

Towards Robust Data Association and Feature Modeling for Concurrent Mapping and Localization

John J. Leonard, Paul M. Newman, Richard J. Rikoski
MIT Dept. of Ocean Engineering
Cambridge, MA 02139, U.S.A

José Neira, Juan. D. Tardós
Universidad de Zaragoza
50015 Zaragoza, Spain

Abstract

One of the most challenging aspects of concurrent mapping and localization (CML) is the problem of data association. Because of uncertainty in the origins of sensor measurements, it is difficult to determine the correspondence between measured data and features of the scene or object being observed, while rejecting spurious measurements. However, there are many important applications of mobile robots where maps need to be built of complex environments, consisting of composite features, from noisy sensor data. This paper reviews several new approaches to data association and feature modeling for CML that share the common theme of combining information from multiple uncertain vantage points while rejecting spurious data. Our results include: (1) feature-based mapping from laser data using robust segmentation, (2) map-building with sonar data using a novel application of the Hough transform for perception grouping, and (3) a new stochastic framework for making delayed decisions for combination of data from multiple uncertain vantage points. Experimental results are shown for CML using laser and sonar data from a B21 mobile robot.

1 Introduction

The problem of concurrent mapping and localization (CML) for an autonomous mobile robot is stated as follows: starting from an initial position, a mobile robot travels through a sequence of positions and obtains a set of sensor measurements at each position. The goal is for the mobile robot to process the sensor data to produce an estimate of its position while concurrently building a map of the environment. While the problem of CML is deceptively easy to state, it presents many theoretical challenges. The problem is also of great practical importance; if a robust, general-purpose solu-

tion to CML can be found, then many new applications of mobile robotics will become possible.

CML, also referred to as SLAM (simultaneous localization and map building), has been a recurring theme at the series of ISRR Symposia over the years [1, 2, 3, 4]. For example, in his paper for the second ISRR symposium, Brooks [2] was among the first to suggest that a probabilistic approach was necessary to develop robust algorithms for mapping and navigation:

“Mobile robots sense their environment and receive error laden readings. They try to move a certain distance and direction, only to do so approximately. Rather than try to engineer these problems away it may be possible, and may be necessary, to develop map mapping and navigation algorithms which explicitly represent these uncertainties, but still provide robust information. [2]”

The key technical difficulty in performing CML is coping with uncertainty. Three distinct forms of uncertainty – data association uncertainty, navigation error, and sensor noise – work together to present a challenging data interpretation problem. For example, Figures 1 and 2 show the laser and sonar data, respectively, collected by a B21 mobile robot during several back-and-forth traverses of a corridor a few tens of meters long. Figure 3 shows the accumulation of dead-reckoning error during a longer duration traverse of about 500 meters in the MIT “infinite corridor” (shown in Figure 22).

Most successful recent implementations of CML have either been performed with SICK laser scanner data [5, 6] or in environments that consist of isolated “point” objects [7, 8]. However, there are many important applications of mobile robots where maps need to be built of complex environments, consisting of com-



Figure 1: Laser data for a short corridor experiment, referenced to the dead-reckoning position estimate

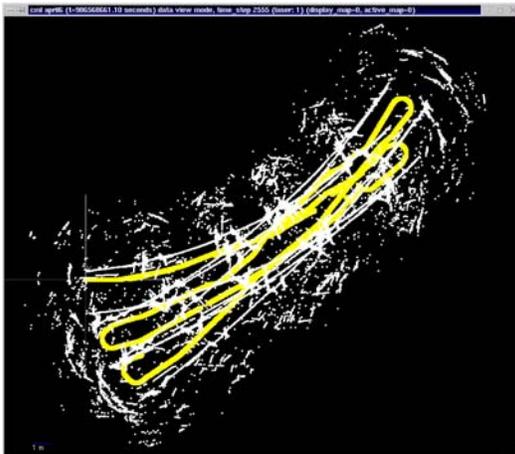


Figure 2: Sonar data for a short corridor experiment, referenced to the dead-reckoning position estimate

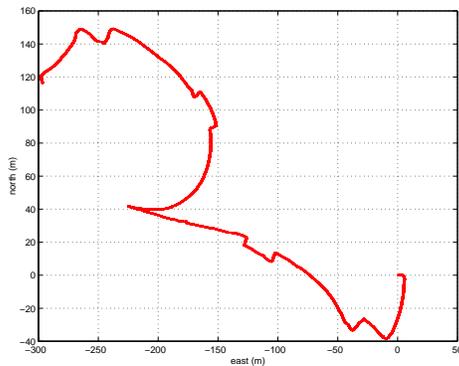


Figure 3: Accumulation of position error relying only on dead-reckoning for a long distance traverse of the B21 mobile robot. The actual vehicle path went down approximately 40 meters, to the left approximately 225 meters, and then back to the origin.

posite features, from noisy sensor data. The goal of our work is to enable autonomous underwater vehicles to navigate autonomously using sonar. Current methods for data association in feature-based CML are unable to cope with sonar because of its sparse and ambiguous nature.

Thrun et al. [5] and Gutmann et al. [6] have developed implementations of CML using laser data that are capable of closing moderately sized loops in real-time. In their work, the representation consists of “raw” sensor data referenced back to a complete trajectory of the vehicle. With this representation, they are able to greatly simplify the data association problem. CML algorithms that use a feature-based representation must explicitly solve the data association problem for each sensor measurement. Given a new sensor measurement, does it correspond to a previously mapped feature, a new feature that should be mapped, or is it spurious and should be ignored?

A key benefit of the SICK laser scanner is that the data from one position can be directly correlated with data taken from a nearby position, to compute the offset in robot position between the two positions. With sonar, the raw data is usually too noisy and ambiguous for this type of approach to work.

Recent work in feature-based CML has shown the importance of maintaining spatial correlations to achieve consistent error bounds [7, 8]. The representation of spatial correlations results in an $O(n^2)$ growth in computational cost [4], motivating techniques to address the map scaling problem through spatial and temporal partitioning [9, 10, 11]. Almost all implementations of feature-based CML to-date have used fairly simple nearest-neighbor gating techniques. A more powerful technique that tests the Joint Compatibility testing of multiple sensor measurements, using a branch and bound algorithm, has been developed by Neira and Tardós [12].

In this paper, we present results from several different new implementations of CML using either sonar or laser data. The results demonstrate feature classification and mapping from multiple uncertain vantage points. Section 2 presents results from a real-time implementation of CML with laser data that uses techniques from robust statistics for line segment extraction. Section 3 presents map-building results with sonar using a novel application of the Hough transform for perception grouping. Experimental results for sonar map-building and laser map-building of the same scene are compared. Section 4 summarizes a new stochastic framework for making delayed decisions to enable combination of data from multiple uncertain vantage points. Sonar data processing results are pre-

sented. Finally, Section 5 draws some conclusions and discusses challenges for future research.

2 “Explore and return” using Laser

This section presents results from use of the CMLKernel – a new, generic, real-time implementation of feature-based CML. Novel characteristics of this implementation include: (1) a hierarchical representation of uncertain geometric relationships that extends the SPMMap framework [13], (2) use of robust statistics to perform extraction of line segments from laser data in real-time, and (3) the integration of CML with a “roadmap” path planning method for autonomous trajectory execution. These innovations are combined to demonstrate the ability for a mobile robot to autonomously return back to its starting position within a few centimeters of precision, despite the presence of numerous people walking through the environment.

The sensors used were a SICK laser scanner and wheel encoders mounted on the B21 vehicle. The floor surface was a combination of sandstone tiles and carpet mats providing alternatively high and low wheel slippage. The exploration stage was manually controlled although it should be emphasized that this was done *without* visual contact with the vehicle. The output of the CMLKernel was rendered in 3D and used as a real-time visualization tool of the robots workspace. This enabled the remote operator to “visit” previously unexplored areas while simultaneously building an accurate geometric representation of the environment. This in itself is a useful application of CML; nevertheless, future experiments will implement an autonomous explore function as well as the existing autonomous return.

To illustrate the accuracy of the CML algorithm the starting position of the robot was marked with four ten-cent coins; the robot then explored its environment and when commanded used the resulting map to return to its initial position and park itself on top of the coins with less than 2cm of error. The duration of the experiment was a little over 20 minutes long with just over 6MB of data processed. The total distance traveled was well in excess of 100m. Videos of various stages of the experiment can be found in various formats at <http://oe.mit.edu/~pnewman>.

Figure 6 shows the environment in which the experiment occurred. The main entrance hall to the MIT campus was undergoing renovation during which large wood-clad pillars had been erected throughout the hallway yielding an interesting, landmark rich and densely populated area.

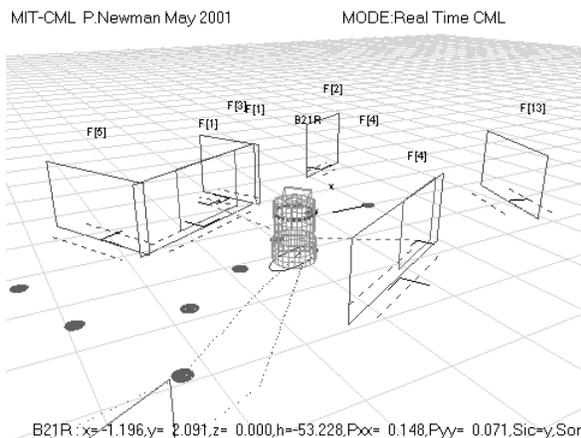


Figure 4: Re-observing an existing feature

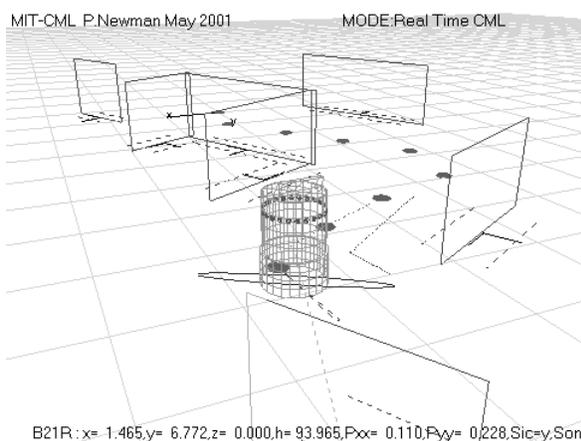


Figure 5: Creating a new feature in the foreground following a rotation



Figure 6: The experiment scene

Figures 4 and 5 show rendered views of the estimated map during the exploration phase of the experiment. In Figure 4 the robot can be seen to be applying a line segment observation of an existing feature. In contrast Figure 5 shows an observation initializing a new feature just after the robot has turned a corner. The dotted lines parallel to the walls are representations of the uncertainty of lateral uncertainty in that wall feature.

The vehicle was started with an initial uncertainty of 0.35 m and as shown in [14] all features will inherit this uncertainty as a limiting lower bound in their own uncertainty. The 1σ uncertainty of the vehicle location is shown as a dotted ellipse around the base of the vehicle.

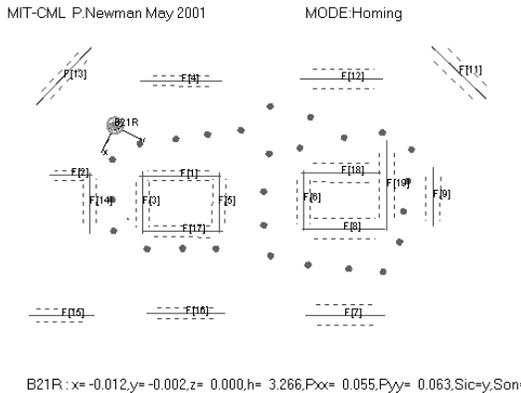


Figure 7: A plan view of the CML map at the end of the experiment. The approximate size of the environment was a 20m by 15m rectangle.

Figure 7 shows an OpenGL view of the estimated map towards the end of the experiment when the robot is executing its homing algorithm. The circles on the ground mark the free space markers that were dropped during the exploration phase of the experiment. The homing command was given when the robot was at the far corner of the hallway. Using the output of the CMLKernel, the robot set the goal marker to be the closest way point. When the algorithm deduces that the vehicle is within an acceptable tolerance ϵ of the present goal marker it sets the goal way-point to be the closest marker that has score less than the present goal marker. This then proceeds until the goal marker is the origin or initial robot position. At this point the goal seeking tolerance ϵ is reduced to 1cm. The CMLKernel spent about thirty seconds commanding small adjustments to the location and pose of the robot before declaring that the vehicle had indeed arrived back at its starting location. Figure 8 and 9 show the starting and finishing positions with respect to the coin markers. As can be seen in these figures the vehicle returned to within an inch of the starting location. Readers are invited to view videos of this experiment and others including navigation in a populated museum at <http://oe.mit.edu/~pnewman>.

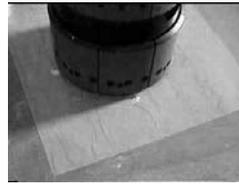


Figure 8: The starting position



Figure 9: The robot position after the completion of the homing leg of the mission

3 Sonar Perceptual Grouping Using the Hough Transform

The data from a standard ring of Polaroid sonar sensors can be notoriously difficult to interpret. This leads many researchers away from a geometric approach to sonar mapping. However, using a physics-based sensor model, the geometric constraints provided by an individual sonar return can be formulated [15]. Each return could originate from various types of features (point, plane, etc.) or could be spurious. For each type of feature, there is a limited range of locations for a potential feature that are possible. Given these constraints, the Hough transform [16] can be used as a voting scheme to identify point and planar features. More detail on this technique will be published in a future report. A somewhat related technique called triangulation-based fusion has been developed by Wijk and Christensen [17] for point objects only. The Hough transform approach is advantageous because it can directly identify specular planar reflectors from sonar data, which is vitally important in typical man-made environments with many smooth walls.

Figure 10 through 13 provide an illustrative result for this approach. The Hough transform is applied to small batches of sonar data (22 positions each) as a pre-filter to look for potential new features in the sonar data. These groupings are then fed into an implementation of CML that uses the SPMMap as the state estimation framework [13], Joint Compatibility for data association [12], and a new technique called Map Joining (to be described in a future report). Figure 14 shows a map of the same environment built from laser data. One can see that sonar map is almost as good as as well as the laser map.

4 Delayed Stochastic Mapping

This section reviews stochastic mapping and its extension to account for temporal correlations [18]. Stochastic mapping is a feature-based concurrent map-

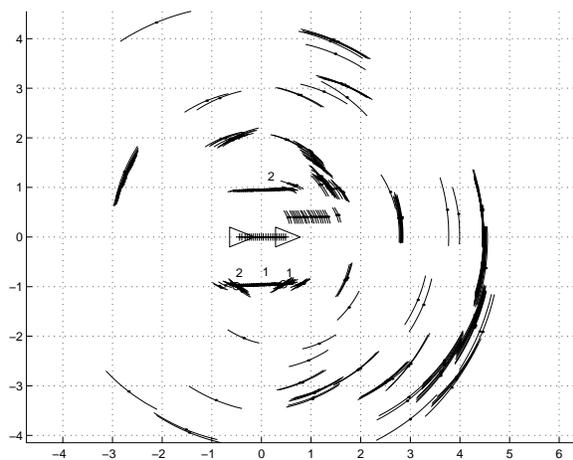


Figure 10: Example of Hough processing to extract point and line features. Sonar returns are processed in a group of twenty-two positions. A voting scheme is performed to find clusters of measurements that hypothesize the existence of point and plane features. For this example, two planes and two points have been found.

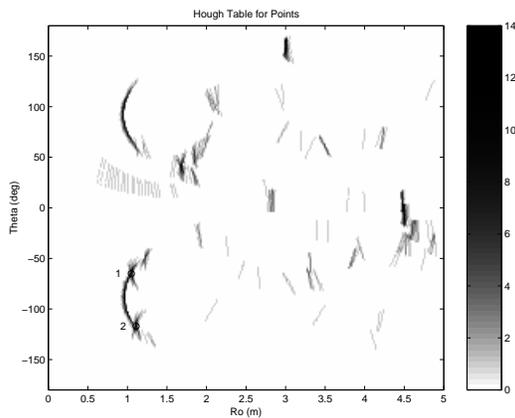


Figure 11: Hough voting table for point features.

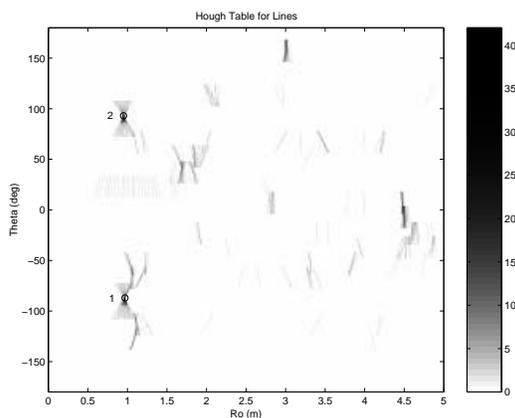


Figure 12: Hough voting table for line features.

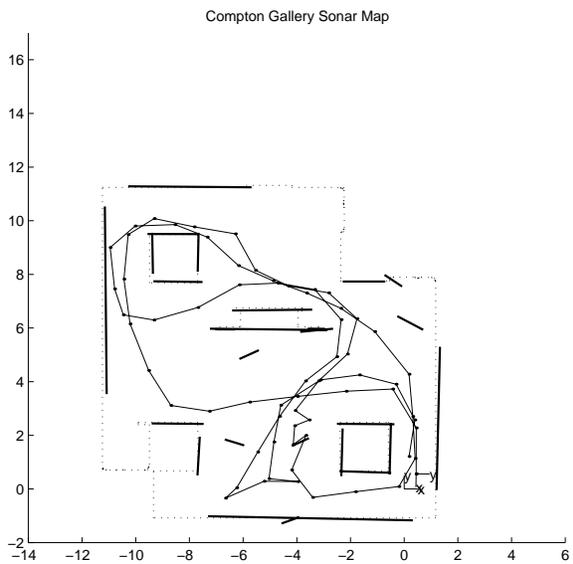


Figure 13: Complete map for the MIT Compton Gallery built from sonar using Hough grouping, Map Joining, and Joint Compatibility.

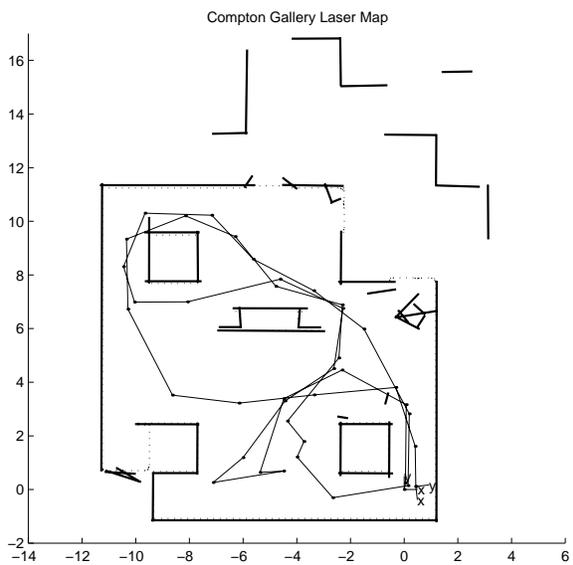


Figure 14: Complete map for the MIT Compton Gallery built from laser data using Robust Statistics, Map Joining, and Joint Compatibility.

ping and localization algorithm that was first published by Smith, Self, and Cheeseman [19] and Moutarlier and Chatila [20]. The method assumes that there are n features in the environment, and that they are static. The true state at time k is designated by $\mathbf{x}[k] = [\mathbf{x}_r[k]^T \ \mathbf{x}_f[k]^T]^T$, where $\mathbf{x}_r[k]$ represents the location of the robot, and $\mathbf{x}_f[k]^T = [\mathbf{x}_{f_1}[k]^T \ \dots \ \mathbf{x}_{f_n}[k]^T]^T$ represent the locations of the environmental features. Let $\mathbf{z}[k]$ designate the sensor measurements obtained at time k , and Z^k designate the set of all measurements obtained from time 0 through time k . The extended Kalman filter to compute recursively a state estimate $\hat{\mathbf{x}}[k|k] = [\hat{\mathbf{x}}_r[k|k]^T \ \hat{\mathbf{x}}_f[k|k]^T]^T$ at each discrete time step k , where $\hat{\mathbf{x}}_r[k|k]^T$ and $\hat{\mathbf{x}}_f[k|k]^T = [\hat{\mathbf{x}}_{f_1}[k|k]^T \ \dots \ \hat{\mathbf{x}}_{f_n}[k|k]^T]^T$ are the robot and feature state estimates, respectively. The stochastic mapping equations are not repeated here, for more detail, see [19, 21].

Data association decisions must be made for each new measurement to determine if (1) it originates from one of the features currently in the map, (2) it originates from a new feature, or (3) it is spurious. In general, the data association problem is exponentially complex [22], and no general solution that can run in real-time has been published. Most published implementations of CML have used variations of “nearest-neighbor” gating techniques [22], however nearest-neighbor gating can be shown to fail in many situations [23, 24]. Many implementations of stochastic mapping will only consider a measurement as a candidate for new feature initialization if it does gate with any existing features. However, such an approach will fail in environments with composite features (extended objects) or many features close to one another. The motivation for delayed stochastic mapping is to be able to consider various hypothesis for the origins of measurements in a computationally efficient manner, including the generation of a new feature with a measurement, even if it gates with an existing feature.

An assumption commonly employed in previous work is that the state of the new feature, $\hat{\mathbf{x}}_{f_{n+1}}[k]$ can be computed using the measurement data available from a single vehicle position, using a feature initialization function $\mathbf{g}(\cdot)$:

$$\hat{\mathbf{x}}_{f_{n+1}}[k] = \mathbf{g}(\hat{\mathbf{x}}[k|k], \mathbf{z}_j[k]). \quad (1)$$

For example, for a sensor providing range and bearing measurements, $\mathbf{z}_j[k] = [r \ \theta]$, the feature initialization function for a point $\mathbf{g}(\cdot)$ takes the following form:

$$\hat{\mathbf{x}}_{f_{n+1}}[k] = \mathbf{g}(\hat{\mathbf{x}}[k|k], \mathbf{z}_j[k]) = \begin{bmatrix} x_r + r \cos(\phi + \theta) \\ y_r + r \sin(\phi + \theta) \end{bmatrix}. \quad (2)$$

The new feature is integrated into the map by expanding the state vector $\hat{\mathbf{x}}[k|k]$ and covariance $\mathbf{P}[k|k]$ as shown below:

$$\hat{\mathbf{x}}[k|k] \leftarrow \begin{bmatrix} \hat{\mathbf{x}}[k|k] \\ \hat{\mathbf{x}}_{f_{n+1}}[k|k] \end{bmatrix}, \quad (3)$$

$$\mathbf{P}[k|k] \leftarrow \begin{bmatrix} \mathbf{P}_{rr}[k|k] & \mathbf{P}_{rf}[k|k] & \mathbf{P}_{rf_{n+1}}[k|k] \\ \mathbf{P}_{fr}[k|k] & \mathbf{P}_{ff}[k|k] & \mathbf{P}_{ff_{n+1}}[k|k] \\ \mathbf{P}_{f_{n+1}r}[k|k] & \mathbf{P}_{f_{n+1}f}[k|k] & \mathbf{P}_{f_{n+1}f_{n+1}}[k|k] \end{bmatrix}, \quad (4)$$

where

$$\mathbf{P}_{f_{n+1}f_{n+1}}[k|k] = \mathbf{G}_x \mathbf{P}[k|k] \mathbf{G}_x^T + \mathbf{G}_z \mathbf{R}[k] \mathbf{G}_z^T, \quad (5)$$

$$\begin{bmatrix} \mathbf{P}_{f_{n+1}r}[k|k] & \mathbf{P}_{f_{n+1}f}[k|k] \end{bmatrix} = \begin{bmatrix} \mathbf{P}_{f_{n+1}r}[k|k] \\ \mathbf{P}_{f_{n+1}f}[k|k] \end{bmatrix}^T = \mathbf{G}_x \mathbf{P}[k|k], \quad (6)$$

\mathbf{G}_x is the Jacobian of \mathbf{g} with respect to the state vector and \mathbf{G}_z is the Jacobian of \mathbf{g} with respect to the measurement.

To be able to perform feature initialization from multiple vantage points, we need to expand the state vector and to account for temporal correlations. To achieve this, the representation is expanded to add a number of previous vehicle locations to the state vector. We refer to these states as trajectory states. Each time the vehicle moves, the previous vehicle location is added to the state vector. We use the notation $\hat{\mathbf{x}}_{t_i}[k]$ to refer to the estimate of the state (position) of the robot at time i given all information up to time k . The complete trajectory of the robot for time step 0 through time step $k-1$ is given by the vector $\hat{\mathbf{x}}_t[k] = [\hat{\mathbf{x}}_{t_0}[k]^T \ \hat{\mathbf{x}}_{t_1}[k]^T \ \hat{\mathbf{x}}_{t_2}[k]^T \ \dots \ \hat{\mathbf{x}}_{t_{k-1}}[k]^T]^T$. The complete state vector is:

$$\hat{\mathbf{x}}[k|k] = \begin{bmatrix} \hat{\mathbf{x}}_r[k|k] \\ \hat{\mathbf{x}}_{t_0}[k] \\ \hat{\mathbf{x}}_{t_1}[k] \\ \hat{\mathbf{x}}_{t_2}[k] \\ \vdots \\ \hat{\mathbf{x}}_{t_{k-1}}[k] \\ \hat{\mathbf{x}}_{f_1}[k] \\ \hat{\mathbf{x}}_{f_2}[k] \\ \hat{\mathbf{x}}_{f_3}[k] \\ \vdots \\ \hat{\mathbf{x}}_{f_{n-1}}[k] \\ \hat{\mathbf{x}}_{f_n}[k] \end{bmatrix}. \quad (7)$$

The associated covariance matrix is:

$$\mathbf{P}[k|k] = \begin{bmatrix} \mathbf{P}_{rr}[k|k] & \mathbf{P}_{rt}[k|k] & \mathbf{P}_{rf}[k|k] \\ \mathbf{P}_{tr}[k|k] & \mathbf{P}_{tt}[k|k] & \mathbf{P}_{tf}[k|k] \\ \mathbf{P}_{fr}[k|k] & \mathbf{P}_{ft}[k|k] & \mathbf{P}_{ff}[k|k] \end{bmatrix}. \quad (8)$$

New trajectory states are added to the state vector each time step by defining a new trajectory state $\hat{\mathbf{x}}_{t_k}[k] = \hat{\mathbf{x}}_r[k|k]$ and adding this to the state vector:

$$\hat{\mathbf{x}}[k|k] \leftarrow \begin{bmatrix} \hat{\mathbf{x}}_r[k|k] \\ \hat{\mathbf{x}}_{t_0}[k] \\ \hat{\mathbf{x}}_{t_1}[k] \\ \hat{\mathbf{x}}_{t_2}[k] \\ \vdots \\ \hat{\mathbf{x}}_{t_{k-1}}[k] \\ \hat{\mathbf{x}}_{t_k}[k] \\ \hat{\mathbf{x}}_f[k] \end{bmatrix}. \quad (9)$$

The state covariance is expanded as follows:

$$\mathbf{P}[k|k] \leftarrow \begin{bmatrix} \mathbf{P}_{rr} & \mathbf{P}_{rt_0} & \cdots & \mathbf{P}_{rt_k} & \mathbf{P}_{rf} \\ \mathbf{P}_{t_0r} & \mathbf{P}_{t_0t_0} & \cdots & \mathbf{P}_{t_0t_k} & \mathbf{P}_{t_0f} \\ \vdots & \vdots & \ddots & \vdots & \vdots \\ \mathbf{P}_{t_{k-1}r} & \mathbf{P}_{t_{k-1}t_0} & \cdots & \mathbf{P}_{t_{k-1}t_k} & \mathbf{P}_{t_{k-1}f} \\ \mathbf{P}_{t_k r} & \mathbf{P}_{t_k t_0} & \cdots & \mathbf{P}_{t_k t_k} & \mathbf{P}_{t_k f} \\ \mathbf{P}_{fr} & \mathbf{P}_{ft_0} & \cdots & \mathbf{P}_{ft_k} & \mathbf{P}_{ff} \end{bmatrix}. \quad (10)$$

where $\mathbf{P}_{t_k t_i} = \mathbf{P}_{r t_i}$, $\mathbf{P}_{t_k f} = \mathbf{P}_{r f}$, and $\mathbf{P}_{t_k t_k} = \mathbf{P}_{r r}$. The growth of the state vector in this manner increases the computational burden, however it is straightforward to delete old vehicle trajectory states and associated terms in the covariance, once all the measurements from a given time step have been either processed or discarded.

This process of adding past states is similar to a fixed-lag Kalman smoother [25]. In a fixed-lag smoother, states exceeding a certain age are automatically removed. In our approach, states are added and removed based on the data processing requirements of the stochastic mapping process. Unlike the fixed-lag smoother, states are not necessarily removed in the order in which they are added.

With the addition of prior vehicle states to the state vector, it now becomes possible to initialize new features using measurements from multiple time steps. For example, consider the initialization of a new feature using two measurements, $\mathbf{z}[k_1]$ and $\mathbf{z}[k_2]$, taken at time steps k_1 and k_2 . The state of the new feature can be computed using a feature initialization function involving data from multiple time steps:

$$\hat{\mathbf{x}}_{f_{n+1}} = g(\hat{\mathbf{x}}_{t_{k_1}}[k], \hat{\mathbf{x}}_{t_{k_2}}[k], [\mathbf{z}[k_1]^T \ \mathbf{z}[k_2]^T]^T). \quad (11)$$

For example, in two-dimensions if each measurement is a range-only sonar measurement, then the function $g(\cdot)$ represents a solution for the intersection of two circles. The covariance for the new feature is initialized in a similar fashion as shown above in Equations 4 to 6, except that the Jacobian matrix \mathbf{G}_x will contain

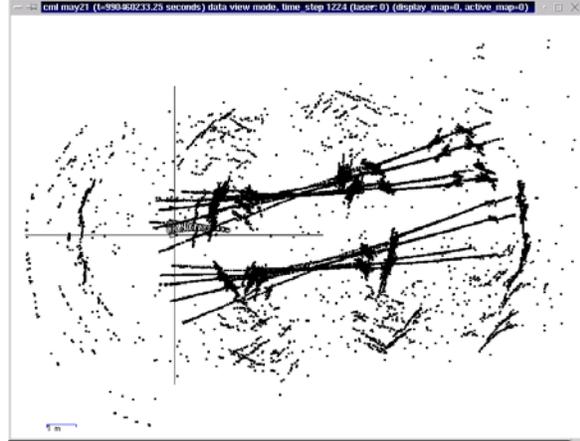


Figure 15: Raw data for corridor experiment, referenced to odometric position estimate.

additional non-zero terms corresponding to the trajectory states and the Jacobian matrix \mathbf{G}_z . The procedure is the same if the feature initialization function $g(\cdot)$ is a function of measurements from more than two time steps. New feature initialization can also be performed using non-linear least squares [26] performed on many measurements, instead of using an explicit function $g(\cdot)$.

To provide improved stability, the addition of new features to the state vector can be delayed to occur only when the initializing Jacobians indicate that the new feature estimate is well-conditioned. By examining the different possible initialization sets and choosing the Jacobian with the smallest values, the most stable initialization can be determined. In addition, one can incorporate an adaptive motion control step to direct the robot to move to a better vantage point that will yield a more stable initialization. By considering second-order derivatives, the robot can determine the optimal direction to move in order to obtain data that will yield the most stable initialization of a new feature.

Once a new feature is initialized, the map can be updated using all other previously obtained measurements that can be associated with the new feature. We call this procedure a "batch update". It allows the maximum amount of information to be extracted from all past measurements. It also provides a means to incrementally build up composite models of more complex objects [18].

Some illustrative results for delayed decision making are presented in Figures 15 to 18, which show the results for processing of data in an MIT corridor. Further details can be found in a forthcoming paper.

We believe that this methodology provides a new

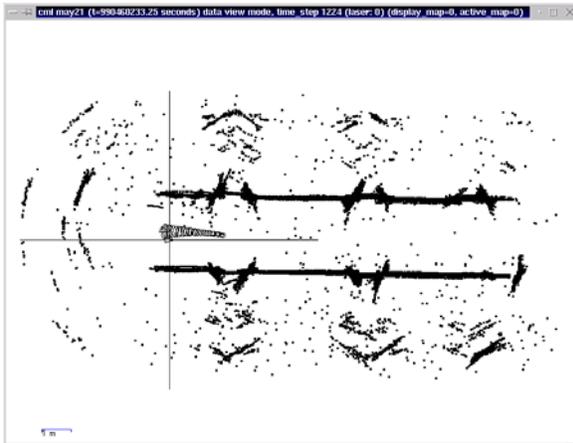


Figure 16: Raw data for corridor experiment, referenced to CML position estimate.

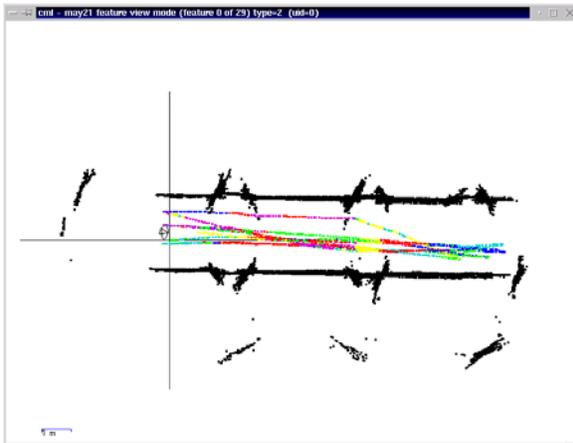


Figure 17: Sonar measurements that were matched to features for corridor experiment, displayed as points.

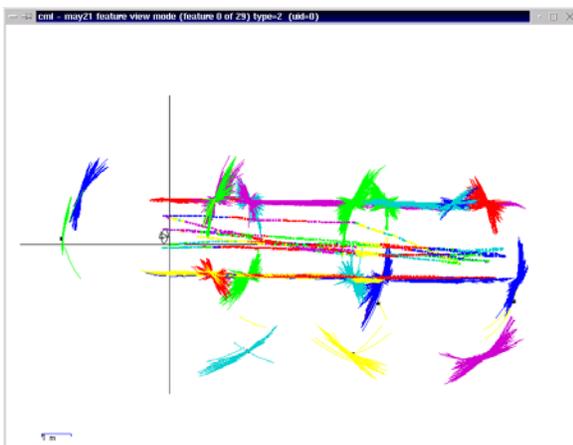


Figure 18: Sonar measurements that were matched to features for corridor experiment, displayed as circular arcs.

generic framework for improved feature modeling and classification using delayed decision making. The ability to perform a batch update using many previous measurements provides a facility for making delayed data association decisions. If there is ambiguity about the correspondence between measurements and features, decisions can be postponed until additional information becomes available. Feature extraction is also simplified. The initialization of complex features in situations with high ambiguity can be greatly simplified by considering a batch of data obtained at multiple time steps. Efficient, non-stochastic perceptual grouping methods such as the Hough technique described above in Section 3 can be used to screen the data and make preliminary association decisions that can later be confirmed with delayed stochastic gating, and then applied via batch updating.

5 Conclusion

This paper has considered the development of improved data association and feature modeling techniques for CML, using delayed decision making. Experimental results have been shown for both Polaroid sonar and SICK laser scanner data from a B21 robot, operating in the corridors of MIT, using several new data association and feature modeling techniques.

The ultimate goal of our research is to create a robust, consistent, convergent, and computationally efficient real-time algorithm for CML for large-scale environments. As a challenge for our research, we have set ourselves the show-term goal of autonomously performing CML in the largest building of MIT campus — the “infinite corridor” and its adjacent hallways — using standard Polaroid sonar data and/or laser data. Figure 21 shows a preliminary result that we have achieved with sonar only, showing a map consisting of about 800 features (walls, door moldings, etc.) built from a single pass around a large loop (travel distance of about 200 meters, duration of 26 minutes). Features were initialized using the output of the Hough perceptual grouping technique described above in Section 3. The processing time for this map is slightly greater than real time on an 800 MHz machine. Much further work is necessary to reach our goal of achieving convergence, while operating in real-time and maintaining consistent error bounds. We feel that the map scaling problem can be addressed with spatial and temporal partitioning techniques, and that the biggest remaining roadblocks are the problems of data association and modeling of complex features. We believe that the techniques presented in this paper offer significant help towards reaching this goal.

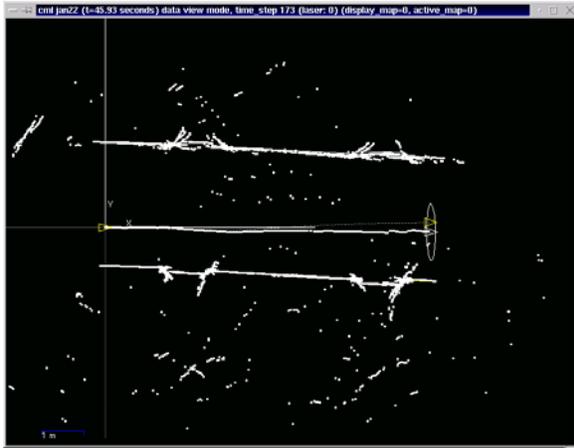


Figure 19: Raw sonar data for the beginning of the experiment.

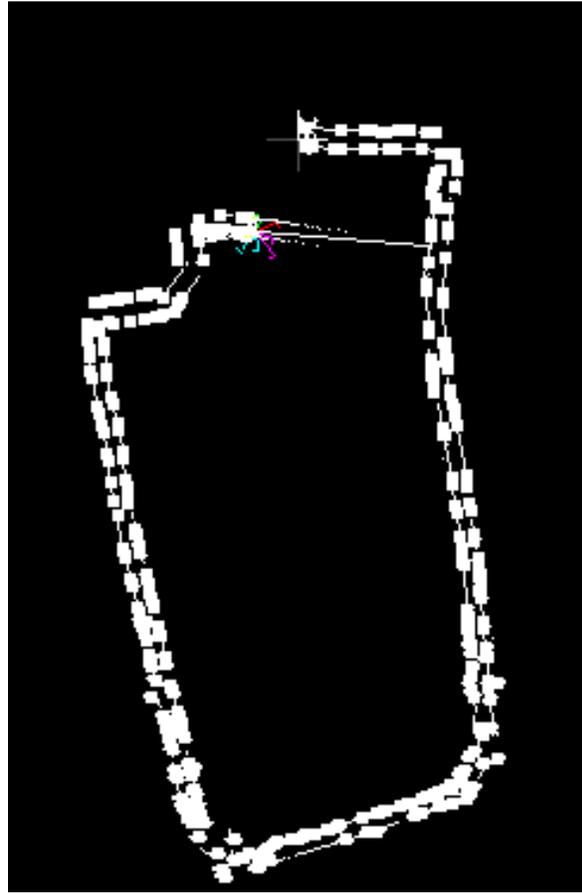


Figure 21: Map constructed for a complete loop. The robot is approximately back at its starting position at the end of the trajectory.

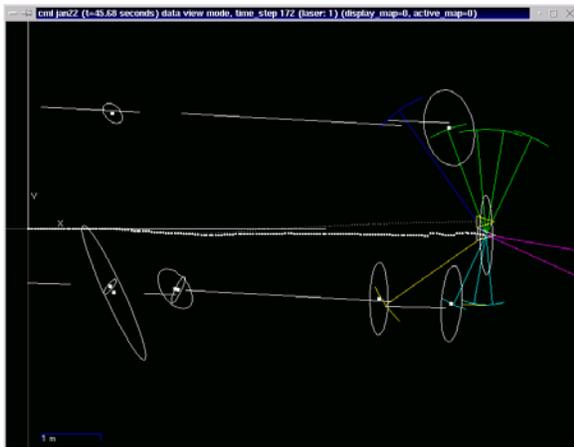


Figure 20: Point and line features (with 3σ error ellipses for point targets) for the beginning of the experiment.



Figure 22: Schematic of the MIT campus, showing the “infinite corridor” that connects buildings 7, 3, 10, 4, and 8. The data for Figure 21 was taken in buildings 5, 7, 3, and 1.

Acknowledgments

This research has been funded in part by the Henry L. and Grace Doherty Assistant Professorship in Ocean Utilization, NSF Career Award BES-9733040, the MIT Sea Grant College Program under grant NA86RG0074 (project RCM-3), and Ministerio de Educación y Cultura of Spain, grant PR2000-0104, and by the Spain-US Commission for Educational and Scientific Exchange (Fulbright), grant 20079.

References

- [1] R. Chatila. Mobile robot navigation space modeling and decisional processes. In *Third Int. Symposium Robotics Research*. MIT Press, 1985.
- [2] R. A Brooks. Aspects of mobile robot visual map making. In *Second Int. Symp. Robotics Research*, pages 287–293, Tokyo, Japan, 1984. MIT Press.
- [3] R. Smith and P. Cheeseman. On the representation and estimation of spatial uncertainty. *International Journal of Robotics Research*, 5(4):56, 1987.
- [4] P. Moutarlier and R. Chatila. Stochastic multi-sensory data fusion for mobile robot location and environment modeling. In *5th Int. Symposium on Robotics Research*, pages 207–216, 1989.
- [5] S. Thrun. An online mapping algorithm for teams of mobile robots. *Int. J. Robotics Research*, 20(5):335–363, May 2001.
- [6] J-S. Gutmann and K. Konolige. Incremental mapping of large cyclic environments. In *International Symposium on Computational Intelligence in Robotics and Automation*, 1999.
- [7] J. A. Castellanos and J. D. Tardos. *Mobile Robot Localization and Map Building: A Multisensor Fusion Approach*. Kluwer Academic Publishers, Boston, 2000.
- [8] M. W. M. G. Dissanayake, P. Newman, H. F. Durrant-Whyte, S. Clark, and M. Csorba. An experimental and theoretical investigation into simultaneous localization and map building. In *Sixth International Symposium on Experimental Robotics*, pages 265–274, March 1999.
- [9] A. J. Davison. *Mobile Robot Navigation Using Active Vision*. PhD thesis, University of Oxford, 1998.
- [10] J. J. Leonard and H. J. S. Feder. A computationally efficient method for large-scale concurrent mapping and localization. In D Koditschek and J. Hollerbach, editors, *Robotics Research: The Ninth International Symposium*, pages 169–176, Snowbird, Utah, 2000. Springer Verlag.
- [11] J. Guivant and E. Nebot. Optimization of the simultaneous localization and map building algorithm for real time implementation. *IEEE Transactions on Robotic and Automation*, 17(3):242–257, June 2001.
- [12] J. Neira and J.D. Tardós. Data association in stochastic mapping using the joint compatibility test. *IEEE Trans. on Robotics and Automation*, page (To appear), 2001.
- [13] J. A. Castellanos, J. M. M. Montiel, J. Neira, and J. D. Tardos. The SPmap: A probabilistic framework for simultaneous localization and map building. *IEEE Trans. Robotics and Automation*, 15(5):948–952, 1999.
- [14] M. W. M. G. Dissanayake, P. Newman, H. F. Durrant-Whyte, S. Clark, and M. Csorba. A solution to the simultaneous localization and map building (slam) problem. *IEEE Transactions on Robotic and Automation*, 17(3):229–241, June 2001.
- [15] J. J. Leonard and H. F. Durrant-Whyte. *Directed Sonar Sensing for Mobile Robot Navigation*. Boston: Kluwer Academic Publishers, 1992.
- [16] D. Ballard and C. Brown. *Computer Vision*. Prentice-Hall, 1982.
- [17] O. Wijk and H. Christensen. Triangulation based fusion of sonar data with application in robot pose tracking. *IEEE Trans. Robotics and Automation*, 16(6):740–752, December 2000.
- [18] J. J. Leonard and R. Rikoski. Incorporation of delayed decision making into stochastic mapping. In D. Rus and S. Singh, editors, *Experimental Robotics VII*, Lecture Notes in Control and Information Sciences. Springer-Verlag, 2001.
- [19] R. Smith, M. Self, and P. Cheeseman. Estimating uncertain spatial relationships in robotics. In I. Cox and G. Wilfong, editors, *Autonomous Robot Vehicles*, pages 167–193. Springer-Verlag, 1990.
- [20] 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.
- [21] H. J. S. Feder, J. J. Leonard, and C. M. Smith. Adaptive mobile robot navigation and mapping. *Int. J. Robotics Research*, 18(7):650–668, July 1999.
- [22] Y. Bar-Shalom and T. E. Fortmann. *Tracking and Data Association*. Academic Press, 1988.
- [23] I. J. Cox and J. J. Leonard. Modeling a dynamic environment using a Bayesian multiple hypothesis approach. *Artificial Intelligence*, 66(2):311–344, April 1994.
- [24] J. Neira and J. D. Tardos. Data association in stochastic mapping: the fallacy of the nearest neighbor, 2000. Workshop on mobile robot navigation and mapping (part of ICRA 2000).
- [25] B. D. O. Anderson and J. B. Moore. *Optimal filtering*. Englewood Cliffs, N.J.: Prentice-Hall, 1979.
- [26] O. Faugeras. *Three-Dimensional Computer Vision: A Geometric Viewpoint*. MIT Press, 1993.