Concurrent Mapping and Localization for Autonomous Underwater Vehicles

Christopher M. Smith[1]   John J. Leonard[1]   Andrew A. Bennett[1]   Christopher Shaw[2]

[1] Department of Ocean Engineering
Massachusetts Institute of Technology
Cambridge, Massachusetts 02139, U.S.A.

[2] Naval Undersea Warfare Center
Newport, Rhode Island 02841, U.S.A.

ABSTRACT

Although traditional tracking and multitarget tracking have examined the problems of estimating multiple targets in an uncertain environment, the addition of navigational uncertainty results in additional complexities which are not addressed by these methodologies. Concurrent mapping and localization refers to the use of environmental cues to both build a map and reduce navigational uncertainty. This new technology promises to enhance greatly the performance of unmanned underwater vehicles (UUVs) navigating in unknown and unmapped environments. After examining background to this problem, we describe the process of concurrent mapping and localization. An implementation plan for realizing the capability of concurrent mapping and localization on a field vehicle is presented. Key research issues are highlighted and progress in these areas is presented.

I. INTRODUCTION

Traditional navigation techniques are inadequate for many envisioned missions for UUVs. Dead-reckoning and inertial navigation systems are insufficient for medium to long duration (extent) missions, because navigation error grows without bound. GPS navigation is usually impractical because the act of surfacing to get a position fix can conflict with mission objectives. Long baseline (LBL) acoustic navigation requires that an array of beacons has been deployed and calibrated, which is often not possible.

To overcome these limitations, techniques for non-tradition navigation (NTN) have been proposed. The basic idea of NTN is to match measurements of one or more geophysical properties, such as bathymetry, gravity, or magnetic field, to an *a priori* environment map. If there is sufficient spatial variation in the parameter(s) being measured, there is potential to reduce navigation uncertainty without surfacing or the use of acoustic beacons [14, 15, 5, 2]. However, techniques proposed to date have relied upon the availability of an accurate *a priori* map of one or more geophysical parameters in the operating area. The goal of our research is to enable navigation for UUVs in areas where an *a priori* map of the environment is not available.

This capability would enable a UUV to perform repeatable, geographical tasks in an unknown environment. For example, consider a mission that requires a UUV to conduct a survey of an area, looking for distinctive objects, and, once the survey is complete, to revisit selected features for more detailed imaging. The question we wish to pose for such a mission is: could the UUV integrate the observations that it obtains during the course of its mission to enable accurate revisitation of feature locations in scenarios where INS and dead-reckoning are by themselves insufficient?

A set of feature location estimates can basically be thought of as a map. The challenge is to combine INS/dead-reckoning information with sensor observations of features to build such a map of feature location estimates. One can envision construction of either a single, globally-referenced map or a network of local maps.

The problem of simultaneously building a map and navigating relative to that map is a major outstanding problem in the robotics research community. A seminal technique for concurrent mapping and localization, called the stochastic map, was published by Smith, Self, and Cheeseman [12]. The stochastic map consists of a single state vector that represents the estimates of the vehicle and feature locations and an associated covariance matrix. As the vehicle moves around its

Report Documentation Page

Public reporting burden for the collection of information is estimated to average 1 hour per response, including the time for reviewing instructions, searching existing data sources, gathering and maintaining the data needed, and completing and reviewing the collection of information. Send comments regarding this burden estimate or any other aspect of this collection of information, including suggestions for reducing this burden, to Washington Headquarters Services, Directorate for Information Operations and Reports, 1215 Jefferson Davis Highway, Suite 1204, Arlington VA 22202-4302. Respondents should be aware that notwithstanding any other provision of law, no person shall be subject to a penalty for failing to comply with a collection of information if it does not display a currently valid OMB control number.

1. REPORT DATE **1997**	2. REPORT TYPE	3. DATES COVERED **00-00-1997 to 00-00-1997**

4. TITLE AND SUBTITLE **Concurrent Mapping and Localization for Autonomous Underwater Vehicles**	5a. CONTRACT NUMBER
	5b. GRANT NUMBER
	5c. PROGRAM ELEMENT NUMBER
6. AUTHOR(S)	5d. PROJECT NUMBER
	5e. TASK NUMBER
	5f. WORK UNIT NUMBER

7. PERFORMING ORGANIZATION NAME(S) AND ADDRESS(ES) **Massachusetts Institute of Technology,Department of Ocean Engineering,Cambridge,MA,02139**	8. PERFORMING ORGANIZATION REPORT NUMBER

9. SPONSORING/MONITORING AGENCY NAME(S) AND ADDRESS(ES)	10. SPONSOR/MONITOR'S ACRONYM(S)
	11. SPONSOR/MONITOR'S REPORT NUMBER(S)

12. DISTRIBUTION/AVAILABILITY STATEMENT
Approved for public release; distribution unlimited

13. SUPPLEMENTARY NOTES
In Undersea Defence Technology, pages 338-342. 1997. U.S. Government or Federal Rights License

14. ABSTRACT

Although traditional tracking and multitarget tracking have examined the problems of estimating multiple targets in an uncertain environment, the addition of navigational uncertainty results in additional complexities which are not addressed by these methodologies. Concurrent mapping and localization refers to the use of environmental cues to both build a map and reduce navigational uncertainty. This new technology promises to enhance greatly the performance of unmanned underwater vehicles (UUVs) navigating in unknown and unmapped environments. After examining background to this problem, we describe the process of concurrent mapping and localization. An implementation plan for realizing the capability of concurrent mapping and localization on a eld vehicle is presented. Key research issues are highlighted and progress in these areas is presented.

15. SUBJECT TERMS

16. SECURITY CLASSIFICATION OF:			17. LIMITATION OF ABSTRACT	18. NUMBER OF PAGES	19a. NAME OF RESPONSIBLE PERSON
a. REPORT **unclassified**	b. ABSTRACT **unclassified**	c. THIS PAGE **unclassified**	**Same as Report (SAR)**	**7**	

environment, taking measurements of environmental features, the stochastic map is updated using an extended Kalman filter. Moutarlier and Chatila [9] and Rencken [11] have implemented suboptimal versions of the stochastic map using land robots.

The major problems encountered by the stochastic map are the failure of the extended Kalman filter to properly track the highly nonlinear transformations involved in geometric estimation and the fact that the technique scales (at best) quadratically with the number of features present. Also, Smith, Self, and Cheeseman ignored errors which may arise from ambiguity in the source of sensor data. Some of these drawbacks have recently been considered by Uhlmann [16] in a theoretical investigation of concurrent mapping and localization. However, there are many important issues to address in future research to realize this capability on-board a UUV.

II. PROBLEM STATEMENT

A. WHAT IS CONCURRENT MAPPING AND LOCALIZATION?

It is easiest to understand the problems of performing concurrent mapping and localization by considering the progression of problems from the classical tracking problem. In the traditional tracking problem, we need to estimate the positions of a set of known features using measurements. The sensing location is known. We have stochastic models of target behavior and the measurement process. When we take a measurement, we know what we are measuring; specifically, we know which feature we are measuring.

Multi-target tracking takes this problem a step further, assuming that the features are unknown. We are not given the number or type of features present, so that we are estimating what features are present in addition to the positions of these postulated features. There is also data association uncertainty. In other words, when we take a measurement, we have no *a priori* information about what feature, if any, we are measuring.

Concurrent mapping and localization is further complicated by allowing the vehicle position to be uncertain. At first blush, this does not seem likely to add fundamental difficulties to multi-target tracking. The difficulties become apparent however, when we consider the solutions to these problems. The traditional tracking problem is the realm of the Kalman filter and its variants (such as the extended Kalman filter). The essential point is that a recursive algorithm is used to efficiently (in a computational sense) obtain a good estimate of the feature positions given initial state information and a history of measurements.

Our approach to concurrent mapping and localization derives from a tracking technique called multiple hypothesis tracking [10, 1]. In order to ensure theoretically consistent estimates, we look for explanations for measurements and prospective tracks. A measurement must come from a single prospective feature or a new feature, or be a false alarm. Each prospective feature must give rise to a measurement or undergo a skipped measurement (a miss). It is often the case that current information does not provide enough explanatory power to confidently make a decision about how to assign measurements to features. Therefore, decisions are delayed until enough corroborative evidence is developed to make assignment unambiguous (or moot). However, by delaying decisions in this way, we have not one estimate of the feature positions, but a set of competing hypotheses about the current state of the world.

When vehicle navigational uncertainty is added to multiple hypothesis tracking problem, the result is that there are multiple competing estimates of the current vehicle position, as well as the feature positions.

B. THE PROCESS OF CM&L

The process of performing concurrent mapping and localization as described above can be separated into several steps: state projection, measurement gating, hypothesis formation, state updating, and assignment decision-making. Each of these steps is described below. Figure 1 presents the process of concurrent mapping and localization in flowchart form.

B.1. STATE PROJECTION

State projection is the prediction of a vehicle or feature state in the future given an assumed uncertain current state. State projection using a plant model for the feature or vehicle is well-understood. A possible complication is that there may be multiple models for the vehicle or a feature. For example, we may have a model for steady motion and one for maneuvering. The result is that each current state produces a set of possible projected states. The state projection takes care of estimating not only the position of the vehicle or feature, but also the growth of uncertainty during the period when no measurements are taken.

B.2. GATING MEASUREMENTS

Once we have sets of possible projected states for the vehicle and features, we can start deciding how to match up features and measurements. In order to reduce the number of possible matches, a gating process is used to reject unlikely matches early on. Each possible projected vehicle state defines a validation region. If a measurement overlaps with the validation gate, there is a possible match between the feature estimated by the state and that measurement. Each possible projected state produces a set of gated states which indicate possible matches with measurements. Figure 2 illustrates the process of gating.

B.3. HYPOTHESIS FORMATION

A hypothesis is a possible match linking a vehicle (sensor), a relative position measurement, and a feature track. Once the possible projected states have gated the measurements, their decisions are summarized as a set of possible hypotheses. If any combination of a possible projected state of the vehicle and a possible projected state for a given feature are consistent with a measurement, a hypothesis explaining the measurement as coming from that feature is made. Additional hypotheses are made to cover the possibilities that a measurement may correspond to a previously unmodeled feature, that a measurement is a false alarm, or that a feature gives rise to no observed measurement.

B.4. UPDATING STATE ESTIMATES

Each possible gated state for each feature can then produce a set of updated state estimates. There can be more than one of these updated states for each gated state because there can be more than one possible projected state for the vehicle which is consistent with the measurement and that gated state. Figure 3 illustrates the process of updating a gated state for a feature. The process of updating vehicle state is a bit more complex, since multiple measurements are being evaluated at the same time. Updating vehicle state estimates therefore depends on global assignment decisions.

B.5. THE ASSIGNMENT DECISION

As stated above, the reason for pursuing a multiple hypothesis tracking approach to the problem is to make theoretically consistent assignment decisions. This means having an explanation for what happens to each postulated feature and where each measurement comes from. A hypothesis provides an explanation for a single measurement and/or a single feature. A global assignment is a set of coherent exhaustively explanatory hypotheses. In other words, it is a set of hypotheses which explains every measurement and feature, and each measurement and feature is only explained by one hypothesis in the set. At each time step, there may be a number of viable assignments. The primary method of deciding which possible states represent the real world, the underlying basis for our estimation, is to compare the likelihoods of global assignments. When an assignment is discounted in favor of other more likely assignments, the hypotheses which depend upon it, along with the possible states which depend on those hypotheses are pruned from the trees of possible vehicle and feature estimates, so that, in time, a single track is left for the vehicle and each prospective feature track.

III. IMPLEMENTATION PLAN

Our investigation of concurrent mapping and localization is directed to support the non-traditional navigation requirements of the LDUUV and 21UUV programs of the Naval Undersea Warface Center (NUWC) in Newport, Rhode Island, funded by the Office of Naval Research [4]. The implementation plan has four components: algorithm development in simulation, testing with real data, experimentation in controlled environments, and implementation at sea on-board NUWC UUV systems.

The first step in the implementation of concurrent mapping and localization (CM&L) is to create a simulation environment that captures the essential elements of the problem: navigation error, sensor noise, assignment ambiguity, and spurious data. Additional desirable features include simulation of the UUV platform dynamics, ocean environment dynamics, and forward-look sonar beam geometry. The simulation is written in C++ on Silicon Graphics workstations, and makes heavy use of generic data structures and algorithms supplied by the Standard Template Library. An example run of the simulator is shown in Figure 4.

The next stage of implementation is to test the code by post-processing of real data. To do this, we are collaborating with Carpenter and Medeiros in the collection of sonar data sets and navigation data in a realistic operational environment in Naragansett Bay [6]. The sensors available for the tests will include NUWC's high resolution array (HRA) forward-look sonar, a Klein side-scan sonar, an Imagenex sector-scan sonar, an EDO doppler sonar, and an inertial navigation system.

In addition, work is underway at MIT to create a dynamic sonar positioning system, to permit experimentation in a controlled tank setting. Earlier work in such an environment that applied multiple hypothesis tracking to acoustic scene reconstruction with accurate navigation is reported in Moran *et al.* [7].

Finally, our plan is to implement CM&L at sea on-board one of the NUWC UUV Systems. Initial tests will focus on performing a "navigation reset" relative to a single feature (a passive "beacon" or a distinctive bottom feature) during a UUV run. Subsequent work will attempt to implement the full process of building a map and using that map to navigate. Ultimately, our intention is to move beyond point-like, compact features to consider use of natural terrain features such as contours [3] for concurrent mapping and localization.

IV. RESEARCH ISSUES

Several research issues pose particular challenges in formulating and implementing concurrent mapping and localization. In this section, we examine several of these issues, consider how they are currently being addressed by the state of the art, and indicate our progress in resolving them.

A. FEATURE REPRESENTATION, DETECTION, AND TRACKING

There are a number of problems dealing with how to implement an algorithm which refers to environmental features. What kind of features should be used? How can they be represented? How can they be detected and measured? How should models for new features be initialized? How can features be tracked? How should we update a feature estimate given some new information? This host of issues is referred to as feature ontology, or feature representation, detection, and tracking.

There are a number of factors making feature ontology difficult to come to grips with. Perhaps the most important is the representation of uncertainty. Any model developed will contain uncertainty due to ambiguity in data association, vehicle navigational uncertainty, and measurement errors. If we are to use feature estimates as the basis of vehicle control decisions (such as vehicle navigation and path planning), we need not only feature estimates, but accurate assessments of the uncertainty of these estimates. In the extreme case, improper uncertainty estimation can lead to divergent estimates.

Feature ontology is affected by three fundamental algorithmic factors: the underlying representation of the feature, the sensor or sensors used to measure the feature, and the estimation process for updating feature estimates. The deep interconnectedness of these three factors makes generalizing ontological solutions difficult.

A third factor complicating feature representation is its relation to map representation. Feature representation schemes can be broadly separated into two categories: cell-based and feature-based representations. For cell-based feature representations, space is divided into cells, and sensor returns are probabilistically backscattered onto the cells. In such a case, feature representation and map representation are indistinct. In a feature-based representation, distinct features are represented on a map. In this case, the feature and map representations are separable, so that a given feature representation may be used with several different kinds of maps.

Finally, feature ontology can be complicated by the used of multiple vehicle or feature models. Using multiple models is an attractive method for dealing with non-parametric choices in feature representations. For example, a moving target may be estimated using two models, one for a straight trajectory and the other for maneuvering. Developing the ontological basis for horizontal model fusion, particularly when uncertainty representation is necessary, can simplify feature detection and tracking, but at the cost of complexity in feature representation and estimation.

There has been a fair amount of preliminary work in feature representation. However, most of this research is of more restricted scope, and does not readily extend itself to feature ontology for concurrent mapping and localization. Stewart [13] is a recent example of a cell-based representation. He includes the effects of uncertainty, but depends on field-like features; the problems of detection and tracking are not addressed. Kamgar-Parse [3] considers feature-based tracking using an *a priori* map, but does not deal with feature detection and track initialization. Smith, *et al.* [12] develop uncertainty representation within a Kalman framework, but fail to generalize feature representation and detection. Their algorithm is also sensitive to the highly nonlinear transforms used in relative positioning.

Our research extends current feature representation, detection, and tracking technology in two ways. First, we are formulating a general approach to feature and map representation, so that different ontologies can be used transparently in the larger concurrent mapping and localization framework. Second, we are developing an appropriate and accurate uncertainty representation for a number of representational approaches.

We have extended cell-based representation of bathymetric fields to include detection and tracking with a point sensor (such as an altimeter). This representation is currently being modified for feature-based representation of objects within a cell-based map representation. We have generalized feature-based representations to include uncertainty with a Bayesian estimation scheme. Finally, we are linking feature ontology to the rest of the concurrent mapping and localization problem, and, in particular, with data association decisions.

B. DATA ASSOCIATION

Data association refers to the problems of identifying what features are being sensed, associating individual measurements with postulated features, and coherently explaining features and measurements over time. This problem is difficult because it is inherently computationally intense; it may be difficult to gate and cluster the measurements; it is difficult to combine uncertainties from measurement errors, vehicle navigational uncertainty, and feature estimation; feature and vehicle models themselves may be uncertain; and the algorithm must be robust to occasional bad data association decisions.

Uhlmann [16] lays the groundwork for making batch data association decisions. However, his approach does not apply to a recursive algorithm, it retains extreme computational intensity, and it is an incomplete solution to data association. Mori *et al.* [8] have provided a theoretical basis for using multiple hypothesis tracking for data association. However, their formulation does not account for navigational uncertainty, which adds significant intrinsic complexity to the problem.

We have extended multiple hypothesis tracking to incorporate navigational uncertainty. This algorithm is being modified to reduce the computational load while maintaining robustness. We are also formulating an explicit probabilistic technique for incorporating model uncertainty into data association decisions. Finally, we are developing techniques to allow nowcasting for control and planning.

C. STATE ESTIMATION

The problem of state estimation is to update an estimate of vehicle and feature estimates properly in the presence of new information. Although this is a somewhat standard problem, there are several factors introduced by concurrent mapping and localization which complicate it. First, accurate estimation of error covariance is important because it will form the basis for data association and high-level control decisions. Second, there may be unknown correlations in the estimates and measurements. This can lead to filter overconfidence, causing inaccurate uncertainty estimation or even divergence. Third, since relative position (range and bearing) measurements are being used, the filter is highly nonlinear. It is important that in accommodating these nonlinearities, the covariance estimate is not unacceptably degraded. Finally, the plant models used to predict vehicle and feature behavior may themselves be uncertain.

The usual approach to a state estimation problem of this sort is the extended Kalman filter. Other, less-widely-used techniques include iterated filters, higher-order filters, and stochastic linearizations. Although these alternative approaches may improve performance (at some computational penalty) for many situations, we will contrast our approach to the most-used implementation, the extended Kalman filter. The problems with the extended Kalman filter, and the other approaches, is that nonlinearities degrade their error covariance estimation, they assume known correlations, and they do not address the possibility of model uncertainty.

We are modifying the extended Kalman filter to improve its robustness to nonlinearities. We handle the issue of correlation uncertainty by adaptively managing the correlation information contained in the vehicle and feature estimates. Finally, we are providing an explicit method for handling model uncertainty, both for estimate updating and decision-making. Model uncertainty metrics are also used to restrict information flow within the filter. This is to prevent model uncertainty from spreading between models.

V. CONCLUSIONS AND FUTURE RESEARCH

Current tracking technology has addressed the problem of tracking multiple targets, but breaks down when vehicle navigation contains uncertainty. By incorporating navigational uncertainty into the multi-target tracking problem, concurrent mapping and localization extends the capability of UUVs in situations where navigation errors are significant. Using environmental cues to provide navigation information grounded in the global reference frame, this technology enables vehicles to perform repeatably accurate geographical tasks despite the limitations of INS and dead-reckoning and without resorting to a calibrated, predeployed array.

We have defined the problem of concurrent mapping and localization and the process by which it is accomplished. Due to the added complications of vehicle navigation uncertainty, this extension is not trivial. We also outlined a plan for implementing concurrent mapping and localization in simulation, utilizing real data, in a controlled environment, and at sea. We have identified the primary research issues in concurrent mapping and localization as feature representation, detection, and tracking, data association, and state estimation.

We have focused our research to this point on implementing concurrent mapping and localization in simulation. There remain significant issues to be addressed, particularly the computational efficiency and robustness of the algorithm. We are currently improving the software implementation of the algorithm and beginning the process of collecting data for testing.

REFERENCES

[1] 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.

[2] E. Geyer, P. Creamer, J. D'Appolito, and R. Gains. Characteristics and capabilities of navigation systems for unmanned untethered submersibles. In *Proc. Int. Symp. on Unmanned Untethered Submersible Technology*, pages 320–347, 1987.

[3] B. Kamgar-Parse, L. Rosenblum, F. Pipitone, L. Davis, and J. Jones. Toward an automated system for a correctly registered bathymetric chart. *IEEE J. Ocean Engineering*, 14(4):314–325, October 1989.

[4] E. Levine, D. Connors, R. Shell, T. Gagliardi, and R. Hanson. Oceanographic mapping with navy's large diameter uuv. *Sea Technology*, pages 49–57, 1995.

[5] M. B. May. Gravity navigation. In *Record of the 1978 Position Location and Navigation Symposium*, pages 212–218, San Diego, CA, USA, November 1978.

[6] M. Medeiros and R. Carpenter. High resolution array signal processing for AUVs. In *AUV 96*, pages 10–15, 1996.

[7] B. A. Moran, J. J. Leonard, and C. Chryssostomidis. Curved shape reconstruction using multiple hypothesis tracking. *IEEE J. Ocean Engineering*, 1996. To Appear.

[8] S. Mori, C. Chong, E. Tse, and R. Wishner. Tracking and classifying multiple targets without a priori identification. *IEEE Transactions on Automatic Control*, AC-31(5), May 1986.

[9] 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.

[10] D. B. Reid. An algorithm for tracking multiple targets. *IEEE Transactions on Automatic Control*, AC-24(6), December 1979.

[11] W. D. Rencken. Concurrent localisation and map building for mobile robots using ultrasonic sensors. In *Proc. IEEE Int. Workshop on Intelligent Robots and Systems*, pages 2192–2197, Yokohama, Japan, 1993.

[12] R. Smith, M. Self, and P. Cheeseman. Estimating uncertain spatial relationships in robotics. In I. Cox and G. Wilfong, editors, *Autonomous Robot Vehicles*. Springer-Verlag, 1990.

[13] W. K. Stewart. *Multisensor Modeling Underwater with Uncertain Information*. PhD thesis, Massachusetts Institute of Technology, 1988.

[14] S. T. Tuohy, J. J. Leonard, J. G. Bellingham, N. M. Patrikalakis, and C. Chryssostomidis. Map based navigation for autonomous underwater vehicles. *International Journal of Offshore and Polar Engineering*, 6(1):9–18, March 1996.

[15] C. Tyren. Magnetic anomalies as a reference for ground-speed and map-matching navigation. *The Journal of Navigation*, 35(2):242–254, May 1982.

[16] J. Uhlmann. *Dynamic Map Building and Localization: New Theoretical Foundations*. PhD thesis, University of Oxford, 1995.
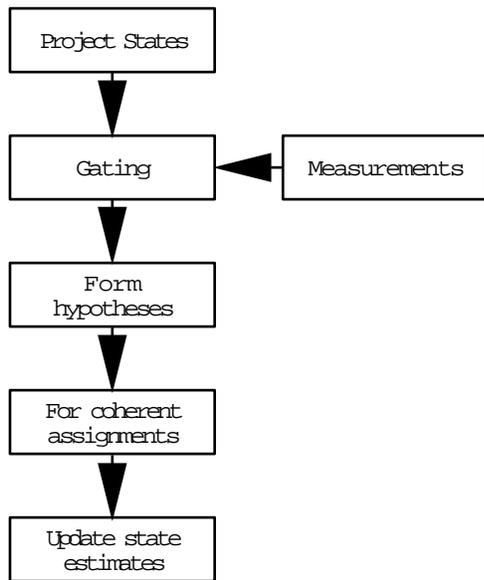
Figure 1. Process flowchart for a concurrent mapping and localization decision cycle. Because measurements involve two stochastic models, the vehicle and the feature, measurement gating and state updating must be separated.
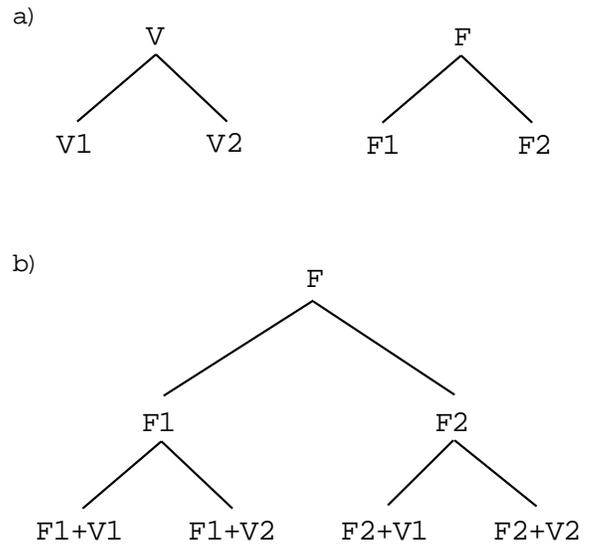

Figure 2. Gating measurements. Two prospective features, F1 and F2, are begin tracked. Three measurements, Z1, Z2, and Z3, are obtained in the current timestep. a) Feature estimate uncertainty is used to gate the measurements. Measurements Z1 and Z2 possibly arise from F1. Measurements Z2 and Z3 possibly arise from feature F2. b) Feature tracks are grown to accommodate each of the gated possibilities.

Figure 3. Track tree growth for state updating. a) A vehicle V and a feature F are hypothesized to be linked by a measurement. Each has two proposed states. b) During updating, all state combinations consistent with the hypothesis are combined to provide possible updated states. For example, feature state F1 can be updated with the measurement assuming either vehicle state V1 or V2 is correct. Vehicle state updating must be delayed until global assignments have beed formed.
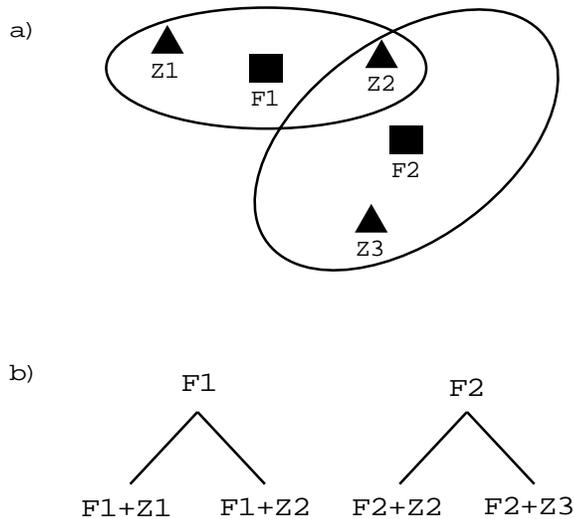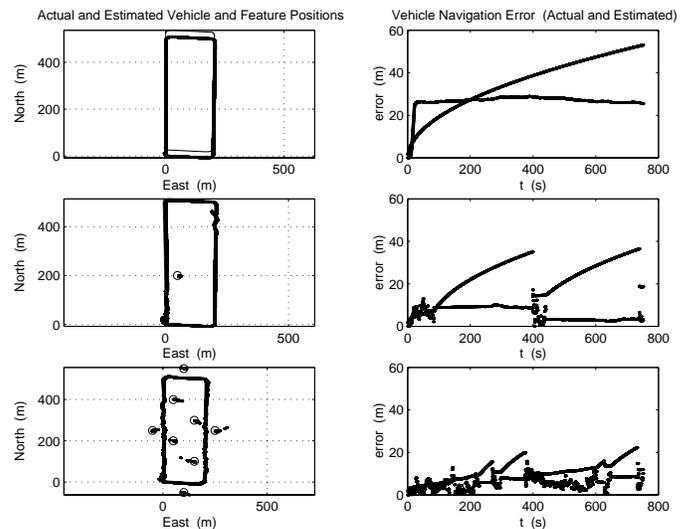

Figure 4. Simulation sample run showing the benefit of using environmental cues to reduce navigational error. In each of the three runs, the vehicle travels in a rectangle; the number of observable features changes in each run. The right graphs show the estimated (thick) and actual (thin) vehicle tracks and the estimated (points) and actual (circles) feature positions. The left graphs show estimated (upper) and actual vehicle navigation error. In the first run, only dead-reckoning is available for navigation information. In the second run, one feature is present. In the third run, eight features are present. The presence of features not only reduces the navigation error, but also improves the uncertainty estimate.