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

John J. Leonard[1], Paul M. Newman[1], Richard J. Rikoski[1], José Neira[2], and Juan D. Tardós[2]

[1] Massachusetts Institute of Technology, Cambridge MA 02139, USA
[2] 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. 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 a 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 solution 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 (Brooks, 1984; Chatila, 1985; Moutarlier, Chatila, 1989; Smith, Cheeseman, 1987). For example, in his paper for the second ISRR symposium, Brooks (Brooks, 1984) 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 (Brooks, 1984)."

R.A. Jarvis and A. Zelinsky (Eds.): Robotics Research, STAR 6, pp. 7–20, 2003

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".

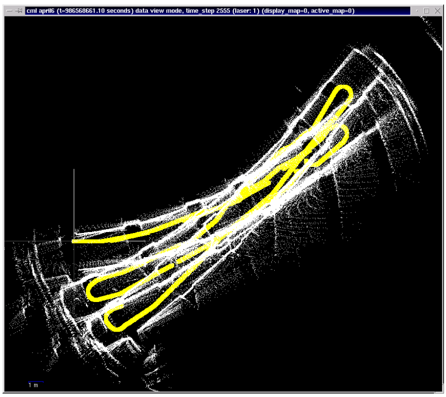

Fig. 1. Laser data for a short corridor experiment, referenced to the dead-reckoning position estimate

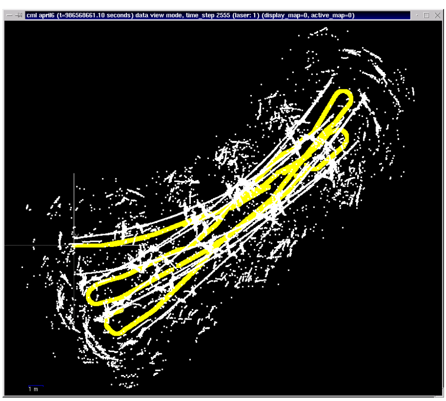

Fig. 2. Sonar data for a short corridor experiment, referenced to the dead-reckoning position estimate

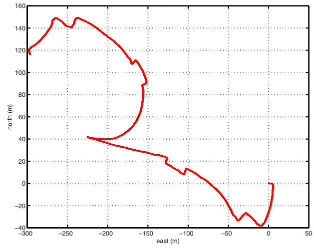

Fig. 3. Accumulation of position error relying only on dead-reckoning for a long distance traverse of the B21 mobile robot. The actual vehicle path went down approximately 40 meters, to the left approximately 225 meters, and then back to the origin.

Most successful recent implementations of CML have either been performed with SICK laser scanner data (Gutmann, Konolige, 1999; Thrun, 2001) or in environments that consist of isolated "point" objects (Castellanos, Tardos, 2000; Dissanayake, Newman, Durrant-Whyte, Clark, Csorba, 2001). 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. 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.

Gutmann, Konolige (1999); Thrun (2001) 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 (Castellanos, Tardos, 2000; Dissanayake et al., 2001). The representation of spatial correlations results in an $\mathcal{O}(n^2)$ growth in computational cost (Moutarlier, Chatila, 1989), motivating techniques to address the map scaling problem through spatial and temporal partitioning (Davison, 1998; Guivant, Nebot, 2001; Leonard, Feder, 2000). 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 of multiple sensor measurements, using a branch and bound algorithm, has been developed (Neira, Tardós, 2001).

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

2 "Explore and return" using Laser

This section presents results from use 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 SPMap framework (Castellanos, Tardos, 2000), (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 system was rendered in 3D and used as a real-time visualization tool of the robots workspace. This enabled the remote operator to "visit" previously un-explored 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

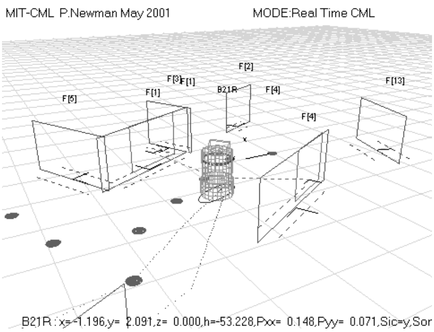

Fig. 4. Re-observing an existing feature

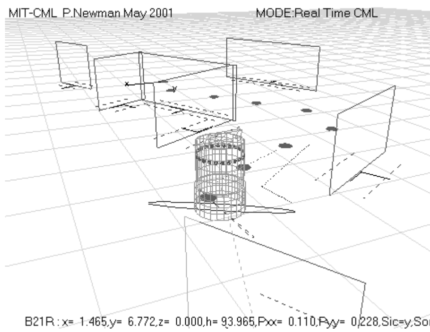

Fig. 5. Creating a new feature in the foreground following a rotation

wood-clad pillars had been erected throughout the hallway yielding an interesting, landmark rich and densely populated area.

Fig. 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 Dissanayake et al. (2001) 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.

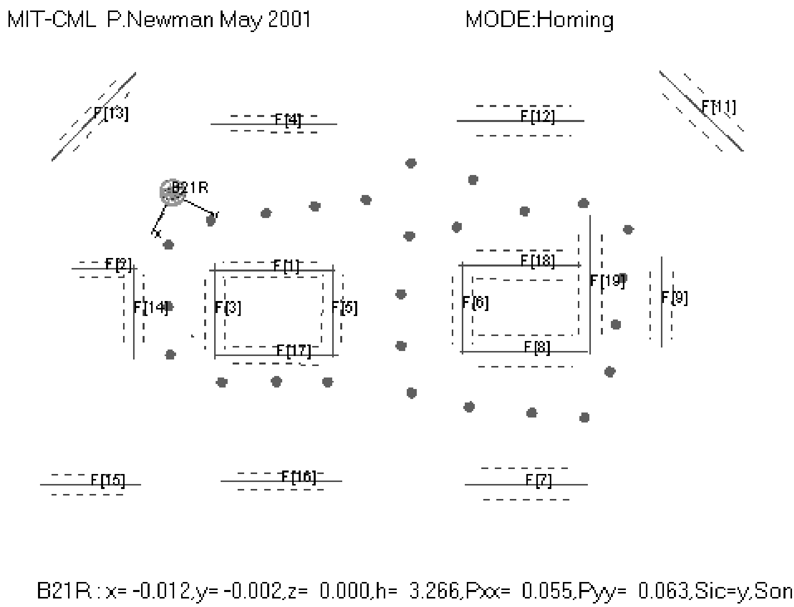

Fig. 7. A plan view of the CML map at the end of the experiment. The approximate size of the environment was a 20m by 15m rectangle.

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 CML algorithm, 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 program 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`.

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

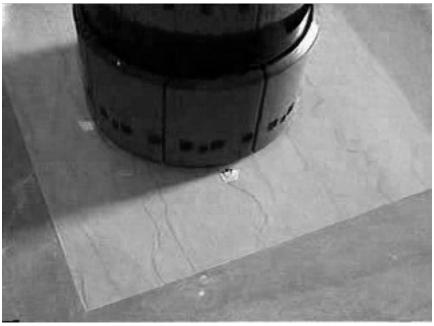

Fig. 8. The starting position

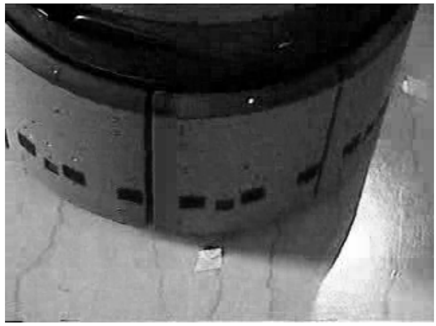

Fig. 9. The robot position after the completion of the homing leg of the mission

mapping. However, using a physics-based sensor model, the geometric constraints provided by an individual sonar return can be formulated (Leonard, Durrant-Whyte, 1992). 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 (Ballard, Brown, 1982) can be used as a voting scheme identify point and planar features. More detail on this technique can be found in Tardós, Neira, Newman, Leonard (2002). A related technique called triangulation-based fusion has been developed in Wijk, Christensen (2000) for point objects only. Figure 10 through 12 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 SPMap as the state estimation framework (Castellanos, Tardos, 2000), Joint Compatibility for data association (Neira, Tardós, 2001), and a new technique called Sequential Map Joining (Tardós et al., 2002). Figure 13 shows a map of the same environment built from laser data. One can see that sonar map is almost as good as the laser map.

4 Delayed Stochastic Mapping

Stochastic mapping is a feature-based concurrent mapping and localization algorithm that was first published in Moutarlier, Chatila (1989); Smith, Self, Cheeseman (1990). 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 $\boldsymbol{x}[k] = [\boldsymbol{x}_r[k]^T \ \boldsymbol{x}_f[k]^T]^T$, where $\boldsymbol{x}_r[k]$ represents the location of the robot, and $\boldsymbol{x}_f[k]^T = [\boldsymbol{x}_{f_1}[k]^T \ \ldots \ \boldsymbol{x}_{f_n}[k]^T]^T$ represent the locations of the environmental features. Let $\boldsymbol{z}[k]$ 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 $\hat{\boldsymbol{x}}[k|k] = [\hat{\boldsymbol{x}}_r[k|k]^T \ \hat{\boldsymbol{x}}_f[k]^T]^T$ at each

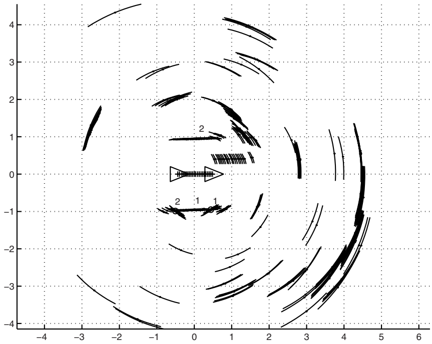

Fig. 10. Example of Hough processing to extract point and line features. Sonar returns are processed in a group of twenty-two positions. A voting scheme is performed to find clusters of measurements that hypothesize the existence of point and plane features. For this example, two planes and two points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

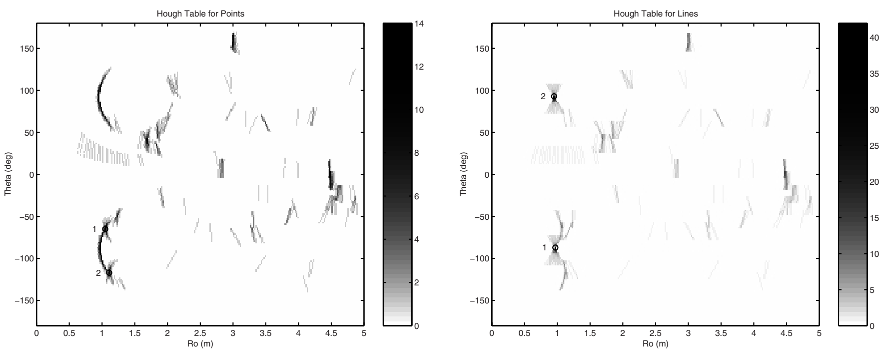

Fig. 11. Hough voting table for point (left) and line (right) features.

discrete time step k, where $\hat{\boldsymbol{x}}_r[k|k]^T$ and $\hat{\boldsymbol{x}}_f[k]^T = [\hat{\boldsymbol{x}}_{f_1}[k]^T \ldots \hat{\boldsymbol{x}}_{f_n}[k]^T]^T$ are the robot and feature state estimates, respectively. The stochastic mapping equations are not repeated here, for more detail, see Feder, Leonard, Smith (1999); Smith et al. (1990).

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 (Bar-Shalom, Fortmann, 1988), and no general solution that can run in real-time has been published. The motivation for delayed stochastic mapping is to be able to consider various hypothesis for the origins of measurements in a computationally efficient manner.

An assumption commonly employed in previous work is that the state of the new feature, $\hat{\boldsymbol{x}}_{f_{n+1}}[k]$ can be computed using the measurement data available from

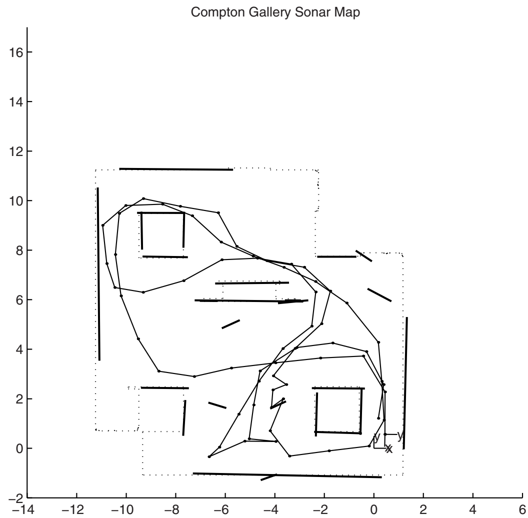

Fig. 12. Complete map for the MIT Compton Gallery built from sonar using Hough grouping, Map Joining, and Joint Compatibility.

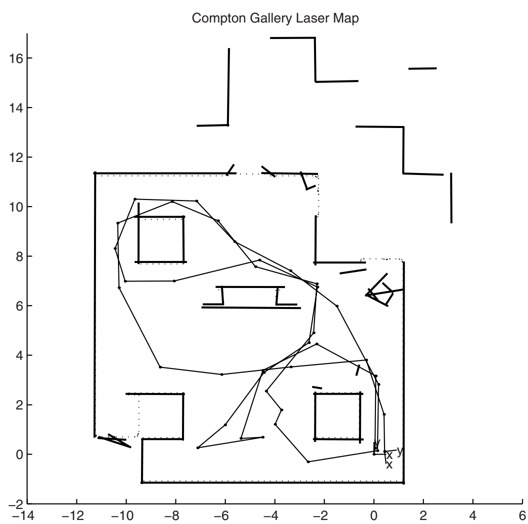

Fig. 13. Complete map for the MIT Compton Gallery built from laser data using Robust Statistics, Map Joining, and Joint Compatibility.

a single vehicle position, using a feature initialization function $\boldsymbol{g}(\cdot)$:

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

To be able to perform feature initialization from multiple vantage points, 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 $\hat{\boldsymbol{x}}_{t_i}[k]$ 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 $\hat{\boldsymbol{x}}_t[k] = [\hat{\boldsymbol{x}}_{t_0}[k]^T \ \hat{\boldsymbol{x}}_{t_1}[k]^T \ \hat{\boldsymbol{x}}_{t_2}[k]^T \ \ldots \ \hat{\boldsymbol{x}}_{t_{k-1}}[k]]^T$. The complete state vector is:

$$\hat{\boldsymbol{x}}[k|k] = \begin{bmatrix} \hat{\boldsymbol{x}}_r[k|k] \\ \hat{\boldsymbol{x}}_t[k] \\ \hat{\boldsymbol{x}}_f[k] \end{bmatrix} = \begin{bmatrix} \hat{\boldsymbol{x}}_r[k|k] \\ \hat{\boldsymbol{x}}_{t_0}[k] \\ \hat{\boldsymbol{x}}_{t_1}[k] \\ \hat{\boldsymbol{x}}_{t_2}[k] \\ \vdots \\ \hat{\boldsymbol{x}}_{t_{k-1}}[k] \\ \hat{\boldsymbol{x}}_{f_1}[k] \\ \hat{\boldsymbol{x}}_{f_2}[k] \\ \hat{\boldsymbol{x}}_{f_3}[k] \\ \vdots \\ \hat{\boldsymbol{x}}_{f_{n-1}}[k] \\ \hat{\boldsymbol{x}}_{f_n}[k] \end{bmatrix}. \tag{2}$$

The associated covariance matrix is:

$$\boldsymbol{P}[k|k] = \begin{bmatrix} \boldsymbol{P}_{rr}[k|k] & \boldsymbol{P}_{rt}[k|k] & \boldsymbol{P}_{rf}[k|k] \\ \boldsymbol{P}_{tr}[k|k] & \boldsymbol{P}_{tt}[k|k] & \boldsymbol{P}_{tf}[k|k] \\ \boldsymbol{P}_{fr}[k|k] & \boldsymbol{P}_{ft}[k|k] & \boldsymbol{P}_{ff}[k|k] \end{bmatrix}. \tag{3}$$

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

$$\hat{\boldsymbol{x}}[k|k] \leftarrow \begin{bmatrix} \hat{\boldsymbol{x}}_r[k|k] \\ \hat{\boldsymbol{x}}_{t_0}[k] \\ \vdots \\ \hat{\boldsymbol{x}}_{t_{k-1}}[k] \\ \hat{\boldsymbol{x}}_{t_k}[k] \\ \hat{\boldsymbol{x}}_f[k] \end{bmatrix}. \tag{4}$$

The state covariance is expanded as follows:

$$\boldsymbol{P}[k|k] \leftarrow \begin{bmatrix} \boldsymbol{P}_{rr} & \boldsymbol{P}_{rt_0} & \dots & \boldsymbol{P}_{rt_k} & \boldsymbol{P}_{rf} \\ \boldsymbol{P}_{t_0r} & \boldsymbol{P}_{t_0t_0} & \dots & \boldsymbol{P}_{t_0t_k} & \boldsymbol{P}_{t_0f} \\ \vdots & \vdots & \ddots & \vdots & \vdots & \vdots \\ \boldsymbol{P}_{t_{k-1}r} & \boldsymbol{P}_{t_{k-1}t_0} & \dots & \boldsymbol{P}_{t_{k-1}t_k} & \boldsymbol{P}_{t_{k-1}f} \\ \boldsymbol{P}_{t_kr} & \boldsymbol{P}_{t_kt_0} & \dots & \boldsymbol{P}_{t_kt_k} & \boldsymbol{P}_{t_kf} \\ \boldsymbol{P}_{fr} & \boldsymbol{P}_{ft_0} & \dots & \boldsymbol{P}_{ft_k} & \boldsymbol{P}_{ff} \end{bmatrix}, \tag{5}$$

where $\boldsymbol{P}_{t_kt_i} = \boldsymbol{P}_{rt_i}$, $\boldsymbol{P}_{t_kf} = \boldsymbol{P}_{rf}$, and $\boldsymbol{P}_{t_kt_k} = \boldsymbol{P}_{rr}$. 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 (Anderson, Moore, 1979). 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, $\boldsymbol{z}[k_1]$ and $\boldsymbol{z}[k_2]$, 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:

$$\hat{\boldsymbol{x}}_{f_{n+1}} = g(\hat{\boldsymbol{x}}_{t_{k_1}}[k], \hat{\boldsymbol{x}}_{t_{k_2}}[k], [\boldsymbol{z}[k_1]^T \; \boldsymbol{z}[k_2]^T]^T). \tag{6}$$

For example, in two-dimensions if each measurement is a range-only sonar measurement, then the function $\boldsymbol{g}(\cdot)$ represents a solution for the intersection of two circles. The procedure is the same if the feature initialization function $\boldsymbol{g}(\cdot)$ is a function of measurements from more than two time steps.

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 (Leonard, Rikoski, 2001). The method has been implemented as part of an integrated framework for real-time CML, which incorporates delayed state management, perceptual grouping, multiple vantage point initialization, batch updating, and feature fusion. Some illustrative results for this approach are presented in Figures 14 to 17, which show the results for processing of data in an MIT corridor. Further details can be found in Leonard, Rikoski, Newman, Bosse (2002). These experiments used the Hough Transform for sonar perceptual grouping as described above in Section 3.

This methodology provides a new generic framework for improved feature modeling and classification. The ability to perform a batch update using many previous

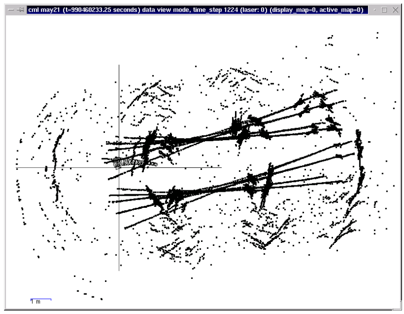

Fig. 14. Raw sonar data for corridor experiment, referenced to odometry.

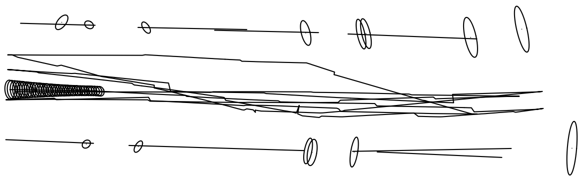

Fig. 15. CML estimated trajectory for corridor scene and estimated map consisting of points and line segments. Three-sigma error bounds are shown for the location of points.

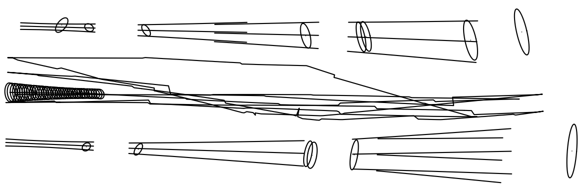

Fig. 16. Same plot as in Fig. 15 but with three-sigma error bounds for lines added.

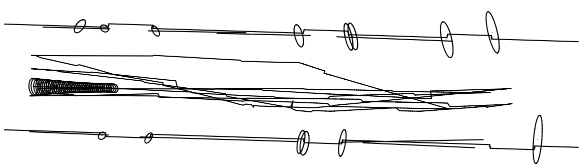

Fig. 17. Same plot as in Fig. 15 but with hand-measured model overlaid.

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. 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.

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, Ministerio de Educación y Cultura of Spain, grant PR2000-0104, and the Spain-US Commission for Educational and Scientific Exchange (Fulbright), grant 20079.

References

Anderson BDO, Moore JB (1979) Optimal filtering. Prentice-Hall, Englewood Cliffs, N.J.

Ballard D, Brown C (1982) Computer vision. Prentice-Hall.

Bar-Shalom Y, Fortmann TE (1988) Tracking and data association. Academic Press.

Brooks RA (1984) Aspects of mobile robot visual map making. In Second Int. Symp. Robotics Research, pp 287–293, Tokyo, Japan. MIT Press.

Castellanos JA, Tardos JD (2000) Mobile robot localization and map building: A multisensor fusion approach. Kluwer Academic Publishers, Boston.

Chatila R (1985) Mobile robot navigation space modeling and decisional processes. In Third Int. Symposium Robotics Research. MIT Press.

Davison AJ (1998) Mobile robot navigation using active vision. PhD thesis, University of Oxford.

Dissanayake MWMG, Newman P, Durrant-Whyte HF, Clark S, Csorba M (2001) A solution to the simultaneous localization and map building (slam) problem. IEEE Transactions on Robotic and Automation, 17(3):229–241.

Feder HJS, Leonard JJ, Smith CM (1999) Adaptive mobile robot navigation and mapping. Int. J. Robotics Research, 18(7):650–668.

Guivant J, Nebot E (2001) Optimization of the simultaneous localization and map building algorithm for real time implementation. IEEE Transactions on Robotic and Automation, 17(3):242–257.

Gutmann JS, Konolige K (1999) Incremental mapping of large cyclic environments. In International Symposium on Computational Intelligence in Robotics and Automation.

Leonard JJ, Durrant-Whyte HF (1992) Directed sonar sensing for mobile robot navigation. Kluwer Academic Publishers, Boston.

Leonard JJ, Feder HJS (2000) A computationally efficient method for large-scale concurrent mapping and localization. In Koditschek D, Hollerbach J (eds) Robotics Research: The Ninth International Symposium, pp 169–176, Snowbird, Utah. Springer Verlag.

Leonard JJ, Rikoski R (2001) Incorporation of delayed decision making into stochastic mapping. In Rus D, Singh S (eds) Experimental Robotics VII, Lecture Notes in Control and Information Sciences. Springer-Verlag.

Leonard JJ, Rikoski RJ, Newman PM, Bosse M (2002) Mapping partially observable features from multiple uncertain vantage points. Int. J. Robotics Research. To Appear.

Moutarlier P, Chatila R (1989) Stochastic multi-sensory data fusion for mobile robot location and environment modeling. In 5th Int. Symposium on Robotics Research, pp 207–216.

Neira J, Tardós J (2001) Data association in stochastic mapping using the joint compatibility test. IEEE Trans. on Robotics and Automation, 17(6):890–897.

Smith R, Cheeseman P (1987) On the representation and estimation of spatial uncertainty. International Journal of Robotics Research, 5(4):56.

Smith R, Self M, Cheeseman P (1990) Estimating uncertain spatial relationships in robotics. In Cox I, Wilfong G (eds) Autonomous Robot Vehicles, pp 167–193. Springer-Verlag.

Tardós J, Neira J, Newman PM, Leonard JJ (2002) Robust mapping and localization in indoor environments using sonar data. Int. J. Robotics Research. To Appear.

Thrun S (2001) An online mapping algorithm for teams of mobile robots. Int. J. Robotics Research, 20(5):335–363.

Wijk O, Christensen H (2000) Triangulation based fusion of sonar data with application in robot pose tracking. IEEE Trans. Robotics and Automation, 16(6):740–752.