and J. A. Kurien, "Acting under uncertainty: Discrete Bayesian models for mobile-robot navigation," in *Proc. IEEE/RSJ Int. Conf. Intell. Robots Syst.*, 1996, pp. 963–972.

[24] K. Konolige and J.-S. Gutmann, "Incremental mapping of large cyclic environments," in *Proc. Int. Symp. Comput. Intell. Robot. Autom.*, 1999, pp. 318–325.

[25] D. Kortenkamp and T. Weymouth, "Topological mapping for mobile robots using a combination of sonar and vision sensing," in *Proc. 12th Nat. Conf. Artif. Intell.*, 1994, pp. 979–984.

[26] B. Kuipers and P. Beeson, "Bootstrap learning for place recognition," in *Proc. AAAI Nat. Conf. Artif. Intell.*, 2002, pp. 174–180.

[27] B. J. Kuipers and Y.-T. Byun, "A robot exploration and mapping strategy based on a semantic hierarchy of spatial representations," *J. Robot. Auton. Syst.*, vol. 8, pp. 47–63, 1991.

[28] J. J. Leonard and H. F. Durrant-Whyte, "Simultaneous map building and localization for an autonomous mobile robot," in *Proc. IEEE Int. Workshop Intell. Robots Syst.*, 1991, pp. 1442–1447.

[29] B. Lisien, D. Morales, D. Silver, G. Kantor, I. Rekleitis, and H. Choset, "Hierarchical simultaneous localization and mapping," in *Proc. IEEE/RSJ Int. Conf. Intell. Robots Syst.*, 2003, pp. 448–453.

[30] F. Lu and E. Milios, "Globally consistent range scan alignment for environment mapping," *Auton. Robots*, vol. 4, pp. 333–349, Apr. 1997.

[31] M. J. Matarić, "A distributed model for mobile robot environment—Learning and navigation," Master's thesis, MIT, Artif. Intell. Lab., Cambridge, MA, 1990.

[32] E. Menegatti, T. Maeda, and H. Ishiguro, "Image-based memory for robot navigation using properties of the omnidirectional images," *Robot. Auton. Syst.*, vol. 47, no. 4, pp. 251–267, 2004.

[33] E. Menegatti, M. Zoccarato, E. Pagello, and H. Ishiguro, "Image-based Monte Carlo localization with omnidirectional images," *Robot. Auton. Syst.*, vol. 48, no. 1, pp. 17–30, 2004.

[34] J. Modayil, P. Beeson, and B. Kuipers, "Using the topological skeleton for scalable global metrical map building," in *Proc. IEEE/RSJ Int. Conf. Intell. Robots Syst.*, 2004, pp. 1530–1536.

[35] M. Montemerlo and S. Thrun, "Simultaneous localization and mapping with unknown data association using FastSLAM," in *Proc. IEEE Int. Conf. Robot. Autom.*, 2003, pp. 14–19.

[36] M. Montemerlo, S. Thrun, D. Koller, and B. Wegbreit, "FastSLAM: A factored solution to the simultaneous localization and mapping problem," in *Proc. AAAI Nat. Conf. Artif. Intell.*, 2002, pp. 593–598.

[37] H. P. Moravec, "Sensor fusion in certainty grids for mobile robots," *AI Mag.*, vol. 9, pp. 61–74, 1988.

[38] K. Murphy, "Bayesian map learning in dynamic environments," in *Proc. Adv. Neural Inf. Process. Syst.*, 1999, pp. 1015–1021.

[39] J. Neira and J. D. Tardós, "Data association in stochastic mapping using the joint compatibility test," *IEEE Trans. Robot. Autom.*, vol. 17, no. 6, pp. 890–897, Dec. 2001.

[40] A. Nijenhuis and H. Wilf, *Combinatorial Algorithms*, 2 ed. New York: Academic, 1978.

[41] D. Pierce and B. Kuipers, "Map learning with uninterpreted sensors and effectors," *Artif. Intell.*, vol. 92, pp. 169–229, 1997.

[42] A. Ranganathan and F. Dellaert, "Inference in the space of topological maps: An MCMC-based approach," in *Proc. IEEE/RSJ Int. Conf. Intell. Robots Syst.*, 2004, pp. 1518–1523.

[43] E. Remolina and B. Kuipers, "Toward a general theory of topological maps," *Artif. Intell.*, vol. 152, no. 1, pp. 47–104, 2004.

[44] F. Savelli and B. Kuipers, "Loop-closing and planarity in topological map-building," in *Proc. IEEE/RSJ Int. Conf. Intell. Robots Syst.*, 2004, pp. 1511–1518.

[45] H. Shatkay and L. Kaelbling, "Learning topological maps with weak local odometric information," in *Proc. Int. Joint Conf. Artif. Intell.*, 1997, pp. 920–929.

[46] R. Simmons and S. Koenig, "Probabilistic robot navigation in partially observable environments," in *Proc. Int. Joint Conf. Artif. Intell.*, 1995, pp. 1080–1087.

[47] R. Smith, M. Self, and P. Cheeseman, "Estimating uncertain spatial relationships in robotics," in *Autonomous Robot Vehicles*, I. Cox and G. Wilfong, Eds. New York: Springer-Verlag, 1990, pp. 167–193.

[48] ——, "A stochastic map for uncertain spatial relationships," in *Autonomous Mobile Robots: Perception, Mapping, and Navigation*, S. S. Iyengar and A. Elfes, Eds. Los Alamitos, CA: IEEE Computer Society Press, 1991, vol. 1, pp. 323–330.

[49] B. Stewart, J. Ko, D. Fox, and K. Konolige, "The revisiting problem in mobile robot map building: A hierarchical Bayesian approach," in *Proc. Conf. Uncertainty Artif. Intell.*, 2003, pp. 551–558.

[50] S. Thrun, "Learning metric-topological maps for indoor mobile robot navigation," *Artif. Intell.*, vol. 99, no. 1, pp. 21–71, 1998.

[51] ——, "A probabilistic online mapping algorithm for teams of mobile robots," *Int. J. Robot. Res.*, vol. 20, no. 5, pp. 335–363, 2001.

[52] S. Thrun, D. Fox, and W. Burgard, "A probabilistic approach to concurrent mapping and localization for mobile robots," *Mach. Learning*, vol. 31, pp. 29–53, 1998.

[53] S. Thrun, S. Gutmann, D. Fox, W. Burgard, and B. Kuipers, "Integrating topological and metric maps for mobile robot navigation: A statistical approach," in *Proc. AAAI Nat. Conf. Artif. Intell.*, 1998, pp. 989–995.

[54] L. Tierney and J. B. Kadane, "Accurate approximations for posterior moments and marginal distributions," *J. Amer. Statist. Assoc.*, vol. 81, pp. 82–86, 1986.

[55] N. Tomatis, I. Nourbakhsh, and R. Siegwart, "Hybrid simultaneous localization and map building: Closing the loop with multi-hypotheses tracking," in *Proc. IEEE Int. Conf. Robot. Autom.*, 2002, pp. 2749–2754.

[56] B. Yamauchi and R. Beer, "Spatial learning for navigation in dynamic environments," *IEEE Trans. Syst., Man, Cybern. B, Special Issue Learning Auton. Robots*, vol. 26, no. 3, pp. 496–505, Jun. 1996.

[57] B. Yamauchi and P. Langley, "Place recognition in dynamic environments," *J. Robot. Syst.*, vol. 14, no. 2, pp. 107–120, Feb. 1997.

Ananth Ranganathan (M'02) is currently working toward the Ph.D. degree in computer science with the College of Computing, Georgia Institute of Technology, Atlanta. He received the B.Tech degree in computer science from the Indian Institute of Technology, Roorkee, India, in 2002.

His research focuses on probabilistic methods in robotics and machine learning. Previously, he has published papers in the areas of behavior-based architectures for robot navigation and robot motion planning.

Emanuele Menegatti (M'00) received the "Laurea" in physics from the University of Padua, Padua, Italy, in 1998, the M.Sc. degree in artificial intelligence from the University of Edinburgh, Edinburgh, U.K., in 2000, and the Ph.D. degree in computer science from the University of Padua in 2002.

From 2003 to 2004, he was a Postdoctoral Researcher with the University of Padua. Since 2005, he has been an Assistant Professor of Computer Science with the University of Padua. In 2004, he was a Visiting Researcher for some months at both Osaka University, Osaka, Japan, and the Georgia Institute of Technology, Atlanta. His research interests are in the field of robot vision and, in particular, in omnidirectional vision. His current research is aimed at developing a distributed vision system in which mobile robots can cooperate with static sensors.

Frank Dellaert (M'96) received the Ph.D. degree in computer science from Carnegie Mellon University, Pittsburgh, PA, in 2001, the M.Sc. degree in computer science and engineering from Case Western Reserve University, Cleveland, OH, in 1995, and the degree of "Burgerlijk Ingenieur" in electrical engineering from the Catholic University of Leuven, Leuven, Belgium, in 1989.

Currently, he is an Assistant Professor in the College of Computing, Georgia Institute of Technology, Atlanta. His research focuses on probabilistic methods in robotics and vision. He has applied Markov-chain Monte Carlo sampling methodologies in a variety of novel settings, from multitarget tracking to the correspondence problem. Before that, with D. Fox and S. Thrun, he introduced the Monte Carlo localization method for robot localization, which is now a standard and popular tool in mobile robotics. He has published over 50 articles in journals and refereed conference proceedings, as well as several book chapters.