

A Probabilistic Online Mapping Algorithm for Teams of Mobile Robots

Sebastian Thrun
School of Computer Science
Carnegie Mellon University
Pittsburgh, PA

Abstract

We present an efficient probabilistic algorithm for the concurrent mapping and localization problem that arises in mobile robotics. The algorithm addresses the problem in which a team of robots builds a map online, while simultaneously accommodating errors in their odometry. At the core of the algorithm is a technique that combines fast maximum likelihood map growing with a Monte Carlo localizer that uses particle representations. The combination of both yields an online algorithm that can cope with large odometric errors typically found when mapping environments with cycles. The algorithm can be implemented distributedly on multiple robot platforms, enabling a team of robots to cooperatively generate a single map of their environment. Finally, an extension is described for acquiring three-dimensional maps, which capture the structure and visual appearance of indoor environments in 3D.

1 Introduction

The problem of robot mapping has received considerable attention over the past two decades. Mapping is the problem of generating models of robot environments from sensor data. The importance of the problem is due to the fact that many successful mobile robot systems require maps for their operation. Examples include the Helpmate hospital delivery robot [43], Simmons's Xavier robot [79], and various museum tour-guide robots [5, 36, 63, 86]. Consequently, over the last decade, there has been a flurry of work on map building for mobile robots (see e.g., [11, 50, 69, 88]). Some approaches seek to devise abstract, topological descriptions of robot environments [4, 13, 47, 57], whereas others generate more detailed, metric maps [11, 25, 32, 54, 59]. Naturally, the problem of mapping lends itself nicely to multi-robot solutions, where multiple robots collaborate and jointly explore unknown environments.

In this paper, we will be interested in the problem of building detailed metric maps online, while the robot is in motion. The online characteristic is important for

a range of practical problems, such as the robot exploration problem, where mapping is constantly interleaved with decision making as to where to move next [8, 80]. To map an environment, a robot has to cope with two types of sensor noise: Noise in perception (e.g., range measurements), and noise in odometry (e.g., wheel encoders). Because of the latter, the problem of mapping creates an inherent localization problem, which is the problem of determining the location of a robot relative to its own map. The mobile robot mapping problem is therefore often referred to as the *concurrent mapping and localization problem (CML)* [51], or as the *simultaneous localization and mapping problem (SLAM)* [9, 21]. In fact, errors in odometry render the errors of individual features in the map dependent even if the measurement noise is independent, which suggests that mapping is a high-dimensional statistical estimation problem, often with tens of thousands of dimensions.

Past work on the concurrent mapping and localization problem has led to two distinct families of approaches, each of which possesses unique advantages and deficiencies:

1. The first family, for historical reasons often referred to as *SLAM* algorithms [10, 21, 51], relies on the extended Kalman filter to estimate joint posteriors over maps and robot poses. These approaches provide a sound, online solution to the mapping problem when applicable. However, the use of unimodal Gaussian posteriors inherent in the extended Kalman filter stipulates a range of restrictions. The most important restriction is that the correct associations between measurements and features in the map must be known. For this to be the case, features in the environment must possess unique signatures that enable the robot to disambiguate them. In practice, this is rarely the case (see [15] for a notable exception). Many implementations therefore rely on maximum likelihood data association methods, which tend to work well when pose errors are small [54], but are brittle in the face of larger errors. Such is usually the case in *cyclic* environments, where a robot’s pose error can grow without bounds when traversing a cycle. In particular, we note that SLAM algorithms are incapable of generating some of the larger maps shown in this paper.
2. The second family of approaches is based on Dempster’s *expectation maximization*, or *EM* algorithm [20]. Examples include [7, 19, 76, 77, 88]. EM approaches overcome the data association problem by performing hill-climbing search in the space of all maps in a way that constantly refines the estimated data association. They have successfully been applied to mapping large cyclic environments with highly ambiguous features. However, they are inherently batch algorithms, requiring multiple passes through the entire data set. As a consequence, they are not applicable to online mapping problems.

This paper seeks to combine the best of both worlds. We propose an online algorithm that builds maps from raw range measurements, without the need for exact data as-

sociation. Our work addresses four important open problems in online mapping: the problem of mapping with unknown data associations and raw sensor measurements, the problem of mapping cyclic environments, the problem of generating maps collaboratively with multiple robots, and the problem of generating three-dimensional maps of building interiors.

Our approach, which was originally described in [87], is similar to recent work by Gutmann and Konolige [33]. The core of this paper presents a statistical framework that combines an incremental maximum likelihood estimator with a posterior pose estimator. More specifically, maps are generated online by finding the most likely continuation of the previous maps under the most recent sensor measurement. However, while online, such a methodology is prone to failure when mapping a cycle, due to the larger pose error that the robot might have accrued. If the error in pose estimation is large, backwards-correction of past estimates might be required. To overcome this problem, we propose a second estimator for the residual uncertainty in the robot's pose. This estimator calculates a full posterior over poses (but not maps!), using particle representations. When closing the cycle, this estimate is used to determine the robot's pose relative to its own map, and to correct pose estimates backwards in time to resolve inconsistencies. The resulting mapping algorithm is shown to perform favorably in a range of challenging mapping problems. Moreover, it is easily extended into a modular multi-robot mapping approach, where multiple robots integrate sensor measurements into a single, global map. However, it also possesses a range of limitations, which will be discussed in due course.

Building on the basic algorithm, this paper also describes an extension that generates 3D maps, reflecting the structural and optical appearance of a building's interior. 3D maps are generated from a panoramic camera and a laser range sensor mounted perpendicular to the robot's motion direction. They are represented in a standard virtual reality format (VRML), thereby enabling people to interactively 'fly through' an environment from a remote location.

The paper is organized as follows. We begin with introducing the basic statistical framework, which encompasses virtually all major probabilistic mapping algorithms in the literature. From there, we develop the specific set of estimators, making explicit the various assumptions and approximations necessary to obtain a working online algorithm. The algorithm is then extended into a scalable multi-robot mapping algorithm, which under specific initial conditions is shown to enable teams of robots to build accurate maps. Finally, we describe our approach to generating maps in 3D and show results of mapping a corridor segment in our university building. The paper is concluded with a discussion of limitations, related work, and research opportunities that arise from this work.

2 Bayesian Methods for Concurrent Mapping and Localization

2.1 Bayes Filters

Bayes filters comprise a popular class of Bayesian estimation algorithms that are in widespread use of robotics and engineering. They are also at the core of virtually any successful CML/SLAM algorithm [85]. Special versions of Bayes filters are known as Kalman filters [41], hidden Markov models [68, 67], particle filters [22, 23, 52, 66], and dynamic belief networks [16, 44, 73]. The classical probabilistic occupancy grid-mapping algorithm [25, 59] is also a (binary) Bayes filter.

Bayes filters address the problem of estimating the state of a time-varying (dynamic) system—such as a mobile robot—from sensor measurements. Let us denote the state of the environment by x ; the state at time t will be denoted x_t . Furthermore, let us denote measurements by z , and the measurement at time t by z_t . In time-varying systems (and in robotics in particular), the change of state is often a function of the controls. Controls will be denoted by u . The control at time t , u_t , determines the change of state in the time interval $(t, t + 1]$. The set of all measurements and controls will be referred as *data* and denoted by d :

$$d^t = \{z_0, u_0, z_1, u_1, \dots, z_t\} \quad (1)$$

Without loss of generality, we assume that controls and measurements are alternated. Following standard notation, we adopt the subscript t to denote to an event at time t , and the superscript t to denote to the set of all events up to and including time t .

The Bayes filter calculates the posterior probability $p(x_t|d^t)$ of the state x_t conditioned on the data leading up to time t . It does so recursively, i.e., the estimate $p(x_t|d^t)$ is calculated from the previous posterior $p(x_{t-1}|d^{t-1})$ using only the most recent control u_{t-1} and sensor measurement z_t . Once processed, data can be discarded. Hence the time required to calculate $p(x_t|d^t)$ does not depend on the size of the data set. This online property makes Bayes filters well suited for a range of robotic estimation problem, such as the concurrent mapping and localization problem discussed here.

The recursive update equation of the Bayes filter is as follows:

$$p(x_t|d^t) = p(z_t|x_t) \int p(x_t|u_{t-1}, x_{t-1}) p(x_{t-1}|d^{t-1}) dx_{t-1} \quad (2)$$

Thus, the desired posterior $p(x_t|d^t)$ is obtained from the estimate one time step earlier, $p(x_{t-1}|d^{t-1})$, after incorporating the control u_{t-1} and the measurement z_t . The nature of the conditional probabilities $p(z_t|x_t)$ and $p(x_t|u_{t-1}, x_{t-1})$, which play an essential role in this recursive estimator, will be discussed in Section 2.3 further below.

We will now derive Equation (2), similar to derivations that can be found in textbooks on probabilistic inference [73]. Bayes filters owe their validity a specific assumption concerning the nature of the environment, called the *Markov assumption*. The Markov assumption states that given knowledge of the current state, the future is independent of the past. In particular, this implies that the posterior estimate $p(x_t|d^t)$ is a sufficient statistics of past data, with regards to prediction. This assumption is true for environments in which x is the only state of the environment that might possibly affect multiple measurements. In static environments, the state comprises the robot pose and the shape of the environment itself. This suggests that the Markov assumption holds true in the concurrent mapping and localization problem. Clearly, in crowded environments the assumption is violated unless people are included in the state x . This case is therefore not addressed in the framework presented here—see [5] for techniques to accommodate people in the context of state estimation.

To derive Bayes filters, we notice that the desired posterior can be written as

$$p(x_t|d^t) = p(x_t|z_t, u_{t-1}, d^{t-1}) \quad (3)$$

Applying Bayes rule leads to

$$= \eta p(z_t|x_t, u_{t-1}, d^{t-1}) p(x_t|u_{t-1}, d^{t-1}) \quad (4)$$

where η is a normalizer. The Markov assumption allows us to simplify the term $p(z_t|x_t, u_{t-1}, d^{t-1})$, which leads to the more convenient form

$$= \eta p(z_t|x_t) p(x_t|u_{t-1}, d^{t-1}) \quad (5)$$

Next, we integrate the second term over the previous state x_{t-1} , which gives us

$$= \eta p(z_t|x_t) \int p(x_t|u_{t-1}, d^{t-1}, x_{t-1}) p(x_{t-1}|u_{t-1}, d^{t-1}) dx_{t-1} \quad (6)$$

Exploiting the Markov assumption once again, we obtain the simplified form

$$= \eta p(z_t|x_t) \int p(x_t|u_{t-1}, x_{t-1}) p(x_{t-1}|d^{t-1}) dx_{t-1} \quad (7)$$

which completes the derivation of the Bayes filter, as stated in Equation (2).

2.2 Concurrent Mapping and Localization

In concurrent mapping and localization, the state x comprises the (unknown) map of the environment and the pose of the robot. Let us denote the map by m and the pose by s . Then we have

$$x = \langle m, s \rangle \quad (8)$$

This gives rise to the following filter, which describes the problem of concurrent mapping and localization in Bayesian terminology:

$$p(s_t, m_t | d^t) = \eta p(z_t | s_t, m_t) \int \int p(s_t, m_t | u_{t-1}, s_{t-1}, m_{t-1}) p(s_{t-1}, m_{t-1} | d^{t-1}) ds_{t-1} dm_{t-1} \quad (9)$$

This equation is obtained by substituting (8) into (2).

If the environment does not change over time, which is commonly assumed in the literature on concurrent mapping and localization, we have $m_t = m_{t-1}$. Hence, the term $p(s_t, m_t | u_{t-1}, s_{t-1}, m_{t-1})$ is only non-zero if $m_t = m_{t-1}$, and the integration can be omitted:

$$= \eta p(z_t | s_t, m_t) \int p(s_t | u_{t-1}, s_{t-1}, m_{t-1}) p(s_{t-1}, m_{t-1} | d^{t-1}) ds_{t-1} \quad (10)$$

Stripping the map m of its time index leads to

$$p(s_t, m | d^t) = \eta p(z_t | s_t, m) \int p(s_t | u_{t-1}, s_{t-1}, m) p(s_{t-1}, m | d^{t-1}) ds_{t-1} \quad (11)$$

It is also common, but not essential, to assume that the effect of robot motion does not depend on the map m . This leads to the recursive estimator:

$$p(s_t, m | d^t) = \eta p(z_t | s_t, m) \int p(s_t | u_{t-1}, s_{t-1}) p(s_{t-1}, m | d^{t-1}) ds_{t-1} \quad (12)$$

This important equation is at the heart of all state-of-the-art concurrent mapping and localization algorithms. Virtually all such algorithms can be directly derived from this equation (see also [85]).

Unfortunately, the computation of the joint posterior over poses and maps is usually intractable in the context of concurrent mapping and localization, due to the high dimensionality of the space of all maps. Hence, all existing approaches make further restrictive assumptions. The SLAM family of approaches [21, 50, 62, 82], to date the only approach that implements the full posterior calculation as specified in Equation (12), assumes Gaussian-distributed noise and local-linear perception and motion models. Under these assumptions, the posterior can be computed exactly using the extended Kalman filter, transforming the exponentially hard problem into one that scales quadratically in the size of the environment [61, 62]. However, the Gaussian noise assumption has important ramifications in practice, some of which were already discussed in the introduction to this paper. Among others, it requires that sensor measurements can be uniquely associated with objects in the real world. In addition, handling moderate numbers of features (e.g., 1,000) is typically intractable with Kalman filters. These limitations motivate the development of alternative approaches that do not calculate the full posterior, such as the method presented in this article.

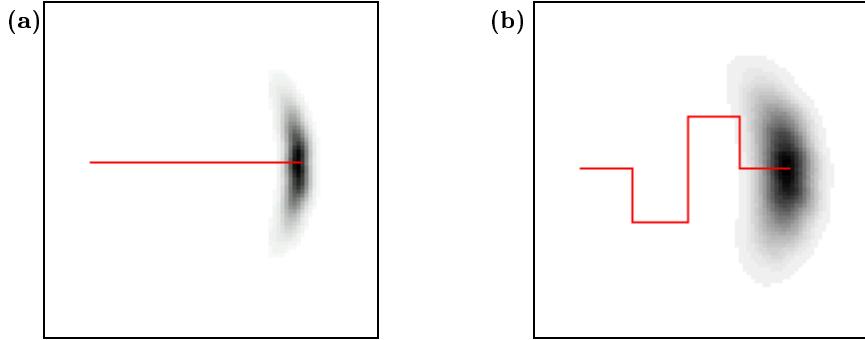


Figure 1: The motion model: Posterior distributions of the robot's pose upon executing the control illustrated by the solid line. The darker a pose, the more likely it is. Poses are projected into 2D, ignoring the robot's heading direction.

2.3 Probabilistic Models

To implement the recursive update equation (12), we need three quantities: a distribution $p(s_0, m)$ for initializing the filter, and the conditional distributions $p(s_t|u_{t-1}, s_{t-1})$ and $p(z_t|s_t, m)$. The latter distributions are often referred to as *motion model* and *measurement model*, respectively, since they physically model robot's motion and its sensors.

The initial distribution in mapping is straightforward. Without loss of generality, the pose s_0 is defined to be the origin of the coordinate system with a heading direction of 0 degrees, that is, $s_0 = \langle 0, 0, 0 \rangle$. The map m is initialized with a uniform prior.

The motion model, $p(s_t|u_{t-1}, s_{t-1})$, specifies the probability that the robot's pose is s_t given that it executed the control u_{t-1} at time $t-1$. Thus, the motion model is a probabilistic generalization of mobile robot kinematics. Classical kinematics specifies where one would expect a robot to see if its control and its motion were perfect. In practice, physical robots often suffer from drift and slippage, which induce deviations from the expected pose that are difficult to predict. Hence, the pose s_t is best modeled by a probability distribution over possible poses that the robot might attain after executing u_{t-1} in s_{t-1} .

In our implementation, we assume that the robot experiences translational and rotational error. For a small motion segment (e.g., the distance a robot traverses in .25 seconds), we assume that these errors can be modeled by two independent Gaussian noise variables, which are added to the commanded translation and rotation. Figure 1 shows two examples of the probabilistic motion model $p(s_t|u_{t-1}, s_{t-1})$. In both cases, the robot's pose prior to executing the control is shown on the left. The robot is then commanded to follow a straight-line motion as indicated. The posterior pose after executing the respective motion command is shown in Figure 1, as indicated by the grayly shaded area. The darker a pose, the more likely it is.

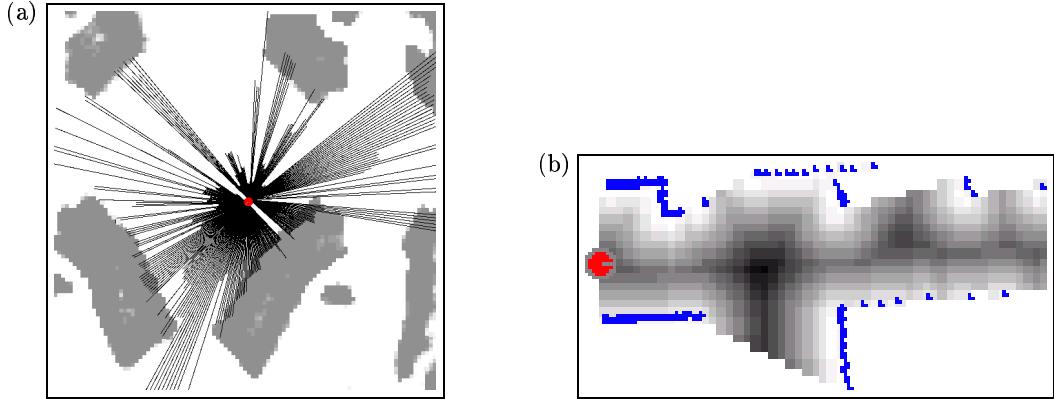


Figure 2: (a) Two laser range scans projected into an occupancy grid map. (b) Probabilistic measurement model for an example range scan, which assigns likelihood in proportion to how close to the scan a non-occluded object was observed.

Obviously, the densities in Figure 1 are only two-dimensional projections of the three-dimensional densities in pose space. A careful comparison of the two diagrams in Figure 1 shows that the control matters. Both examples show a control sequence that lead to the same expected pose, as predicted by classical kinematics. However, the path in the right diagram is longer, inducing more chances of errors as indicated by the increased margin of uncertainty.

The measurement model, $p(z_t|s_t, m)$, physically models the robot's sensors. Under the assumption that the map m and the robot pose s_t is known, the measurement model specifies the probability that z_t is measured. While the mathematical framework put forward in this paper makes no specific assumption on the nature of the sensors, we will consider robots equipped with 2D laser range finders. Such laser range finders are commonly used in mobile robotics. Figure 2a shows a range scan from a bird's eye perspective, of a range finder mounted horizontally on a mobile robot. Each line indicates the range, as measured by the sensor. Also shown is an outline of the environment, which suggests that the range measurements are of high quality.

The measurement model adopted in our implementation is a probabilistic generalization of the rich literature on scan matching [34, 54, 74]. It assumes that each scan induces a local map, which can be conceptually decomposed into three types of areas: free space, occupied space, and occluded space. The same conceptual decomposition applies to the map. Each of the measurements in a range scan z_t can thus fall into three different regions of the map. If it falls into occluded (unknown) space, the probability is uniformly high. If it coincides with occupied space, the probability is high, too. The most interesting case occurs when it falls into free-space, which can be viewed as an inconsistency. The specific probability of the measurement is then given by a function that decreases monotonically with the distance to the nearest

object. The specific function in our implementation is composed by a normal (close range) and a uniform (far range) distribution. It levels off at approximately 1 meter distance, assigning equally low probability to measurements further away. The leveling is important for the gradient ascent search algorithm described further below, as it effectively removes outliers when determining the most likely robot pose.

Figure 2b shows an example. Shown there is a sensor scan (endpoints only) for a robot placed at the left of the diagram (circle). Gray shading indicates the likelihood function: the darker a region, the smaller the likelihood for sensing an object there. Notice that occluded regions are white (and hence incur no penalty).

The likelihood computation is performed symmetrically: First, the likelihood of individual measurements in a scan is evaluated relative to the map. Then, the likelihood of obstacles in the map is calculated relative to the scan. The symmetric computation ensures that occluded measurements are matched to the ‘correct’ wall. The result of the calculation are a collection of probabilities $p(z_t^{[i]}|s_t, m)$, where $z_t^{[i]}$ denotes measurement likelihood obtained from the i -th measurement or map object. The resulting probabilities are then multiplied, assuming conditional independence between the measurements:

$$p(z_t|s_t, m) = \prod_i p(z_t^{[i]}|s_t, m) \quad (13)$$

Our implementation precomputes the nearest neighbor for a grid (e.g., 15 cm resolution). Our implementation also caches trigonometric functions that are expensive to compute, such as sine and the cosine of specific angles. As a result, calculating the measurement model $p(z_t|s_t, m)$ is extremely fast. It takes approximately one millisecond on a 500Mhz Pentium PC for a full scan.

3 Mapping Through Incremental Likelihood Maximization

3.1 Maximum Likelihood Pose Estimation

As pointed out repeatedly already, computing the full posterior (12) over maps and poses is computationally hard. Even discrete approximations scale exponentially with the size of the map. The exponential scaling can be avoided in cases where the data association is known. However, in most cases the data association is unknown and computing the full posterior online appears to be intractable given today’s computers.

We will now discuss a first online algorithm for building maps. This algorithm by itself is non-probabilistic and brittle; however, it forms the basis of a new probabilistic algorithm presented in this paper. The idea is simple, and probably because of its simplicity it is very popular: *Given a scan and an odometry reading, determine*

the most likely pose. Then append the pose and the scan to the map, and freeze it once and forever.

Put mathematically, the idea of the incremental maximum likelihood approach is to calculate a sequence of poses

$$\hat{s}_1, \hat{s}_2, \dots \quad (14)$$

and corresponding maps by maximizing the marginal likelihood of the t -th pose and map relative to the $(t-1)$ -th pose and map. To do so, the approach assumes we have at our fingertips a function for building maps incrementally, which requires knowledge of the robot's poses. Let us denote this function by \hat{m} , and call it the *incremental map updating function*:

$$\hat{m}(s^t, z^t) = \operatorname{argmax}_m p(m|s^t, z^t) \quad (15)$$

Functions of this type have been defined in many different ways. In occupancy grid mapping [25, 59], $\hat{m}(s^t, z^t)$ is simply the update equation for the occupancy grid map. In [90, 55, 32], maps are collections of scans annotated by their pose. Hence, an appropriate map update function \hat{m} simply appends a scan and a pose to the set of scans which comprise the map. This is also the approach taken in this paper.

The availability of the incremental map updating function reduces the problem of estimating posteriors over the product space of maps and poses to the simpler problem of estimating poses only:

$$\begin{aligned} p(s_t, \hat{m}(s^t, z^t) | d^t) &= p(z_t | s_t, \hat{m}(s^{t-1}, z^{t-1})) \\ &\quad \int p(s_t | u_{t-1}, s_{t-1}) p(s_{t-1}, \hat{m}(s^{t-1}, z^{t-1}) | d^{t-1}) ds_{t-1} \end{aligned} \quad (16)$$

This equation is obtained by substituting \hat{m} for m in Equation (12). The incremental maximum likelihood approach calculates at each time step one particular pose, \hat{s}_t , instead of computing the full posterior. From this estimate, it uses the function \hat{m} to obtain a single map. Thus, the posterior (16) can be simplified. Instead of assuming the only the posterior over maps and poses at time $t-1$ is known, we make the (much stronger) assumption that we know the pose and the map: The pose at time $t-1$ is \hat{s}_{t-1} , and the map is $\hat{m}(\hat{s}^{t-1}, z^t)$. Hence we can re-write (16) as follows:

$$p(s_t, \hat{m}(\hat{s}^{t-1}, s_t, z^t) | d^t) = p(z_t | s_t, \hat{m}(\hat{s}^{t-1}, z^{t-1})) p(s_t | u_{t-1}, \hat{s}_{t-1}) \quad (17)$$

In particular, we notice that the integration over poses disappeared (c.f., Equation (12)). The t -th pose is now obtained as the maximum likelihood estimate of (17):

$$\hat{s}_t = \operatorname{argmax}_{s_t} p(z_t | s_t, \hat{m}(\hat{s}^{t-1}, z^{t-1})) p(s_t | u_{t-1}, \hat{s}_{t-1}) \quad (18)$$

The estimated pose \hat{s}_t is then used to generate a new map via the incremental map updating function \hat{m} , and the new map is used from that point on.

To summarize, at any point in time $t - 1$ the robot is given a (non-probabilistic) estimate of its pose \hat{s}_{t-1} and a map $\hat{m}(\hat{s}^{t-1}, z^{t-1})$. As a new control u_{t-1} is executed and a new measurement z_t is taken, the robot determines the most likely new pose \hat{s}_t . It does this by trading off the consistency of the measurement with the map (first probability on the right-hand side in (18)) and the consistency of the new pose with the control action and the previous pose (second probability on the right-hand side in (18)). The result is a pose \hat{s}_t that maximizes the incremental likelihood function. The map is then extended by the new measurement z_t , using the pose \hat{s}_t as the pose at which this measurement was taken. Again, this is a formal description of the intuitive (and popular) idea of building maps incrementally, by determining the most likely pose \hat{s}_t and extending the map accordingly.

3.2 Search in Pose Space

We now discuss how to find the most likely pose \hat{s}_t , that is, how to compute (18). The basic problem is the continuous nature of the space of all robot poses, which makes it impossible to search this space exhaustively. In our approach, \hat{s}_t is found using gradient ascent in log likelihood space. It is common practice to maximize the log likelihood, instead of the likelihood, since the log likelihood is usually mathematically easier to handle. Notice that maximizing the log likelihood is justified by the fact that the logarithm is a strictly monotonic function. Hence, our approach seeks to compute

$$\hat{s}_t = \underset{s_t}{\operatorname{argmax}} \ln[p(z_t|s_t, \hat{m}(\hat{s}^{t-1}, z^{t-1})) p(s_t|u_{t-1}, \hat{s}_{t-1})] \quad (19)$$

This expression can be decomposed into additive terms:

$$= \underset{s_t}{\operatorname{argmax}} \ln p(z_t|s_t, \hat{m}(\hat{s}^{t-1}, z^{t-1})) + \ln p(s_t|u_{t-1}, \hat{s}_{t-1}) \quad (20)$$

$$= \underset{s_t}{\operatorname{argmax}} \sum_i \ln p(z_t^{[i]}|s_t, \hat{m}(\hat{s}^{t-1}, z^{t-1})) + \ln p(s_t|u_{t-1}, \hat{s}_{t-1}) \quad (21)$$

The differentiation of the argument on the right-hand side with respect to the pose s_t leads to the following gradient:

$$\nabla_{s_t} L = \sum_i \nabla_{s_t} \ln p(z_t^{[i]}|s_t, \hat{m}(\hat{s}^{t-1}, z^{t-1})) + \nabla_{s_t} \ln p(s_t|u_{t-1}, \hat{s}_{t-1}) \quad (22)$$

where L denotes the log-likelihood. The gradients on the right-hand side, $\nabla_{s_t} \ln p(z_t^{[i]}|s_t, \hat{m}(\hat{s}^{t-1}, z^{t-1}))$ and $\nabla_{s_t} \ln p(s_t|u_{t-1}, \hat{s}_{t-1})$, are the gradients of the measurement model, and the motion model, respectively. We notice that the motion model is differentiable, and the measurement model is piecewise differentiable. From this follows the existence of the gradients. The specific calculation of the gradient is a mechanical exercise, which is omitted in this paper since the gradients are fairly complex.

Gradient ascent then maximizes the log-likelihood by successively changing the pose s_t in the direction of the gradient.

$$s_t \leftarrow s_t + \alpha \nabla_{s_t} L \quad (23)$$

Here $\alpha > 0$ is a stepsize commonly used in gradient ascent.

Figure 3 shows two examples of maximizing the likelihood of the relative pose of two consecutive scans through gradient ascent. Each column corresponds to one example. The rows correspond to different points in time. As can be seen, 100 iterations are sufficient to obtain a good alignment even if the initial error is huge. In practice, where misalignments are much smaller, we found that usually 20 iterations are sufficient.

In the literature, mapping with incremental likelihood maximization is a popular approach [25, 59, 84, 91, 92]—with the current algorithm being merely a variant of various known algorithms. While most approaches are not phrased in the context of statistical maximum likelihood estimation, they can easily be derived as special cases of Equation (21), using similar probabilistic motion and the measurement models as the ones described in this paper. All these approaches follow the same logic: The robot maintains a single map. When a new measurement is received, a ‘plausible’ pose is estimated and the map is then grown accordingly.

4 Incremental Mapping Using Posteriors

Unfortunately, the incremental maximum likelihood approach suffers a major drawback: It is inapplicable to cyclic environments. This limitation was recognized by Gutmann and Konolige [33] and others [54, 87]. Figure 4 depicts the map obtained in a cyclic environment. Here the map appears to be locally consistent. However, local likelihood maximization still suffers a residual error that accumulates over time. At the time the robot closes the cycle, the error has grown to a reputable amount, making it impossible to generate a map that is locally consistent. The magnitude of the error can grow without bounds, depending on the robot’s odometry and the size of the cycle. Hence, the incremental maximum likelihood approach is brittle.

Technically, the inability to close cycles is caused by two factors. Firstly, the method is unable to revise poses backwards in time. This is necessary when a cycle is closed to regain local consistency. Secondly, the robot maintains only a single hypothesis as to where it is. The unawareness of its own uncertainty may prevent the robot from finding a good match when closing a cycle. Gradient methods only work well if the initial estimate is within the basin of attraction of the maximum likelihood estimate. Since the error can grow without bounds, the initial estimate might be arbitrarily far away from the best match when closing a cycle. The resulting map is then usually unusable.

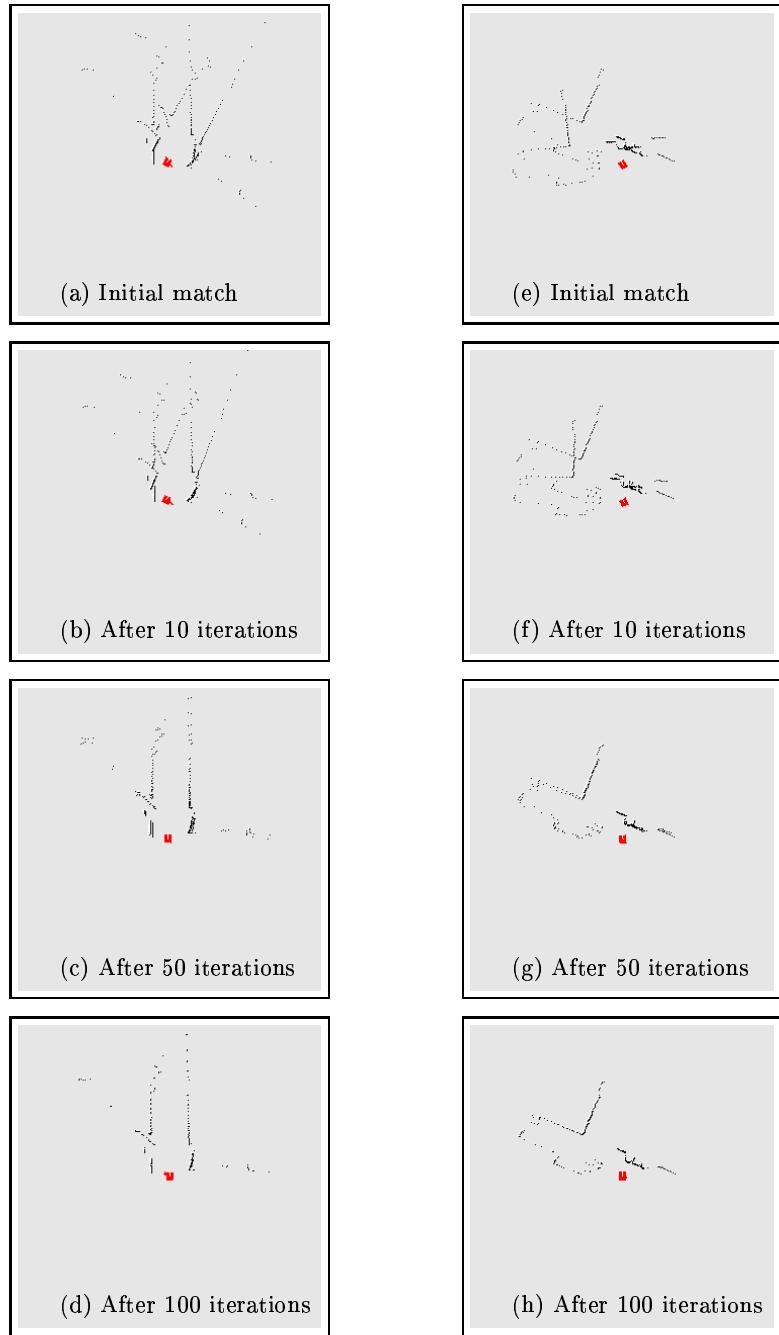


Figure 3: Two examples of gradient ascent for aligning scans (arranged vertically). In both cases, the initial translational error is 10 cm along each axis, and the rotational error is 30 degrees. The gradient ascent algorithm safely recovers the maximum likelihood alignment.

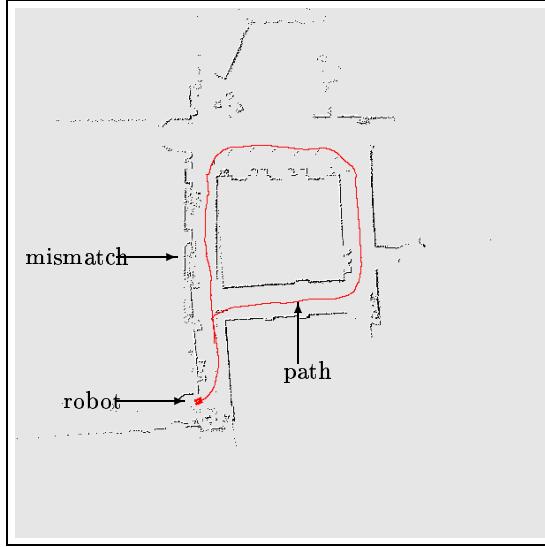


Figure 4: A typical map obtained using the most simple incremental maximum likelihood approach. While the map is locally consistent, the approach fails to close the cycle and leads to inconsistent maps. This result elucidates the limitations of stepwise likelihood maximization, and illustrates the difficulty of mapping cyclic environments.

4.1 Pose Posterior Estimation

To remedy this problem, we will now discuss a hybrid estimation approach. This approach extends the incremental maximum likelihood algorithm by a second estimator, which estimates the full posterior over poses (but not maps!). The posterior over poses is given by

$$p(s_t|d^t, m) = p(z_t|s_t, m) \int p(s_t|u_{t-1}, s_{t-1}) p(s_{t-1}|d^{t-1}, m) ds_{t-1} \quad (24)$$

This equation is obtained by substituting the pose s for x in (2), and conditioning everything on the map m (except for the robot motion). It is essentially equivalent to the Markov localization algorithm, a popular algorithm for mobile robot localization with known maps [6, 40, 64, 81]. However, Markov localization assumes the availability of a complete map of the environment, which is not assumed here.

The key reason for memorizing the full posterior $p(s_t|d^t, m)$, instead of the maximum likelihood estimate \hat{s}_t only, is our desire to accommodate the increasing residual uncertainty when moving into unmapped terrain. As the robot traverses a cycle, the posterior grows increasingly large, and the posterior characterizes this uncertainty.

Formally, the posterior is obtained by substituting the maximum likelihood map $\hat{m}(s^{t-1}, z^{t-1})$ into Equation (24), which gives us

$$\begin{aligned} p(s_t|d^t, \hat{m}(s^{t-1}, z^{t-1})) &= p(z_t|s_t, \hat{m}(s^{t-1}, z^{t-1})) \\ &\quad \int p(s_t|u_{t-1}, s_{t-1}) p(s_{t-1}|d^{t-1}, \hat{m}(s^{t-1}, z^{t-1})) ds_{t-1} \end{aligned} \quad (25)$$

Notice the difference between this estimator and the one in (17). Equation (17) computes the posterior over poses and maps; here we assume knowledge of the map and compute the posterior over poses only. This frees us of the necessity to calculate the most likely pose and instead enables us to maintain a full posterior over all poses. On the other hand, calculating (25) requires integration, whereas (17) does not. Luckily, the integration is well understood in the rich literature on Markov localization, and below we will adopt one of the representations developed there [18, 27, 48], as described in the next section.

Knowledge of the posterior gives us a better way to calculate the maximum likelihood pose \bar{s}_t . In particular, \bar{s}_t is obtained by maximizing the posterior given by (25), instead of the marginal posterior specified in Equation (17):

$$\bar{s}_t = \operatorname{argmax}_{s_t} p(s_t | d^t, \hat{m}(s^{t-1}, z^{t-1})) \quad (26)$$

The resulting pose is generally the same when mapping unexplored terrain. However, it may differ when closing a cycle, at which point the robot might have to compensate a large pose error.

This consideration points out the key advantage of maintaining the full posterior over poses, instead of the most likely pose only. Our approach fundamentally addresses one of the two problems discussed above, namely that of errors growing without bounds when closing a cycle. The posterior estimator can accommodate such errors. In particular, when closing a cycle the posterior provides a spectrum of hypotheses as to where the robot might be. The sample representation, described in turn, provides a set of potential seeds for the gradient ascent search, which makes it possible to find best matches when closing arbitrary large cycles.

4.2 Approximation Using Monte Carlo Localization

Our algorithm implements the posterior calculation using particle filters [22, 23, 52, 66]. Particle filters apply Rubin’s idea of importance sampling [72] to Bayes filters. The resulting algorithm is known in computer vision as *condensation algorithm* [38], and in mobile robotics as *Monte Carlo localization* (MCL) [18, 27, 48]. A similar algorithm has been proposed in the context of Bayes networks as [42] *survival of the fittest*.

Particle filters represent the posteriors by a set of particles (samples). Each particle is a pose that represents a ‘guess’ as to where the robot might be. Particle filters weigh each particle by a non-negative numerical factor, commonly referred to as *importance factor*. Figure 5 shows an example of the particle representation in the context of mobile robot localization. Shown there is a sequence of posterior beliefs for a robot that follows a U-shaped trajectory. Each of the sample sets is an approximation of densities (of the type shown in Figure 1).

The general algorithm is depicted in Table 1. When transitioning from time $t-1$ to time t , new particles are generated by sampling poses from the current posterior

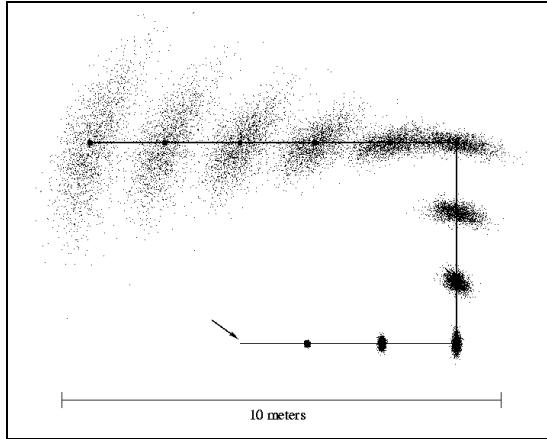


Figure 5: Sample-based approximation for the posterior over poses. Here each density is represented by a set of samples, weighted by numerical importance factors. The Monte Carlo localization algorithm is used to generate the sample sets.

and guessing a successor pose using the motion model. The weight of the particle is then obtained using the measurement model. Particle filters have been shown to approximate the Bayes filter update equation given in (2). As argued in the statistical literature, the particle representation can approximate almost arbitrary posteriors at a convergence rate of $\frac{1}{\sqrt{N}}$ [83]. It is convenient for robotics, since it is easy to implement more general than most alternatives [27], though they are generally less efficient than Kalman filters.

In the context of mapping, the application of particle filters (Monte Carlo localization) for posterior estimation is straightforward. At each point in time, the robot has access to the map built from past data. Another key advantage of the particle representation is that directly facilitates the optimization of (26) using gradient ascent. Our approach performs gradient ascent using each sample as a starting point, then computes the goodness of the result using the obvious likelihood function. If the samples are spaced reasonably densely (which is easily done with only a few dozen samples), one can guarantee that the global maximum of the likelihood function can be found. This differs from the simple-minded approach above, where only a single starting pose is used for hill-climbing search, and which hence might fail to produce the global maximum (and hence a usable map).

4.3 Backwards Correction

As argued above, when closing cycles it is imperative that maps are adjusted backwards in time. The amount of backwards correction is given by the difference Δ_{s_t} :

$$\Delta_{s_t} = \bar{s}_t - \hat{s}_t \quad (27)$$

```

Algorithm particle_filter( $X_{t-1}, u_{t-1}, z_t$ ):
  for  $i = 1$  to  $n$  to
    sample  $x_{t-1}^{(i)}$  from  $X_{t-1}$  according to the importance factors  $w_{t-1}^{(j)}$  in  $X_{t-1}$ 
    sample  $x_t^{(i)} \sim p(x_t | u_{t-1}, x_{t-1}^{(i)})$ 
    set  $w_t^{(i)} = p(z_t | x_t^{(i)})$ 
    add  $\langle x_t^{(i)}, w_t^{(i)} \rangle$  to  $X_t$ 
  endfor
  normalize all weights  $w_t^{(i)}$  in  $X_t$  so that they sum up to 1.
  return  $X_t$ 

```

Table 1: The particle filter algorithm. This algorithm is an approximate version of Bayes filters that uses particles to represent the posterior belief. The input is a weighted particle set X_{t-1} representing $p(x_{t-1} | d^{t-1})$, along with a control u_{t-1} and a measurement z_t . The output is a new weighted set of n particles representing $p(x_t | d^t)$.

where \hat{s}_t is as defined in (21) and \bar{s}_t as in (26), respectively. This expression is the difference between the *incremental* best guess (c.f., our base-line approach) and the best guess using the full posterior.

If $\Delta_{s_t} = 0$, which is typically the case when *not* closing a loop, no backwards correction has to take place. When $\Delta_{s_t} \neq 0$, however, a shift occurred due to a reconnection with a previously mapped area, and poses have to be revised backwards in time.

Our approach does this in three steps:

1. First, the size of the loop is determined by determining the scan in the map that led to the adjustment (this is a trivial side-result in the posterior computation).
2. Second, the error Δ_{s_t} is distributed *proportionally* among all poses in the loop. This computation does *not* yield a maximum likelihood match; however, it places the intermediate poses in a good starting position for subsequent gradient ascent search.
3. Finally, gradient ascent search is applied iteratively for all poses inside the loop, until the map is maximally consistent (maximizes likelihood) under this new constraint arising from the cycle.

These three steps implement an efficient approximation algorithm to the maximum likelihood estimator for the entire loop. Strictly speaking, this approach is not an online algorithm any longer, since the time required for backwards correction is not constant. A remedy might be found by amortizing the costs of closing the loop over the time. However, we found that in practice the entire correction can well be

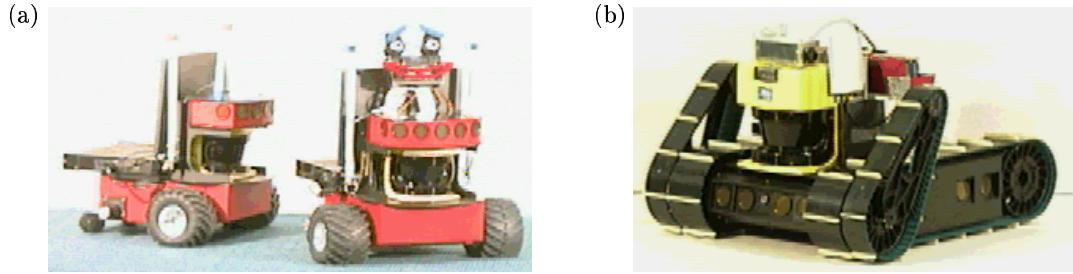


Figure 6: Robots: (a) Two of the pioneer robots used for multi-robot mapping. (b) Urban robot for indoor and outdoor exploration. The urban robot's odometry is extremely poor. All robots have been manufactured by RWI/ISR.

performed between subsequent range measurements, for all environments discussed in the experimental result sections of this paper.

5 Results for Single-Robot Mapping

We validated the basic algorithm in various different environments. The mapping algorithm was at the core of several demonstrations to Government officials in DARPA's Tactical Mobile Robot (TMR) program. The role of CMU's Mercator project was to demonstrate the feasibility of rapidly acquiring accurate maps of building interiors with mobile robots. Our experiments are aimed at illustrating the robustness and accuracy of our approach in environments that traditionally are difficult to map.

The majority of the results discussed below were obtained with the robots shown in Figure 6. Figure 6a shows two RWI Pioneer AT robots, which were modified by replacing two of the wheels with passive casters (the original skid-steering mechanism was unable to turn on carpet). Figure 6b shows a tracked robot known as *urban robot*, developed by ISR. This highly versatile robot has been developed for the DARPA TMR program, for indoor and outdoor missions in rugged terrain.

Our first experiment is concerned with closing the cycle in cyclic environments. Figure 7 shows different stages of mapping, using the very same data as in Figure 4. Here, however, the posterior estimator is used to detect and correct inconsistencies when closing the cycle. The particles are also shown in Figure 7, along with the most likely robot pose. As is easily seen, the resulting map is now consistent. Moments before the cycle is closed, the accumulated error is approximately 50 centimeters, as shown in Figure 7b. When closing the cycle, the most likely pose is adjusted in accordance to the previously mapped area, thereby practically eliminating this error. The backwards correction mechanisms then distributes the adjustment along the cycle, and the new, improved map is generated in less than a second of processing time. Figure 7c shows the final map.

Figure 8 shows a second example, obtained in a cyclic hallway environment that



Figure 7: Incremental algorithm for concurrent mapping and localization. The dots, centered around the robot, indicate the posterior belief which grows over time (a and b). When the cycle is closed as in (c), the posterior becomes small again.

is larger than the previous one. Again, our approach manages to accommodate accumulated odometry errors and arrives at a map sufficiently accurate for navigation. An interesting experiment is shown in Figure 9. Here the overall computational time for mapping was reduced by a factor of five, leading to significantly worse matching results along the way. As Figure 9a indicates, the cumulative error when closing the cycle is substantial; it is approximately 4 meters. When closing the loop, the pose estimate places high likelihood at a location that is actually incorrect. The map at this time shown in Figure 9b is consistent, but does not accurately reflect the geometry of the building. It is important to notice that the estimated most likely pose is perfectly reasonable, as the resulting map is free of any major conflicts. However, as the robot continues to traverse the corridor, such a conflict becomes unavoidable. Consequently, the posterior over poses changed accordingly, and the map is revised. The resulting map is shown in Figure 9c. This map is now topo-

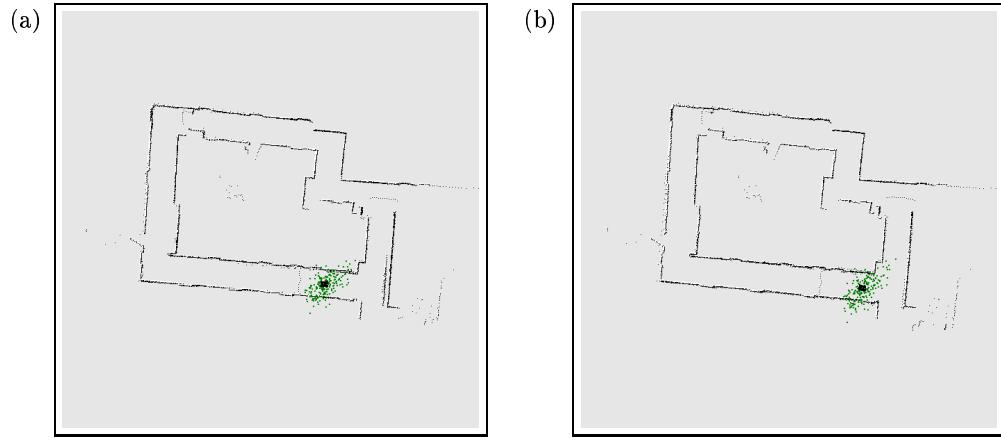


Figure 8: A larger cyclic corridor (approximately 30 meters diameter). Situation (a) before and (b) after closing the cycle.

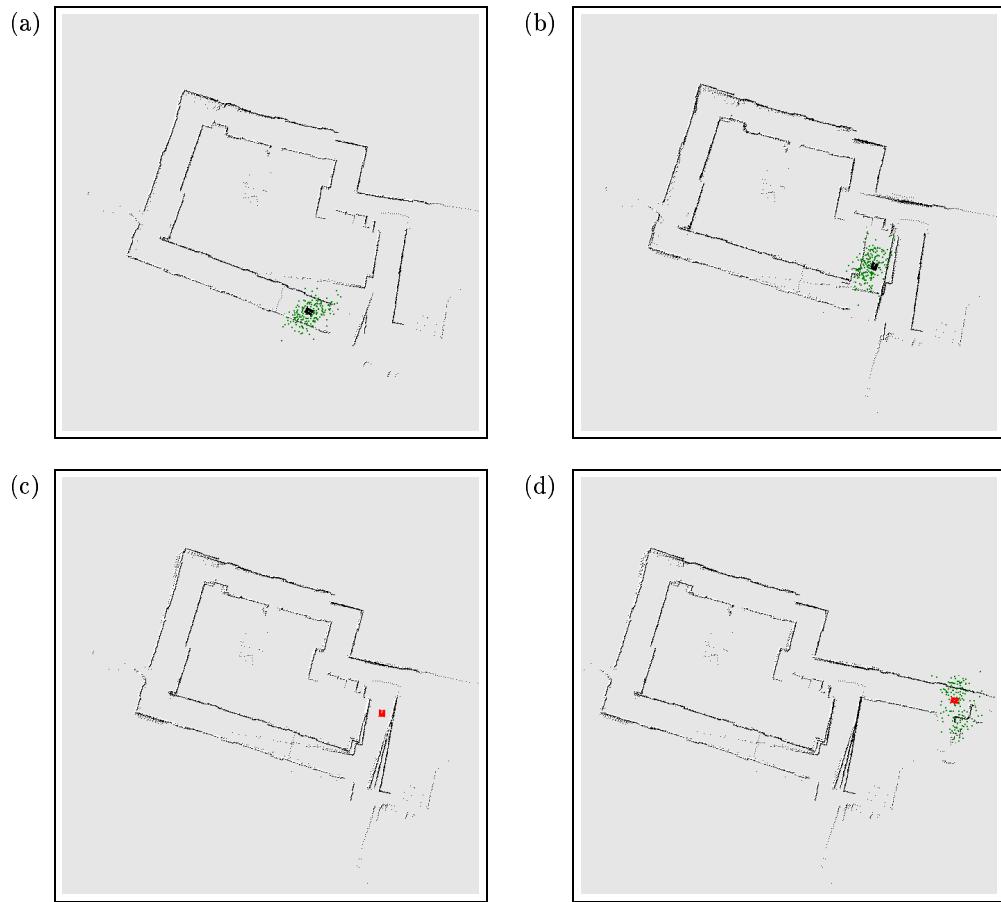


Figure 9: Same data set as in Figure 8, but with additional errors. (a) Situation before closing the cycle. (b) The initial attempt to alleviate the conflict fails. (c) However, the posterior estimate enables the robot to recover a few time steps later. (d) The final map.

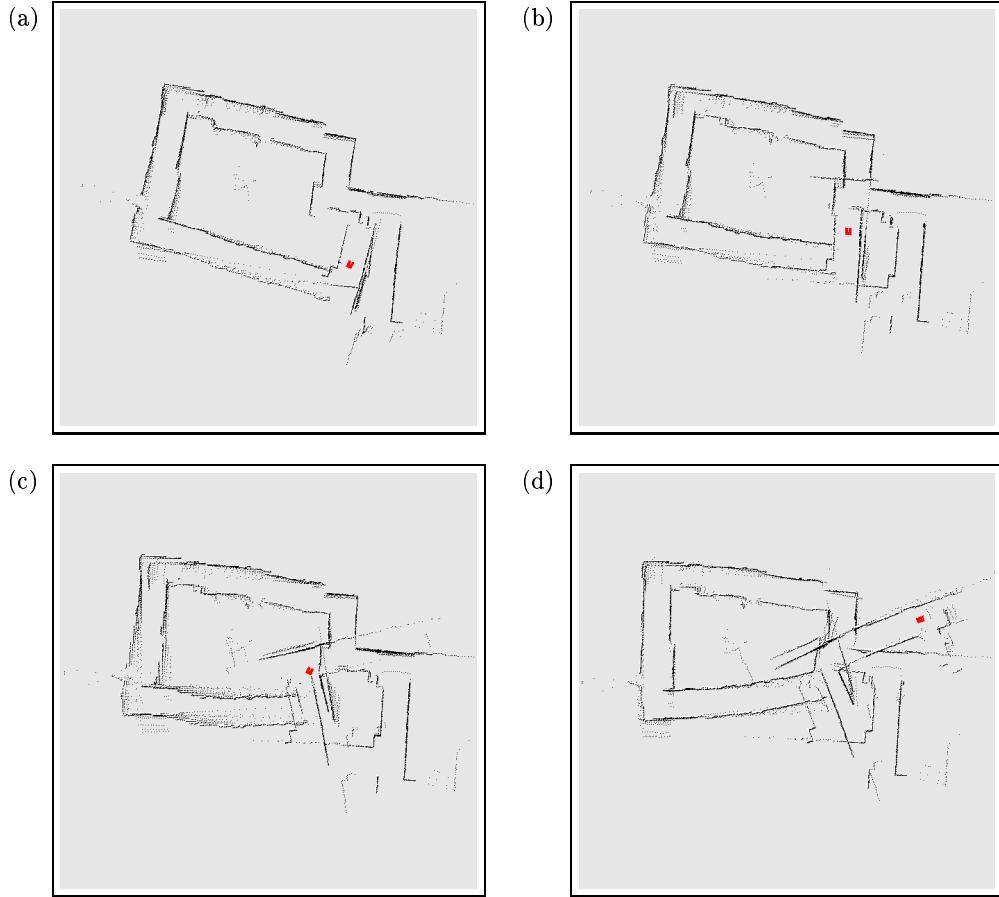


Figure 10: Mapping *without* the probabilistic posterior estimate over poses described in the paper. This result is the outcome of the naive incremental maximum likelihood mapping algorithm, obtained under otherwise identical conditions as in Figure 9. The map shown here is plainly unusable, which is in contrast to the one shown in Figure 9d. This result illustrates the importance of estimating posteriors over poses for the recovery from large odometry errors.

logically consistent, although not quite as accurate as the map obtained with more computational resources available.

The importance of the posterior estimate over poses becomes apparent when running our algorithm without estimating the full posterior, and instead maintaining a single pose estimate, as described in Section 3.1, but with our backwards correction mechanism. This is equivalent to maintaining a single sample only for the pose estimate, at the location of the most likely pose \hat{s} . Figure 10 shows the development of maps over time, obtained under otherwise equivalent conditions as those that led to the results shown in Figure 9. The situations in diagrams a and b of both figures are quite similar: The initial error is large, and the initial map is topologically incorrect. However, without the posterior pose estimator, the approach is unable

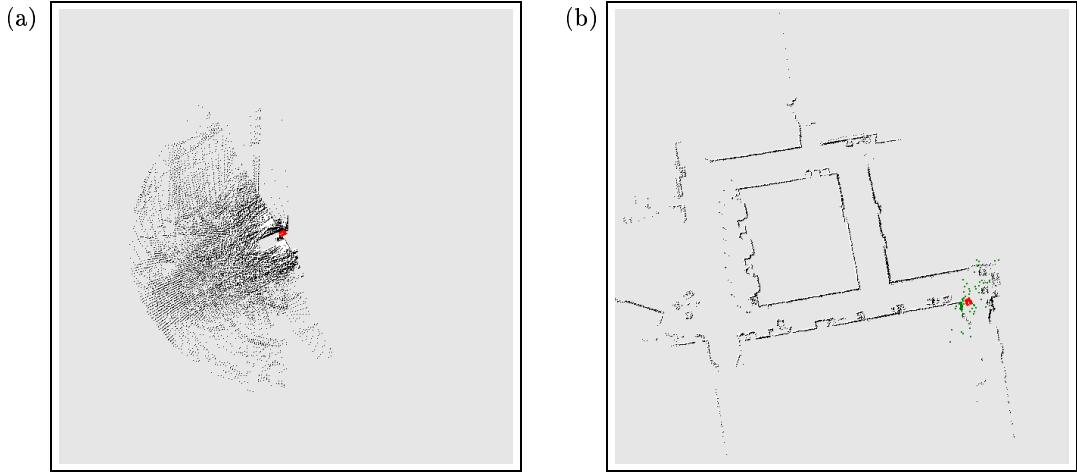


Figure 11: Mapping without odometry: (a) raw data, (b) map, generated online.

to recover from this incorrect alignment and instead generates the maps shown in Figures 9c and d. These maps are clearly incorrect to the extent that they are unusable for robot navigation. This result highlights the importance of estimating the posterior over robot poses during mapping.

To test the robustness of the mapping algorithm in the extreme, we manually deprived the data of its odometry information. Hence, only the laser range data can be used to estimate the robot’s pose. Odometry-free mapping tests the robustness of the algorithm. If an algorithm can generate a consistent map without odometry data, it can accommodate failure of a robot’s wheel encoders or inertial navigation unit (INU). Clearly, odometry-free mapping can only succeed if the environment possesses sufficiently distinguishing features. For example, without odometry it would be impossible to map a wide-open area where all obstacles are beyond the reach of the sensors, or very long, featureless hallways.

Figure 11 shows the result of mapping the cyclic environment *without odometry*. Figure 11a depicts the raw data that, without odometry, is difficult to interpret. The map in Figure 11b has been generated using the statistical technique described in this paper. In particular, the same set of parameters was used to construct this map, as if the odometry data was available. This sheds light on the robustness of our approach.

Mapping with poor odometry has practical significance, as not all mobile robots have good odometry. A particularly inaccurate robot is the urban robot shown in Figure 6b. Figure 12a depicts a raw data set gathered while the robot was autonomously exploring and mapping a military facility in Fort Sam Houston in Texas. While the odometry is quite accurate when the robot moves straight, its tracks introduce up to 100% rotational error when the robot rotates. Thus, odometry alone is unusable even over very short time intervals. The map of this non-cyclic

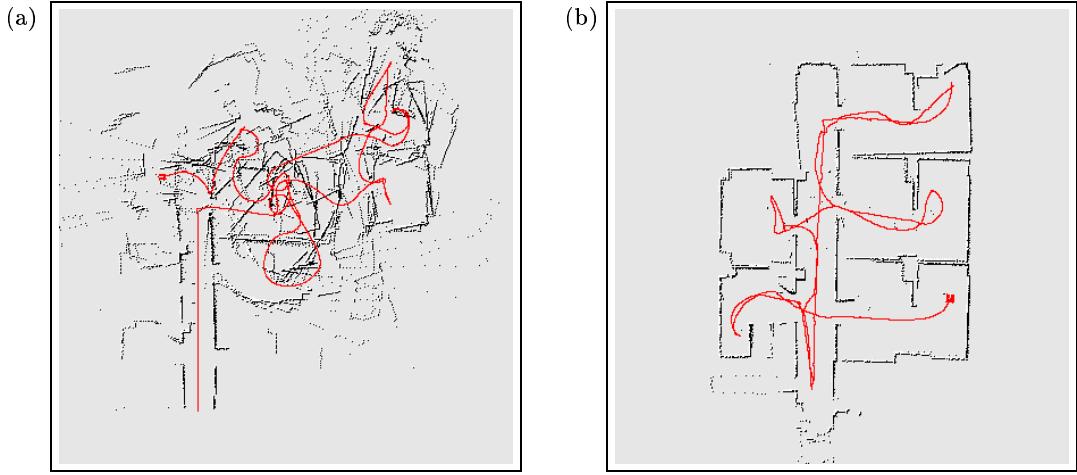


Figure 12: Autonomous exploration and mapping using the urban robot: Raw data and final map, generated in real-time during exploration.

environment is shown in Figure 12b, along with the robot’s exploration path. This map is not perfect, as attested by the fact that the walls are not perfectly aligned. However, it is sufficiently accurate for navigation.

A final result was obtained in a public place, the Tech Museum in San Jose, California. A key feature of this environment is the availability of an accurate blueprint, which enables us to evaluate the accuracy of the map acquired by the robot. The specific map was acquired with a RWI B21 robot. Figure 13a shows an occupancy grid map [25, 59] of a large center area in the museum, extracted from the aligned scans. A comparison with the architectural blueprint of the museum in Figure 13b illustrates that the learned map is quite accurate. Several differences between these maps stem from the fact that the blueprint does not accurately reflect the nature and location of the exhibits in the museum. In this regard, the map learned by the robot is more accurate than the museum’s blueprint.

6 Multi-Robot Mapping

One of the nice features of the probabilistic state estimators is that they make the extension to multi-robot mapping quite straightforward. In the multi-robot mapping problem, multiple robots gather data concurrently in an attempt to build a single, unified map of the environment. The use of multiple robots is advantageous if the speed at which a map is being constructed is important: In the best case, the speed-up by using N robots instead of one is super-unitary: In theory, environments exist where N robots consume approximately $\frac{1}{2N}$ the time it would take a single robot to map the environment. The factor 2 arises from the fact that in some environments, a single robot might waste time traversing known territory to get back to places not



Figure 13: (a) Occupancy grid map and (b) architectural blueprint of a recently constructed museum in San Jose. Notice that the blueprint is inaccurate in certain places.

yet visited. In practice, however, the speedup is usually sub-linear [80].

Our current implementation makes two important restrictive assumptions:

- First, the robots must begin their operation in nearby locations, so that their range scans show substantial overlap.
- Second, the software must be told the approximate relative initial pose of the robots. The accuracy of the initial estimate is bounded by the basin of attraction of the gradient ascent algorithm. A safe error is usually up to 50 cm and 20 degrees in orientation.

Under these conditions, the extension to multi-robot mapping is straightforward. In particular, each robot maintains a local, robot-specific estimate of its pose, while sharing the same map. The estimation of the most likely pose using the incremental maximum likelihood estimator is based on a single, global map that fuses data collected by all robots.

Put mathematically, if we use k to index each particular robot, the posterior over the k -th robot pose is given by

$$p(s_{t,k}|d_k^t, m) = p(z_{t,k}|s_{t,k}, m) \int p(s_{t,k}|u_{t-1,k}, s_{t-1,k}) p(s_{t-1,k}|d_k^{t-1}, m) ds_{t-1,k} \quad (28)$$

This equation generalizes Equation (24) to the multi-robot case with a single global map m . The most likely pose of the k -th robot is then determined by generalizing

Equation (26):

$$\bar{s}_{t,k} = \underset{s_{t,k}}{\operatorname{argmax}} p(s_{t,k} | d_k^t, \hat{m}(s_k^{t-1}, z_k^{t-1})) \quad (29)$$

where the function \hat{m} incrementally enlarges the global map m . This extension exploits the fact that none of the underlying statistics crucially depends on the number of robots. It also exploits an interesting conditional independence that enables us to factorize the posterior distribution of the joint pose space. In particular, the commitment to maximum likelihood maps eliminates possible dependencies between the pose estimates of different robots, which otherwise exist for the full Bayesian solution. Thus, in the general context, the posterior over robot poses would have to be estimated jointly, but here we can estimate it locally on each robot (see [28] for a more detailed discussion). We notice that the pose estimation can be carried out locally, on each robot. This leads to a computationally elegant decomposition of the problem, in which each robot basically performs the same basic algorithm as in single-robot mapping, but map updates are broadcast to all robots.

Figure 14 shows a map built by two autonomously exploring robots. This experiment was carried out in a TMR demonstration that took place in a training facility for fire fighters in Rockville, Maryland. The robots coordinate their exploration using a technique described in [8, 80]. In essence, this technique assigns robots to unexplored areas (frontiers) based on information-theoretic considerations. The map shown in Figure 14 is constructed online, as the robots explore, using the distributed software system outlined above. As is easily seen in Figure 14a through e, the robot quickly maps the confined operating space. Finally, they return to locations that were not mapped perfectly in the first traversal of the building. In this example, the entire operation took a little less than two minutes, and the map was generated online as the robot moved. Figure 15 shows the result of a similar experiment, carried out at Fort Sam Houston. This experiment involved three robots, marked A, B, and C, in an environment more than three times the size. The resulting map was generated in less than three minutes. The paths in Figure 15 show how the robots systematically explore the environment and succeed in building a single, consistent map—despite the poor accuracy of the urban robot (robot B in the figure).

It is possible—though currently not part of our implemented software system—to relax the restrictive assumption of known relative initial poses. If the initial poses of robots relative to each other is unknown, the robots face a global localization problem, which is the problem of localizing robots under global uncertainty. One natural way to solve the problem of determining one robot’s pose relative to another is to globally localize one of the robots in the map built by the other robot. It is well known [18, 28] that Monte Carlo localization (MCL)—which our approach uses for posterior pose estimation—can solve the global localization problem. It does so under one restrictive assumption: That the robot actually be inside the map in which it is being localized. While this assumption is trivially the case for the type problems studied in the localization literature, it is problematic in the context of

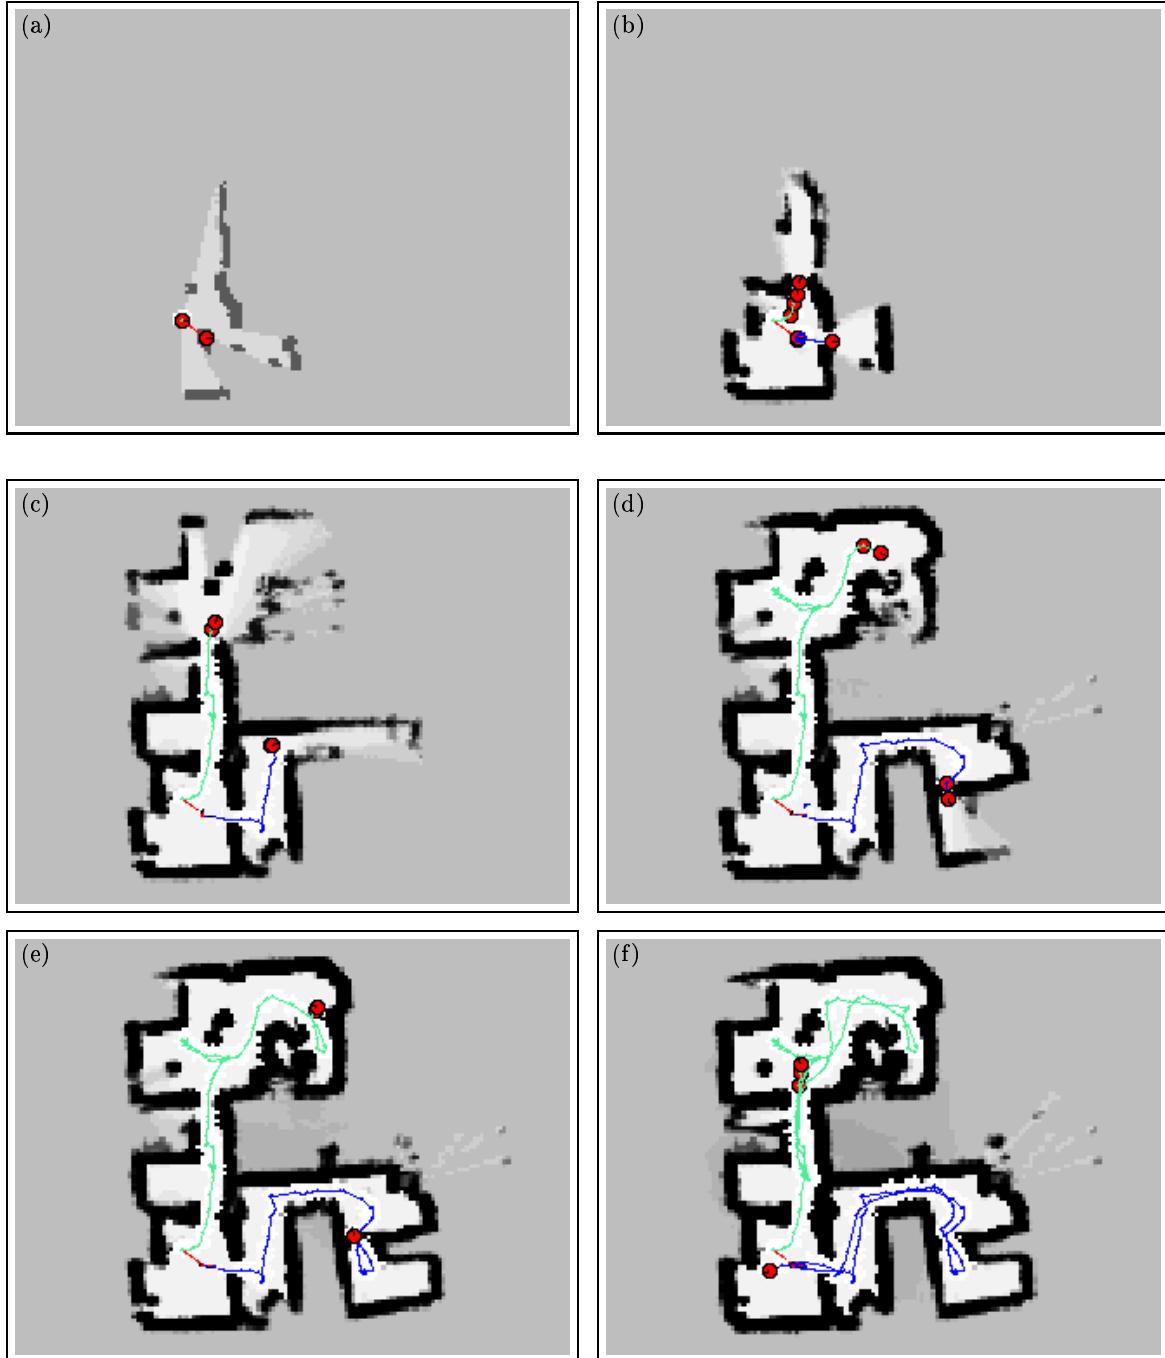


Figure 14: Online mapping during exploration with two robots.

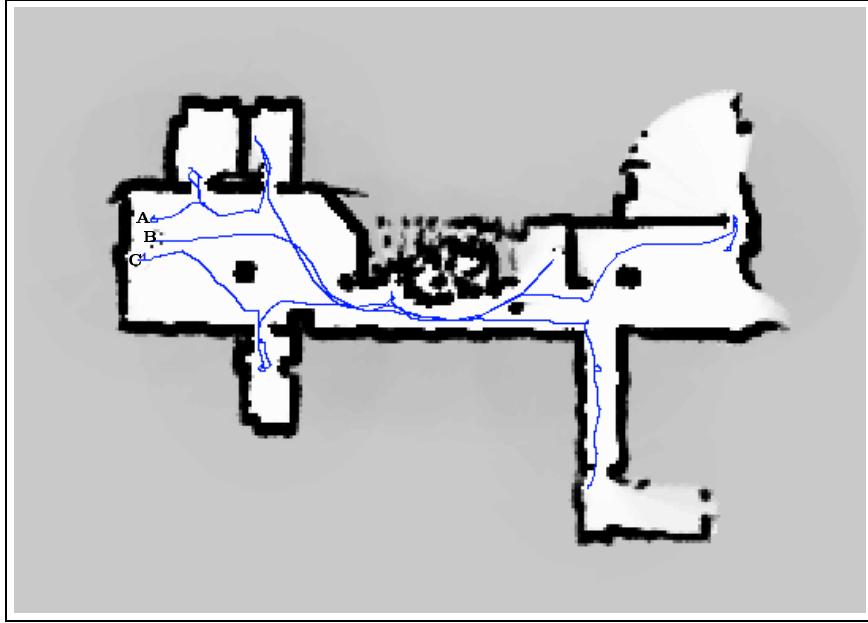


Figure 15: Map built by three autonomously exploring robots. The initial robot poses are on the left as marked by the letters A, B, and C. The robots A and C are pioneers, and the robot B is an urban robot.

multi-robot mapping. In particular, two robots might start out in different, non-overlapping parts of a building, and hence none of the robots operates inside the map built by any of the other robots. Thus, to estimate the poses of robot relative to each other, the robots must be able to estimate the degree of overlap between the areas traversed thus far. Clearly, MCL provides no answer to this hard and important problem.

However, MCL can be applied if it is known that one robot starts in the map built by another. Such a situation is shown in Figure 16. Here one robot is declared the *team leader*, which builds an initial, partial map. The second robot is then placed somewhere *inside* this map, but its pose inside this map is unknown. The particle filter of the second robot is then initialized by a set of particles drawn uniformly across the map of the first robot. The exact same posterior estimator depicted in Table 1 is then applied to globally localize the second robot in the map built by the team leader. Figure 16a shows the posterior after integrating one measurement. In this specific example, the samples have been generated in proximity to the first robot's path, assuming the second robot starts out near a location where the team leader has been before. A few moments later, the relative pose has been determined (Figure 16b) and the robot begins contributing data to the same, global map (Figure 16c).

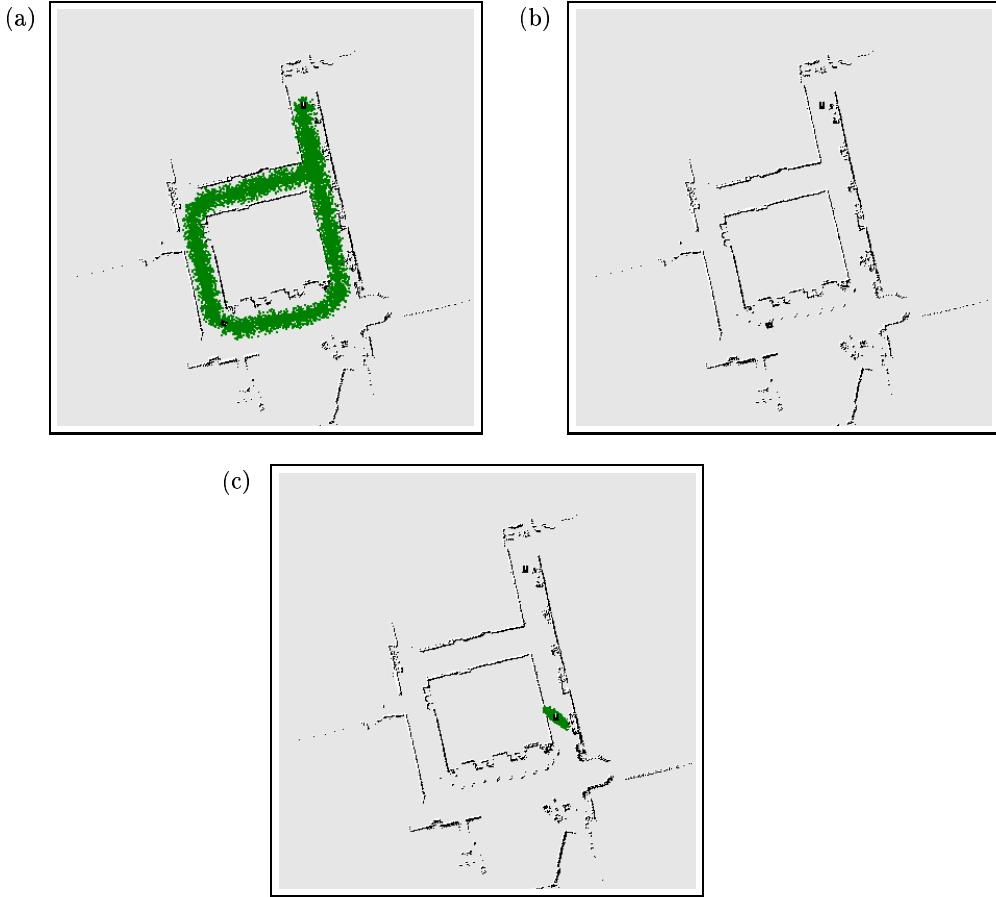


Figure 16: A second robot localizes itself in the map of the first, then contributes to building a single unified map. In (a), the initial uncertainty of the relative pose is expressed by a uniform sample in the existing map. The robot on the left found its pose in (b), and then maintains a sense of location in (c).

7 3D Structural and Texture Mapping

A second extension of the basic mapping algorithm is the idea of three-dimensional maps. Such maps can be informative to people interested in a virtual view of a building interior. For example, they enable one to pay a virtual visit using VR interfaces without actually physically entering the building. This might be useful for architects, real estate agents, and also for people preparing to operate in hazardous environments such as nuclear waste sites.

Figure 17a shows a robot equipped with two laser range finders. One of the laser range finders is pointed forward, to generate a 2D map using the techniques described above. As a result, it provides the robot with accurate pose estimates. The second laser range finder is mounted vertically, so that the light plane is perpendicular to

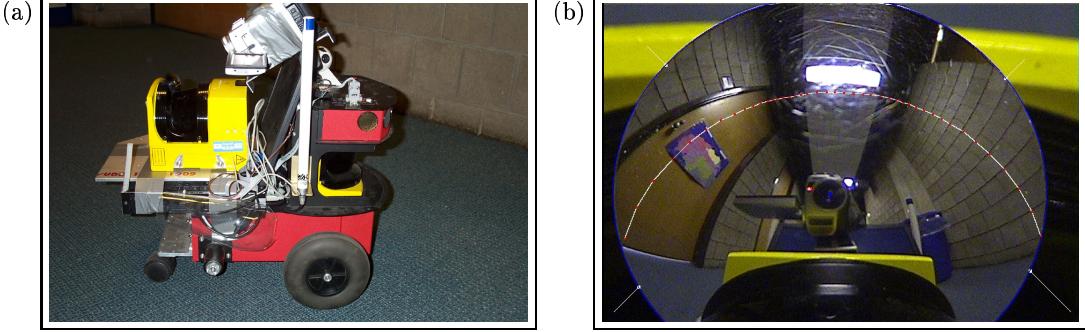


Figure 17: (a) Pioneer robot equipped with 2 laser range finders and panoramic camera used for 3D mapping. (b) Panoramic image acquired by the robot. The curved line indicates the region that corresponds to the vertical slice measured by the laser range finder.

both the first range scanner, and the motion direction of the robot. As a result, this range scanner measures cross-sections of the building as the robot moves. If the robot traverses the environment only once—which is typically the case in exploration—the measurements of the upward-pointed laser do not overlap; hence, there is little hope of using it for localization. However, the data provides information about the 3D structure of the building.

The robot shown in Figure 17a is also equipped with a panoramic camera. The panoramic camera utilizes a small mirror, which provides an image that contains the area covered by the upward-pointed laser range scanner. The mirror is mounted only a few centimeters away from the optical axis of the laser range scanner. Careful calibration of both devices enable us to match range and color measurements along the measurement plane of the vertical range finder, basically ignoring the small disparity between the optical axes of the laser and the camera. Figure 17 shows a panoramic image. The curved line highlights a vertical cross-section of the building, which corresponds to the perceptual field of the upward-pointed laser range finder.

Our approach represents 3D maps as collections of polygons with texture superimposed. Our 3D mapping algorithm proceeds in five steps, all of which can be carried out online.

1. **Pose estimation:** Pose estimation is carried out in 2D, using the techniques described above. The vertical range data is then projected into global coordinates using the maximum likelihood pose \bar{s} .
2. **Raw data smoothing:** Noise in the range measurements is reduced by local smoothing. Neighboring sensor scans are smoothed by replacing them with a locally weighted average, using a smoothing kernel of size 3 by 3. This step yields significantly smoother surfaces in the final map.
3. **Structural Mapping:** The map is then built by searching nearby range scans for nearby measurements. More specifically, for each time τ and index i , for

which the coordinates of the four measurements

$$z_{\tau}^{[i]}, z_{\tau}^{[i+1]}, z_{\tau+1}^{[i]}, z_{\tau+1}^{[i+1]} \quad (30)$$

are less than δ apart in 3D coordinates, these points are assumed to correspond to the same objects. Hence, a polygon is generated with the endpoints corresponding to those four measurements and added to the map. Notice that polygon's are only constructed for measurements recorded right after another, and for neighboring sensor beams. The test of adjacency in 3D space is essential, since large discontinuities often correspond to open space (e.g., an open door).

4. **Texture Mapping:** If a polygon is generated, the image sequence recorded between the times τ and $\tau+1$ are used as texture for the polygon. In particular, our algorithm extracts from the camera image a stripe that corresponds to the perceptual range of the laser. Stripes are then concatenated, which generates a texture image for the entire range scan, and the region corresponding to the area between the i -th and the $(i+1)$ -th range scan are extracted using a straightforward geometrical argument. This texture is then mapped onto the polygon.
5. **Map simplification:** To simplify the maps, the polygons are finally fused by a software package developed by computer graphics researchers [30, 31]. This approach, which our software uses without modification, fuses polygons based on a quadratic error measure of their appearance in the rendered image. We found that the number of polygons can be reduced by a factor of ten without significantly altering the shape of the model. The result of this simplification step is, thus, a reduced complexity model that is also smoother.

The result of this procedure is a collection of textured polygons that can be rendered using a VRML viewer. While in principle, these steps can be carried out online, we remark that our current implementation is offline, due to the lack of a fast frame grabber on our mobile robot.

Figures 18 and 19 show views of a 3D structural map (without texture). The map in Figure 18 has been created from approximately 1 minute of raw data obtained in a straight corridor segment, whereas the ones in Figure 19 include a cycle and correspond to approximately 5 minutes of robot operation. All of these maps are in VRML format, enabling a user to navigate interactively through the building. Figure 19 illustrates the effect of the map simplification step above. The left column shows an un-simplified map, which contains 82,899 polygons. The right column depicts the simplified polygonal model, which contains only 8,289 polygons (10%). The appearance of both is similar; however, rendering the more compact one is an order of magnitude faster.

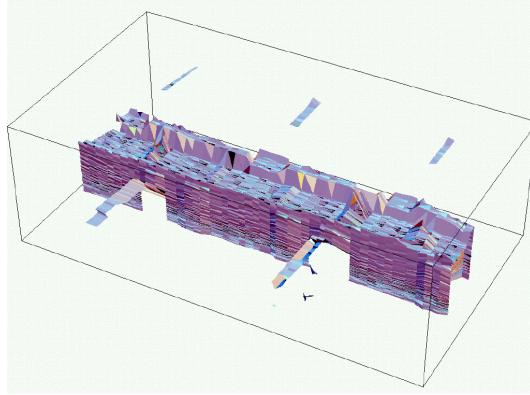


Figure 18: 3D Map.

Figure 20 shows synthesized views of a structural map with texture superimposed, of a hallway segment in an office building. All views in this figure are rendered from the 3D model, using a VRML viewer. Figure 20a shows a view from the outside with the texture of the inside walls. Such a view is unattainable in the physical building. The other three renderings show views from the inside. One of the most notable errors are the lids of the two trash bins, which appear to be sloped. The robot is smaller than those trash bins and therefore unable to see the top. The sloped lid, thus, is the result of the polygonal interpolation. The map in Figure 20 did *not* undergo the map simplification step above and possesses 80,470 polygons.

8 Limitations

The techniques proposed here possess a range of limitations, which shall be discussed in this section. These limitations are the result of the various necessary compromises relative to the true Bayes filter, which computes a full posterior over maps and robot poses.

First and foremost, our approach of using incremental maximum likelihood combined with posterior estimation over poses is clearly less powerful than full posterior estimation. In particular, our approach is unable to map nested cycles, since only a single posterior is maintained. As the size of the cycle increases, the time required for backwards correction will ultimately exceed the available time in an online mapping system. Hence, for very large cycles it may be necessary to delay data processing until the error correction mechanisms succeed. However, this limitation seems not to pose serious obstacles in any of the environments where our approach was applied.

We also believe that our positive results depend on the high accuracy and information-richness of laser range finders. Basically, knowing the poses of the robot at which measurements were taken is sufficient to generate accurate maps. It is un-

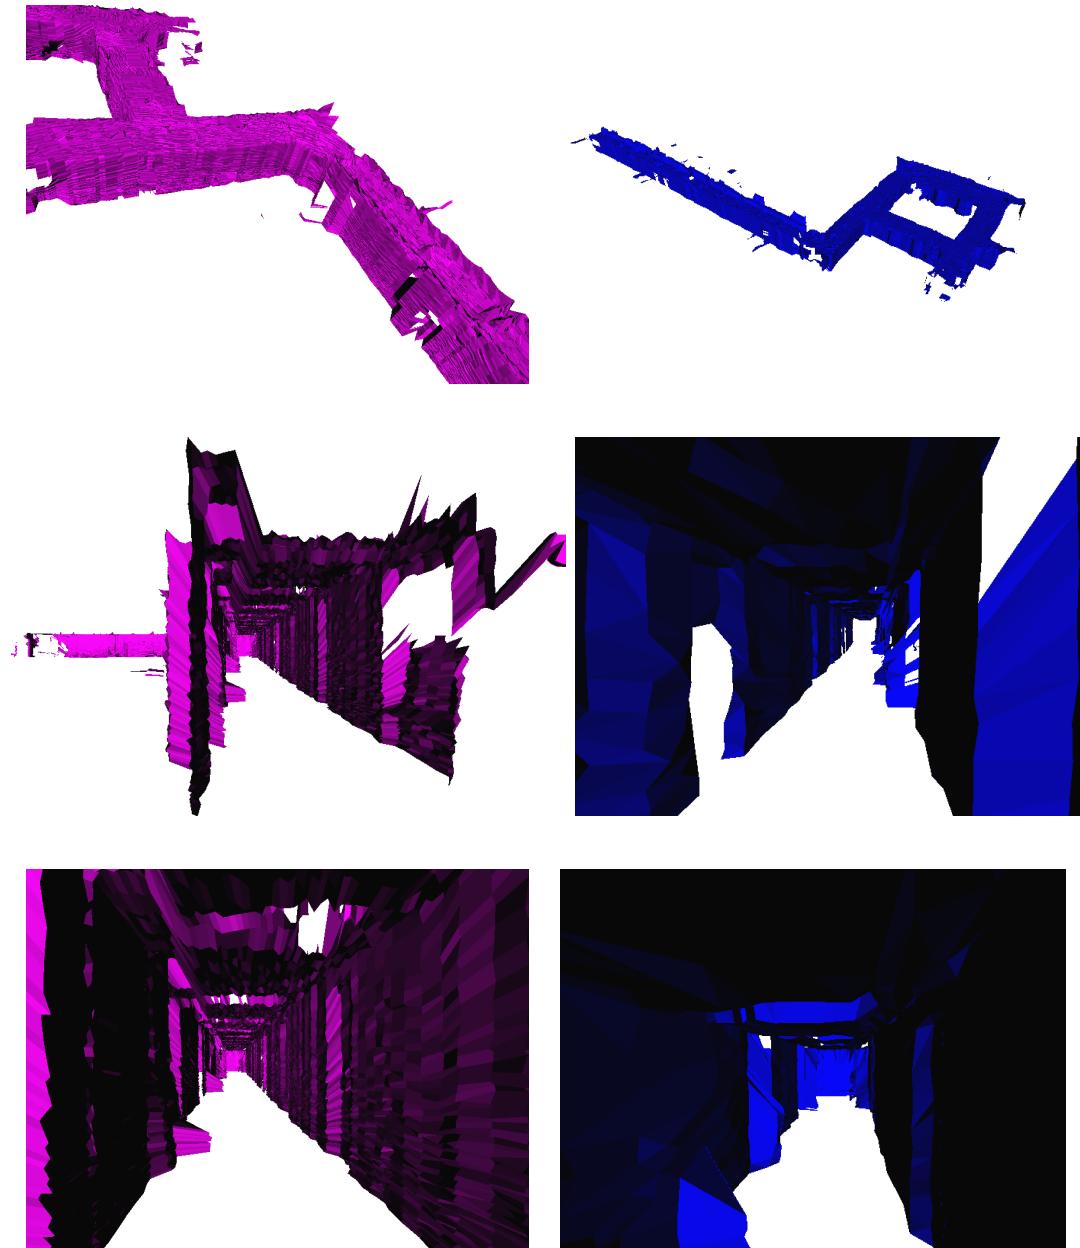


Figure 19: Views of the 3D map, for the high-resolution model (left column) and the lower resolution model (right column).

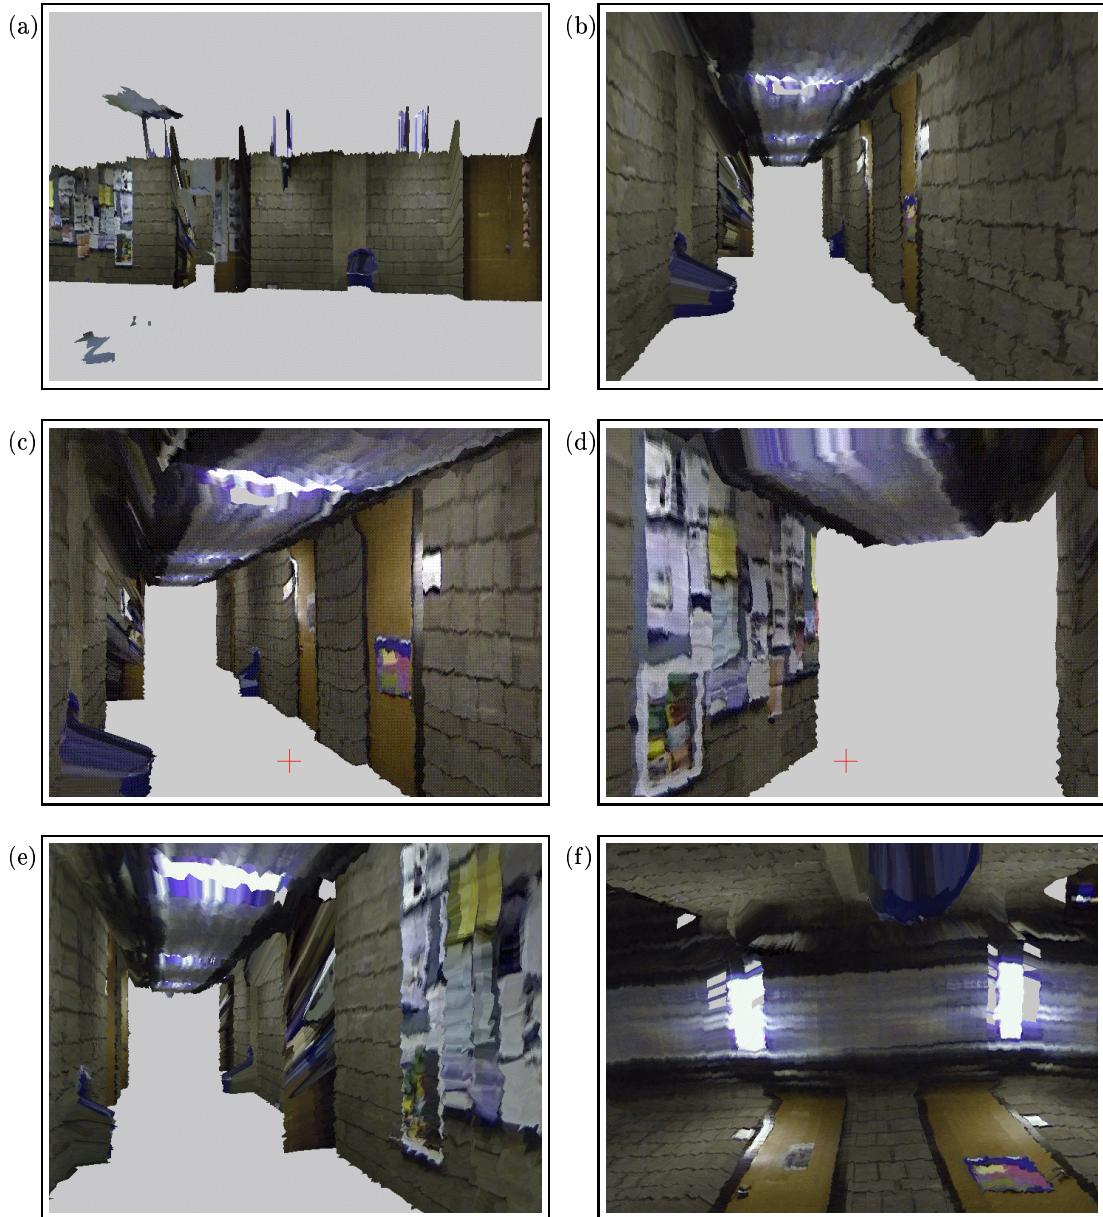


Figure 20: 3D Structural maps with texture. These views have been synthesized using a VRML viewer, from a full-blown polygonal texture model in 3D coordinates. The original renderings are in color and can be viewed at <http://www.cs.cmu.edu/~thrun/3D>.

clear how our approach would degrade if the quality of the measurements was much lower (e.g., sonar sensors or underwater imagery).

In the multi-robot mapping approach we already discussed the restrictive assumptions that the robots must know their relative poses to each other. The most general case, where robots might not even operate in the same part of the environment, is hard and, to our knowledge, a completely open problem. In particular, if robots are placed in arbitrary places without knowledge of their relative pose, they face the difficult decision whether two similarly looking local maps correspond to the same location. Moreover, substantial time might pass by before different local maps can be integrated into a global map, making it difficult to perform mapping online.

Finally, we notice that the 3D mapping approach does not take advantage of the opportunity to use 3D data for pose estimation. It only uses the horizontal laser for pose estimation. As argued above, when traversing the environment once the upward-pointed laser never sees the same object twice, hence the data becomes relatively useless for localization. However, we currently lack a sound approach for integrating data and estimating poses if the same location is traversed more than once.

We also notice that our current texture mapping implementation does not run online, due to the lack of an appropriate frame grabber on our robot. However, we do not view this as a principle limitation of the mapping approach presented here. A recent extension that uses the EM algorithm to generate flat-surface maps can be found in [53].

9 Related Work

The issue of building maps with mobile robots has received considerable attention over the past years. Virtually all state-of-the-art algorithms are probabilistic and can be derived from the Bayes filter for maps, Equation (12).

To date, the most successful attempt to implement the full posterior calculation, as specified in Equation (12), is the SLAM family of algorithms, which employ Kalman filters [9, 10, 49, 50, 61, 62]. This approach goes back to a seminal paper by Smith, Self, and Cheeseman [82]. Recall that Kalman filters represent beliefs by Gaussians; thus, they require $O(N^2)$ parameters to represent the posterior over an N -dimensional space. Calculating the Kalman filter-equivalent of Equation (12) involves matrix multiplication, which can be done in $O(N^2)$ time [58]. This critically limits the number of features that can be mapped (see [51] for a recent attempt to escape this limitation using hierarchies of maps). In practice, this approach has been applied to mapping several hundreds of free parameters [49]. A second limitation arises from the Gaussian noise assumption in Kalman filters that was already discussed at several places in this article. In the context of mapping, unimodal noise can only be achieved if features in the environment can be identified uniquely. The

problem of feature identification is known as the *data association problem*, or the *correspondence problem*. Existing Kalman filter implementations often use ad hoc methods to identify correspondences, and they are brittle if features are confused. This contrasts the work presented here, which does not assume identifiable features. In particular, the maximum likelihood estimator calculates the correspondence of objects in the map and objects detected in the most recent laser scan along the way, enabling our approach to cope with highly ambiguous environments. However, our approach lacks the ability to estimate the full uncertainty of maps and poses, which is a key feature of the Kalman filter approach.

There exist recent extensions of the basic paradigm that rely on Gaussian mixtures of posteriors, instead of a single Gaussian. This leads to an algorithm capable of handling ambiguities in the data association [14, 24]. This work continues a recent trend found in mobile robot localization, where other authors successfully applied similar mixture representations [39, 70, 71]. If this line of research proves successful, it might become a viable alternative to the approach proposed in this paper.

As discussed in the introduction to this paper, a second class of mapping algorithms is based on EM. EM is an algorithm for maximum likelihood estimation that iterates different estimation steps. EM algorithms have been applied for topological mapping [76, 75] and for building metric maps [88]. The latter maps are more related to the maps presented here. EM mapping differs from the maximum likelihood estimator discussed in this paper in that it estimates the entire map all at once, along with all poses. This provides EM with a sound approach to likelihood estimation, which uses all data for estimating the poses and the map. In contrast, the incremental maximum likelihood algorithm uses only the data leading up to a specific pose for pose estimation. However, EM is also inherently batch, and estimating a map with EM requires multiple passes through the entire data set. The algorithm presented here is incremental.

Finally, as pointed out above, there exists a range of algorithms that are variants of the incremental maximum likelihood method described above [26, 46, 47, 56, 65, 89, 91, 92, 93]. Viewed in this light, our maximum likelihood estimator is strikingly similar to a large corpus of existing literature. The key difference between our work and this literature is the integration of posterior pose estimation, which enables our approach to close cycles.

Probably the most related research is that of Gutmann and Konolige [33]. Similar to our algorithm, their approach integrates a maximum likelihood mapping algorithm with a posterior estimator. The posterior estimator is used to identify and correct errors when closing a cycle. Like us, Gutmann and Konolige found that this idea is very robust in practice, leading to highly accurate maps in cyclic environments. The differences between their and our work are on the technical side: The posterior estimator in the current approach is realized using particle filters, whereas Gutmann and Konolige rely on grid-based estimation methods similar to the ones described in [29]. Konolige and colleagues recently extended their approach

to multi-robot mapping, producing similar results to the ones shown here [45].

The vast majority of mapping algorithms in the robotics literature operates in 2D. Recently, Moravec and Martin extended the occupancy grid technique for environment mapping to three-dimensional grids [60]. Moravec and Martin use stereo vision as primary sensor, which makes the extraction of accurate depth information a challenging endeavor. He successfully demonstrated that accurate occupancy grid maps could be learned given sufficient time and memory. His approach requires accurate pose estimation, which in his experiments were provided manually. However, it appears that algorithms like the one reported here could be married with 3D occupancy grids to generate accurate maps even if the robot odometry is poor. A major difference between Moravec and Martin’s work and the work reported here is the nature of the 3D maps. Our representation is (non-probabilistic) polygons, with higher texture resolution than Moravec and Martin’s occupancy grids. The use of polygons makes it easier to render views of the map. However, our approach lacks a methodology for integrating multiple noisy measurements when traversing the same location twice.

In the area of computer vision, several researchers have studied the topic of 3D scene reconstruction. Approaches for 3D modeling can be divided into two categories: Approaches that assume knowledge of the pose of the sensors [1, 2, 3, 17, 78], and approaches that do not [35]. Our approach uses mobile robots for data acquisition; hence our approach falls into the second category due to the inherent noise in robot odometry (even after pose estimation). The majority of existing systems also require human input in the 3D modeling process. Our approach is somewhat related to [37], which reconstructs planar models of indoor environments using stereo vision, using some manual guidance in the reconstruction process to account for the lack of visible structure in typical indoor environments. Related work on outdoor terrain modeling can be found in [12].

10 Discussion

In this paper, we proposed an online algorithm for mapping large environments with mobile robots. At its core is the idea of incrementally generating a single maximum likelihood map, paired with a probabilistic estimator of the residual pose uncertainty. The pose uncertainty is used to identify and eliminate inconsistencies when learning maps in cyclic environments. As a result, even large errors that otherwise might make it impossible to close a cycle can be compensated. The resulting algorithm can generate maps of cyclic indoor environments online. Empirical results suggest that the resulting algorithm is highly efficient and robust in the face of sensor noise. The resulting maps are detailed and accurate.

Building on the basic mapping algorithm, we proposed an extension that enables teams of robots to integrate data into a single, global map in a computationally scalable fashion. Empirical results obtained with up to three simultaneously exploring

robots illustrate that our approach is well suited to generate maps with teams of robots, as long as their relative initial pose is known. We also described a method for constructing three-dimensional maps of indoor environments. These maps are generated using laser range data and visual texture information obtained with a panoramic camera. An implementation of this idea led to detailed 3D maps represented in VRML, which enables people to interactively inspect a building from remote locations.

The discussion of limitations in Section 8 leads to a range of research opportunities. On the statistical side, it appears that our combined estimator yields good results when compared with a (hypothetical) algorithm that calculates the full posterior over maps and poses. It certainly produces maps that are sufficiently accurate for robot navigation. However, it is unclear how the performance of our approach degrades with the inaccuracy of the sensors. For example, it is unclear if sonar sensors are sufficiently accurate to yield good results. Another topic of future research is the development of multi-robot mapping algorithms that make no assumption on the initial conditions. In particular, the assumptions that robots must know their relative initial pose is a key restriction of the current approach, thus warrants further research.

In our opinion, the most striking research opportunity arises from the notion of 3D maps. Three-dimensional maps open new opportunities for mobile robotics, in that they are meaningful and usable for people in ways that 2D maps are not. Ultimately, having access to 3D information should also facilitate a range of hard estimation problems, such as the data association problem in mapping. This is because 3D views carry more information than their 2D counterparts, making it potentially easier to disambiguate different places. At first glance, one might attempt to apply the rich, known statistical machinery for 2D mapping to the three-dimensional case. However, mapping in 3D is computationally more difficult than in 2D, opening new opportunities to generate efficient and sound mapping algorithms.

We believe the approach presented here fills an important gap in the literature. Previous indoor mapping algorithms that could cope with large errors when closing the loop were not incremental, but instead required an elaborate procedure with multiple iterations through all available data. The current approach is an online algorithm. This is specifically important for robot exploration, since exploring robots have to continuously make decisions as to where to explore next. The algorithms presented here were found to be extremely robust and reliable in a range of recent experiments.

Acknowledgment

The author would like to express his gratitude to team members of the Mercator Team and the TMR program, who contributed to this research in various ways. In particular, he would like to thank David Apfelbaum, Greg Armstrong, John Blitch,

Wolfram Burgard, Dieter Fox, Steffen Gutmann, Kurt Konolige, and Reid Simmons for inspiring discussions and feedback on this research.

This research is sponsored by the National Science Foundation (CAREER grant number IIS-9876136 and regular grant number IIS-9877033), and by DARPA-ATO via TACOM (contract number DAAE07-98-C-L032), which is gratefully acknowledged. The views and conclusions contained in this document are those of the author and should not be interpreted as necessarily representing official policies or endorsements, either expressed or implied, of the United States Government or any of the sponsoring institutions.

References

- [1] P.K. Allen and Ioannis Stamos. Integration of range and image sensing for photorealistic 3D modeling. In *Proceedings of the IEEE International Conference on Robotics and Automation (ICRA)*, pages 1435–1440, 2000.
- [2] R. Bajcsy, G. Kamberova, and Lucien Nocera. 3D reconstruction of environments for virtual reconstruction. In *Proc. of the 4th IEEE Workshop on Applications of Computer Vision*, 2000.
- [3] S. Becker and M. Bove. Semiautomatic 3-D model extraction from uncalibrated 2-D camera views. In *Proc. of the SPIE Symposium on Electronic Imaging, San Jose*, 1995.
- [4] J. Borenstein, B. Everett, and L. Feng. *Navigating Mobile Robots: Systems and Techniques*. A. K. Peters, Ltd., Wellesley, MA, 1996.
- [5] W. Burgard, A.B. Cremers, D. Fox, D. Hähnel, G. Lakemeyer, D. Schulz, W. Steiner, and S. Thrun. Experiences with an interactive museum tour-guide robot. *Artificial Intelligence*, 114(1-2):3–55, 1999.
- [6] W. Burgard, D. Fox, D. Hennig, and T. Schmidt. Estimating the absolute position of a mobile robot using position probability grids. In *Proceedings of the Thirteenth National Conference on Artificial Intelligence*, Menlo Park, August 1996. AAAI, AAAI Press/MIT Press.
- [7] W. Burgard, D. Fox, H. Jans, C. Matenar, and S. Thrun. Sonar-based mapping of large-scale mobile robot environments using EM. In *Proceedings of the International Conference on Machine Learning*, Bled, Slovenia, 1999.
- [8] W. Burgard, D. Fox, M. Moors, R. Simmons, and S. Thrun. Collaborative multi-robot exploration. In *Proceedings of the IEEE International Conference on Robotics and Automation (ICRA)*, San Francisco, CA, 2000. IEEE.
- [9] J.A. Castellanos, J.M.M. Montiel, J. Neira, and J.D. Tardós. The SPmap: A probabilistic framework for simultaneous localization and map building. *IEEE Transactions on Robotics and Automation*, 15(5):948–953, 1999.
- [10] J.A. Castellanos and J.D. Tardós. *Mobile Robot Localization and Map Building: A Multisensor Fusion Approach*. Kluwer Academic Publishers, Boston, MA, 2000.
- [11] R. Chatila and J.-P. Laumond. Position referencing and consistent world modeling for mobile robots. In *Proceedings of the 1985 IEEE International Conference on Robotics and Automation*, 1985.
- [12] Y.-Q. Cheng, E. Riseman, X. Wang, R. Collins, and A. Hanson. Three-dimensional reconstruction of points and lines with unknown correspondence across images. *International Journal of Computer Vision*, 2000.
- [13] H. Choset. *Sensor Based Motion Planning: The Hierarchical Generalized Voronoi Graph*. PhD thesis, California Institute of Technology, 1996.

- [14] I.J. Cox and J.J. Leonard. Modeling a dynamic environment using a bayesian multiple hypothesis approach. *Artificial Intelligence*, 66:311–344, 1994.
- [15] T. Dean and R.P. Bonasso. The 1992 AAAI robot exhibition and competition. *AI Magazine*, 14(1):35–48, Spring 1993.
- [16] T.L. Dean and K. Kanazawa. A model for reasoning about persistence and causation. *Computational Intelligence*, 5(3):142–150, 1989.
- [17] P.E. Debevec, C.J. Taylor, and J. Malik. Modeling and rendering architecture from photographs. In *Proc. of the 23rd International Conference on Computer Graphics and Interactive Techniques (SIGGRAPH)*, 1996.
- [18] F. Dellaert, W. Burgard, D. Fox, and S. Thrun. Using the condensation algorithm for robust, vision-based mobile robot localization. In *Proceedings of the IEEE International Conference on Computer Vision and Pattern Recognition*, Fort Collins, CO, 1999. IEEE.
- [19] F. Dellaert, S.M. Seitz, C. Thorpe, and S. Thrun. EM, MCMC, and chain flipping for structure from motion with unknown correspondence. *Machine Learning*, 2000. To appear.
- [20] A.P. Dempster, A.N. Laird, and D.B. Rubin. Maximum likelihood from incomplete data via the EM algorithm. *Journal of the Royal Statistical Society, Series B*, 39(1):1–38, 1977.
- [21] G. Dissanayake, H. Durrant-Whyte, and T. Bailey. A computationally efficient solution to the simultaneous localisation and map building (SLAM) problem. Working notes of ICRA'2000 Workshop W4: Mobile Robot Navigation and Mapping, April 2000.
- [22] A Doucet. On sequential simulation-based methods for bayesian filtering. Technical Report CUED/F-INFENG/TR 310, Cambridge University, Department of Engineering, Cambridge, UK, 1998.
- [23] A. Doucet, J.F.G. de Freitas, and N.J. Gordon, editors. *Sequential Monte Carlo Methods In Practice*. Springer Verlag, New York, 2001.
- [24] H. Durrant-Whyte, S. Majumder, S. Thrun, M. de Battista, and S. Scheding. A bayesian algorithm for simultaneous localization and map building. Submitted for publication, 2001.
- [25] A. Elfes. *Occupancy Grids: A Probabilistic Framework for Robot Perception and Navigation*. PhD thesis, Department of Electrical and Computer Engineering, Carnegie Mellon University, 1989.
- [26] S. Engelson and D. McDermott. Error correction in mobile robot map learning. In *Proceedings of the 1992 IEEE International Conference on Robotics and Automation*, pages 2555–2560, Nice, France, May 1992.
- [27] D. Fox, W. Burgard, F. Dellaert, and S. Thrun. Monte carlo localization: Efficient position estimation for mobile robots. In *Proceedings of the National Conference on Artificial Intelligence (AAAI)*, Orlando, FL, 1999. AAAI.
- [28] D. Fox, W. Burgard, H. Kruppa, and S. Thrun. A probabilistic approach to collaborative multi-robot localization. *Autonomous Robots*, 8(3), 2000.
- [29] D. Fox, W. Burgard, and S. Thrun. Markov localization for mobile robots in dynamic environments. *Journal of Artificial Intelligence Research*, 11:391–427, 1999.
- [30] M. Garland and P. Heckbert. Surface simplification using quadric error metrics. In *Proceedings of SIGGRAPH*, 1997.
- [31] M. Garland and P. Heckbert. Simplifying surfaces with color and texture using quadric error metrics. In *Proceedings IEEE Visualization*, 1998.
- [32] J.-S. Gutmann. Vergleich von algorithmen zur selbstlokalisierung eines mobilen roboters. Master's thesis, University of Ulm, Ulm, Germany, 1996. (in German).

- [33] J.-S. Gutmann and K. Konolige. Incremental mapping of large cyclic environments. In *Proceedings of the IEEE International Symposium on Computational Intelligence in Robotics and Automation (CIRA)*, 2000.
- [34] J.-S. Gutmann and C. Schlegel. Amos: Comparison of scan matching approaches for self-localization in indoor environments. In *Proceedings of the 1st Euromicro Workshop on Advanced Mobile Robots*. IEEE Computer Society Press, 1996.
- [35] S.E. Hakim and P. Boulanger. Sensor based creation of indoor virtual environment models. In *Proc. of the 4th International Conference on Virtual Systems and Multimedia (VSMM)*, 1997.
- [36] I. Horswill. Specialization of perceptual processes. Technical Report AI TR-1511, MIT, AI Lab, Cambridge, MA, September 1994.
- [37] L. Iocchi, K. Konolige, and M. Bajracharya. Visually realistic mapping of a planar environment with stereo. In *Proceedings of the 2000 International Symposium on Experimental Robotics*, Waikiki, Hawaii, 2000.
- [38] M. Isard and A. Blake. Condensation: conditional density propagation for visual tracking. *International Journal of Computer Vision*, 29(1):5–28, 1998. In press.
- [39] P. Jensfelt and S. Kristensen. Active global localisation for a mobile robot using multiple hypothesis tracking. In *Proceedings of the IJCAI Workshop on Reasoning with Uncertainty in Robot Navigation*, pages 13–22, Stockholm, Sweden, 1999. IJCAI.
- [40] L.P. Kaelbling, A.R. Cassandra, and J.A. Kurien. Acting under uncertainty: Discrete bayesian models for mobile-robot navigation. In *Proceedings of the IEEE/RSJ International Conference on Intelligent Robots and Systems*, 1996.
- [41] R. E. Kalman. A new approach to linear filtering and prediction problems. *Trans. ASME, Journal of Basic Engineering*, 82:35–45, 1960.
- [42] K. Kanazawa, D. Koller, and S.J. Russell. Stochastic simulation algorithms for dynamic probabilistic networks. In *Proceedings of the 11th Annual Conference on Uncertainty in AI*, Montreal, Canada, 1995.
- [43] S. King and C. Weiman. Helpmate autonomous mobile robot navigation system. In *Proceedings of the SPIE Conference on Mobile Robots*, pages 190–198, Boston, MA, November 1990. Volume 2352.
- [44] D. Koller and R. Fratkina. Using learning for approximation in stochastic processes. In *Proceedings of the International Conference on Machine Learning (ICML)*, 1998.
- [45] K. Konolige, J.-S. Gutmann, D. Guzzoni, R. Ficklin, and K. Nicewarner. A mobile robot sense net. In *Proceedings of SPIE 3839 Sensor Fusion and Decentralized Control in Robotic Systems II*, Boston, September 1999.
- [46] D. Kortenkamp and T. Weymouth. Topological mapping for mobile robots using a combination of sonar and vision sensing. In *Proceedings of the Twelfth National Conference on Artificial Intelligence*, pages 979–984, Menlo Park, July 1994. AAAI, AAAI Press/MIT Press.
- [47] B. Kuipers and Y.-T. Byun. A robot exploration and mapping strategy based on a semantic hierarchy of spatial representations. *Journal of Robotics and Autonomous Systems*, 8:47–63, 1991.
- [48] S. Lenser and M. Veloso. Sensor resetting localization for poorly modelled mobile robots. In *Proceedings of the IEEE International Conference on Robotics and Automation (ICRA)*, San Francisco, CA, 2000. IEEE.
- [49] J. J. Leonard and H. F. Durrant-Whyte. *Directed Sonar Sensing for Mobile Robot Navigation*. Kluwer Academic Publishers, Boston, MA, 1992.
- [50] J.J. Leonard, H.F. Durrant-Whyte, and I.J. Cox. Dynamic map building for an autonomous mobile robot. *International Journal of Robotics Research*, 11(4):89–96, 1992.

- [51] J.J. Leonard and H.J.S. Feder. A computationally efficient method for large-scale concurrent mapping and localization. In J. Hollerbach and D. Koditschek, editors, *Proceedings of the Ninth International Symposium on Robotics Research*, Salt Lake City, Utah, 1999.
- [52] J. Liu and R. Chen. Sequential monte carlo methods for dynamic systems. *Journal of the American Statistical Association*, 93:1032–1044, 1998.
- [53] Y. Liu, R. Emery, D. Chakrabarti, W. Burgard, and S. Thrun. Using EM to learn 3D models with mobile robots. In *Proceedings of the International Conference on Machine Learning (ICML)*, 2001.
- [54] F. Lu and E. Milios. Globally consistent range scan alignment for environment mapping. *Autonomous Robots*, 4:333–349, 1997.
- [55] F. Lu and E. Milios. Robot pose estimation in unknown environments by matching 2d range scans. *Journal of Intelligent and Robotic Systems*, 18:249–275, 1998.
- [56] M. J. Matarić. A distributed model for mobile robot environment-learning and navigation. Master’s thesis, MIT, Cambridge, MA, January 1990. also available as MIT AI Lab Tech Report AITR-1228.
- [57] M.J. Matarić. Interaction and intelligent behavior. Technical Report AI-TR-1495, Massachusetts Institute of Technology, Artificial Intelligence Laboratory, Cambridge, MA, 1994.
- [58] P.S. Maybeck. The Kalman filter: An introduction to concepts. In I.J. Cox and G.T. Wilfong, editors, *Autonomous Robot Vehicles*. Springer verlag, 1990.
- [59] H. P. Moravec. Sensor fusion in certainty grids for mobile robots. *AI Magazine*, pages 61–74, Summer 1988.
- [60] H.P. Moravec and M.C. Martin. Robot navigation by 3D spatial evidence grids. Mobile Robot Laboratory, Robotics Institute, Carnegie Mellon University, 1994.
- [61] P. Moutarlier and R. Chatila. An experimental system for incremental environment modeling by an autonomous mobile robot. In *1st International Symposium on Experimental Robotics*, Montreal, June 1989.
- [62] P. Moutarlier and R. Chatila. Stochastic multisensory data fusion for mobile robot location and environment modeling. In *5th Int. Symposium on Robotics Research*, Tokyo, 1989.
- [63] I. Nourbakhsh, J. Bobenage, S. Grange, R. Lutz, R. Meyer, and A. Soto. An affective mobile robot with a full-time job. *Artificial Intelligence*, 114(1–2):95–124, 1999.
- [64] I. Nourbakhsh, R. Powers, and S. Birchfield. DERVISH an office-navigating robot. *AI Magazine*, 16(2):53–60, Summer 1995.
- [65] D. Pierce and B. Kuipers. Learning to explore and build maps. In *Proceedings of the Twelfth National Conference on Artificial Intelligence*, pages 1264–1271, Menlo Park, July 1994. AAAI, AAAI Press/MIT Press.
- [66] M. Pitt and N. Shephard. Filtering via simulation: auxiliary particle filter. *Journal of the American Statistical Association*, 94:590–599, 1999.
- [67] L. R. Rabiner. A tutorial on hidden markov models and selected applications in speech recognition. In *Proceedings of the IEEE*. IEEE, 1989. IEEE Log Number 8825949.
- [68] L.R. Rabiner and B.H. Juang. An introduction to hidden markov models. *IEEE ASSP Magazine*, 3(1):4–16, 1986.
- [69] W.D. Rencken. Concurrent localisation and map building for mobile robots using ultrasonic sensors. In *Proceedings of the IEEE/RSJ International Conference on Intelligent Robots and Systems*, pages 2129–2197, Yokohama, Japan, July 1993.
- [70] J. Reuter. Mobile robot self-localization using PDAB. In *Proceedings of the IEEE International Conference on Robotics and Automation (ICRA)*, San Francisco, CA, 2000. IEEE.

[71] S.I. Roumeliotis and G.A. Bekey. Bayesian estimation and kalman filtering: A unified framework for mobile robot localization. In *Proceedings of the IEEE International Conference on Robotics and Automation (ICRA)*, pages 2985–2992, San Francisco, CA, 2000. IEEE.

[72] D.B. Rubin. Using the SIR algorithm to simulate posterior distributions. In M.H. Bernardo, K.M. an DeGroot, D.V. Lindley, and A.F.M. Smith, editors, *Bayesian Statistics 3*. Oxford University Press, Oxford, UK, 1988.

[73] S. Russell and P. Norvig. *Artificial Intelligence: A Modern Approach*. Prentice Hall, Englewood Cliffs, NJ, 1995.

[74] B. Schiele and J. Crowley. A comparison of position estimation techniques using occupancy grids. In *Proceedings of the 1994 IEEE International Conference on Robotics and Automation*, pages 1628–1634, San Diego, CA, May 1994.

[75] H. Shatkay. *Learning Models for Robot Navigation*. PhD thesis, Computer Science Department, Brown University, Providence, RI, 1998.

[76] H. Shatkay and L. Kaelbling. Learning topological maps with weak local odometric information. In *Proceedings of IJCAI-97*. IJCAI, Inc., 1997.

[77] H. Shatkay and L.P. Kaelbling. Learning hidden markov models with geometric information. Technical Report CS-97-04, Computer Science Department, Brown University, Providence, RI, April 1997.

[78] H. Shum, M. Han, and R. Szeliski. Interactive construction of 3D models from panoramic mosaics. In *Proc. of the International Conference on Computer Vision and Pattern Recognition (CVPR)*, 1998.

[79] R. Simmons. Where in the world is xavier, the robot? *Machine Perception*, 5(1), 1996.

[80] R. Simmons, D. Apfelbaum, W. Burgard, M. Fox, D. an Moors, S. Thrun, and H. Younes. Coordination for multi-robot exploration and mapping. In *Proceedings of the AAAI National Conference on Artificial Intelligence*, Austin, TX, 2000. AAAI.

[81] R. Simmons and S. Koenig. Probabilistic robot navigation in partially observable environments. In *Proceedings of IJCAI-95*, pages 1080–1087, Montreal, Canada, August 1995. IJCAI, Inc.

[82] R. Smith, M. Self, and P. Cheeseman. Estimating uncertain spatial relationships in robotics. In I.J. Cox and G.T. Wilfong, editors, *Autonomous Robot Vehnicles*, pages 167–193. Springer-Verlag, 1990.

[83] M.A. Tanner. *Tools for Statistical Inference*. Springer Verlag, New York, 1993. 2nd edition.

[84] S. Thrun. Exploration and model building in mobile robot domains. In E. Ruspini, editor, *Proceedings of the IEEE International Conference on Neural Networks*, pages 175–180, San Francisco, CA, 1993. IEEE Neural Network Council.

[85] S. Thrun. Probabilistic algorithms in robotics. *AI Magazine*, 21(4):93–109, 2000.

[86] S. Thrun, M. Beetz, M. Bennewitz, W. Burgard, A.B. Cremers, F. Dellaert, D. Fox, D. Hähnel, C. Rosenberg, N. Roy, J. Schulte, and D. Schulz. Probabilistic algorithms and the interactive museum tour-guide robot minerva. *International Journal of Robotics Research*, 19(11):972–999, 2000.

[87] S. Thrun, W. Burgard, and D. Fox. A real-time algorithm for mobile robot mapping with applications to multi-robot and 3D mapping. In *Proceedings of the IEEE International Conference on Robotics and Automation (ICRA)*, San Francisco, CA, 2000. IEEE.

[88] S. Thrun, D. Fox, and W. Burgard. A probabilistic approach to concurrent mapping and localization for mobile robots. *Machine Learning*, 31:29–53, 1998. also appeared in *Autonomous Robots* 5, 253–271.

- [89] M. C. Torrance. Natural communication with robots. Master's thesis, MIT Department of Electrical Engineering and Computer Science, Cambridge, MA, January 1994.
- [90] G. Weiß and E.v. Puttkamer. A map based on laserscans without geometric interpretation. In *Proceedings of Intelligent Autonomous Systems, Vol. 4 (IAS-4)*, pages 403–407. IOS Press, Karlsruhe, Germany, 1995.
- [91] B. Yamauchi and R. Beer. Spatial learning for navigation in dynamic environments. *IEEE Transactions on Systems, Man, and Cybernetics - Part B: Cybernetics*, Special Issue on Learning Autonomous Robots, 1996. also located at <http://www.aic.nrl.navy.mil/~yamauchi/>.
- [92] B. Yamauchi, P. Langley, A.C. Schultz, J. Grefenstette, and W. Adams. Magellan: An integrated adaptive architecture for mobile robots. Technical Report 98-2, Institute for the Study of Learning and Expertise (ISLE), Palo Alto, CA, May 1998.
- [93] U.R. Zimmer. Robust world-modeling and navigation in a real world. *Neurocomputing*, 13(2–4), 1996.