

Simultaneous Localization and Mapping: Part I

BY HUGH DURRANT-WHYTE AND TIM BAILEY

The simultaneous localization and mapping (SLAM) problem asks if it is possible for a mobile robot to be placed at an unknown location in an unknown environment and for the robot to incrementally build a consistent map of this environment while simultaneously determining its location within this map. A solution to the SLAM problem has been seen as a “holy grail” for the mobile robotics community as it would provide the means to make a robot truly autonomous.

The “solution” of the SLAM problem has been one of the notable successes of the robotics community over the past decade. SLAM has been formulated and solved as a theoretical problem in a number of different forms. SLAM has also been implemented in a number of different domains from indoor robots to outdoor, underwater, and airborne systems. At a theoretical and conceptual level, SLAM can now be considered a solved problem. However, substantial issues remain in practically realizing more general SLAM solutions and notably in building and using perceptually rich maps as part of a SLAM algorithm.

This two-part tutorial and survey of SLAM aims to provide a broad introduction to this rapidly growing field. Part I (this article) begins by providing a brief history of early developments in SLAM. The formulation section introduces the structure the SLAM problem in now standard Bayesian form, and explains the evolution of the SLAM process. The solution section describes the two key computational solutions to the SLAM problem through the use of the extended Kalman filter (EKF-SLAM) and through the use of Rao-Blackwellized particle filters (FastSLAM). Other recent solutions to the SLAM problem are discussed in Part II of this tutorial. The application section describes a number of important real-world implementations of SLAM and also highlights implementations where the sensor data and software are freely down-loadable for other researchers to study. Part II of this tutorial describes major issues in computation, convergence, and data association in SLAM. These are subjects that have been the main focus of the SLAM research community over the past five years.

History of the SLAM Problem

The genesis of the probabilistic SLAM problem occurred at the 1986 IEEE Robotics and Automation Conference held in San Francisco, California. This was a time when probabilistic

methods were only just beginning to be introduced into both robotics and artificial intelligence (AI). A number of researchers had been looking at applying estimation-theoretic methods to mapping and localization problems; these included Peter Cheeseman, Jim Crowley, and Hugh Durrant-Whyte. Over the course of the conference, many paper table cloths and napkins were filled with long discussions about consistent mapping. Along the way, Raja Chatila, Oliver Faugeras, Randal Smith, and others also made useful contributions to the conversation.

The result of this conversation was a recognition that consistent probabilistic mapping was a fundamental problem in robotics with major conceptual and computational issues that needed to be addressed. Over the next few years a number of key papers were produced. Work by Smith and Cheesman [39] and Durrant-Whyte [17] established a statistical basis for describing relationships between landmarks and manipulating geometric uncertainty. A key element of this work was to show that there must be a high degree of correlation between estimates of the location of different landmarks in a map and that, indeed, these correlations would grow with successive observations.

At the same time Ayache and Faugeras [1] were undertaking early work in visual navigation, Crowley [9] and Chatila and Laumond [6] were working in sonar-based navigation of mobile robots using Kalman filter-type algorithms. These two strands of research had much in common and resulted soon after in the landmark paper by Smith et al. [40]. This paper showed that as a mobile robot moves through an unknown environment taking relative observations of landmarks, the estimates of these landmarks are all necessarily correlated with each other because of the common error in estimated vehicle location [27]. The implication of this was profound: A consistent full solution to the combined localization and mapping problem would require a joint state composed of the vehicle pose and every landmark position, to be updated following each landmark observation. In turn, this would require the estimator to employ a huge state vector (on the order of the number of landmarks maintained in the map) with computation scaling as the square of the number of landmarks.

Crucially, this work did not look at the convergence properties of the map or its steady-state behavior. Indeed, it

was widely assumed at the time that the estimated map errors would not converge and would instead exhibit a random-walk behavior with unbounded error growth. Thus, given the computational complexity of the mapping problem and without knowledge of the convergence behavior of the map, researchers instead focused on a series of approximations to the consistent mapping problem, which assumed or even forced the correlations between landmarks to be minimized or eliminated, so reducing the full filter to a series of decoupled landmark to vehicle filters ([28] and [38] for example). Also for these reasons, theoretical work on the combined localization and mapping problem came to a temporary halt, with work often focused on either mapping or localization as separate problems.

The conceptual breakthrough came with the realization that the combined mapping and localization problem, once formulated as a single estimation problem, was actually convergent. Most importantly, it was recognized that the correlations between landmarks, which most researchers had tried to minimize, were actually the critical part of the problem and that, on the contrary, the more these correlations grew, the better the solution. The structure of the SLAM problem, the convergence result and the coining of the acronym SLAM was first presented in a mobile robotics survey paper presented at the 1995 International Symposium on Robotics Research [16]. The essential theory on convergence and many of the initial results were developed by Csorba [10],

[11]. Several groups already working on mapping and localization, notably at the Massachusetts Institute of Technology [29], Zaragoza [4], [5], the ACFR at Sydney [20], [45], and others [7], [13], began working in earnest on SLAM—also called *concurrent mapping and localization* (CML) at this time—in indoor, outdoor, and subsea environments.

At this time, work focused on improving computational efficiency and addressing issues in data association, or loop closure. The 1999 International Symposium on Robotics Research (ISRR'99) [23] was an important meeting point where the first SLAM session was held and where a degree of convergence between the Kalman-filter-based SLAM methods and the probabilistic localisation and mapping methods introduced by Thrun [42] was achieved. The 2000 IEEE International Conference on Robotics and Automation (ICRA) Workshop on SLAM attracted 15 researchers and focused on issues such as algorithmic complexity, data association, and implementation challenges. The following SLAM workshop at the 2002 ICRA attracted 150 researchers with a broad range of interests and applications. The 2002 SLAM summer school hosted by Henrik Christiansen at KTH in Stockholm attracted all the key researchers together with some 50 Ph.D. students from around the world and was a tremendous success in building the field. Interest in SLAM has grown exponentially in recent years, and workshops continue to be held at both ICRA and IROS. The SLAM summer school ran in 2004 in Toulouse and will run at Oxford, England, in 2006.

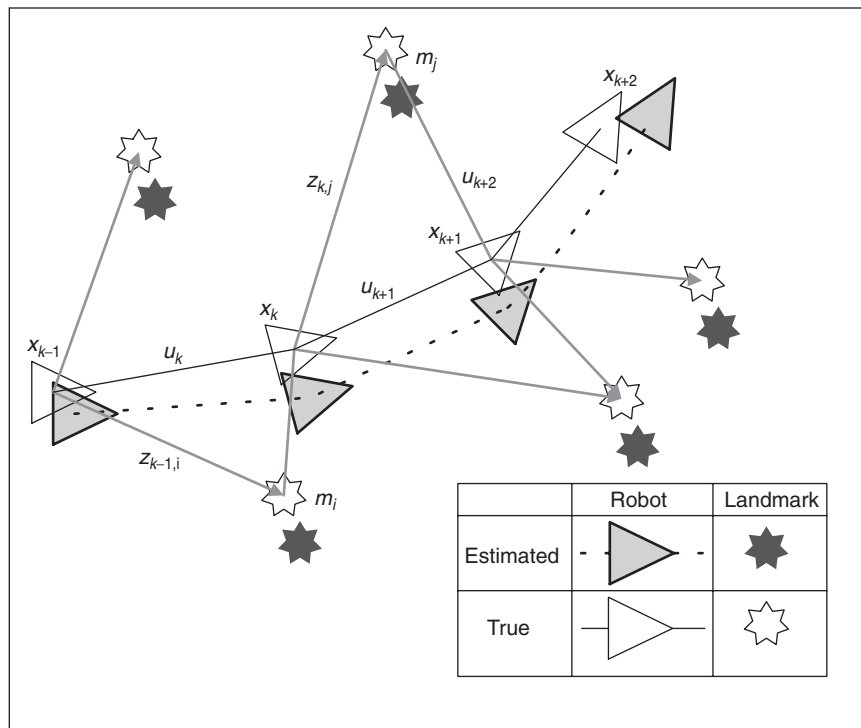


Figure 1. The essential SLAM problem. A simultaneous estimate of both robot and landmark locations is required. The true locations are never known or measured directly. Observations are made between true robot and landmark locations.

Formulation and Structure of the SLAM Problem

SLAM is a process by which a mobile robot can build a map of an environment and at the same time use this map to deduce its location. In SLAM, both the trajectory of the platform and the location of all landmarks are estimated online without the need for any a priori knowledge of location.

Preliminaries

Consider a mobile robot moving through an environment taking relative observations of a number of unknown landmarks using a sensor located on the robot as shown in Figure 1. At a time instant k , the following quantities are defined:

- ◆ \mathbf{x}_k : the state vector describing the location and orientation of the vehicle
- ◆ \mathbf{u}_k : the control vector, applied at time $k - 1$ to drive the vehicle to a state \mathbf{x}_k at time k

- ◆ \mathbf{m}_i : a vector describing the location of the i th landmark whose true location is assumed time invariant
- ◆ \mathbf{z}_{ik} : an observation taken from the vehicle of the location of the i th landmark at time k . When there are multiple landmark observations at any one time or when the specific landmark is not relevant to the discussion, the observation will be written simply as \mathbf{z}_k .

In addition, the following sets are also defined:

- ◆ $\mathbf{X}_{0:k} = \{\mathbf{x}_0, \mathbf{x}_1, \dots, \mathbf{x}_k\} = \{\mathbf{X}_{0:k-1}, \mathbf{x}_k\}$: the history of vehicle locations
- ◆ $\mathbf{U}_{0:k} = \{\mathbf{u}_1, \mathbf{u}_2, \dots, \mathbf{u}_k\} = \{\mathbf{U}_{0:k-1}, \mathbf{u}_k\}$: the history of control inputs
- ◆ $\mathbf{m} = \{\mathbf{m}_1, \mathbf{m}_2, \dots, \mathbf{m}_n\}$ the set of all landmarks
- ◆ $\mathbf{Z}_{0:k} = \{\mathbf{z}_1, \mathbf{z}_2, \dots, \mathbf{z}_k\} = \{\mathbf{Z}_{0:k-1}, \mathbf{z}_k\}$: the set of all landmark observations.

Probabilistic SLAM

In probabilistic form, the simultaneous localization and map building (SLAM) problem requires that the probability distribution

$$P(\mathbf{x}_k, \mathbf{m} | \mathbf{Z}_{0:k}, \mathbf{U}_{0:k}, \mathbf{x}_0) \quad (1)$$

be computed for all times k . This probability distribution describes the *joint* posterior density of the landmark locations and vehicle state (at time k) given the recorded observations and control inputs up to and including time k together with the initial state of the vehicle. In general, a recursive solution to the SLAM problem is desirable. Starting with an estimate for the distribution $P(\mathbf{x}_{k-1}, \mathbf{m} | \mathbf{Z}_{0:k-1}, \mathbf{U}_{0:k-1})$ at time $k-1$, the joint posterior, following a control \mathbf{u}_k and observation \mathbf{z}_k , is computed using Bayes theorem. This computation requires that a state transition model and an observation model are defined describing the effect of the control input and observation respectively.

The *observation model* describes the probability of making an observation \mathbf{z}_k when the vehicle location and landmark locations are known and is generally described in the form

$$P(\mathbf{z}_k | \mathbf{x}_k, \mathbf{m}). \quad (2)$$

It is reasonable to assume that once the vehicle location and map are defined, observations are conditionally independent given the map and the current vehicle state.

The *motion model* for the vehicle can be described in terms of a probability distribution on state transitions in the form

$$P(\mathbf{x}_k | \mathbf{x}_{k-1}, \mathbf{u}_k). \quad (3)$$

That is, the state transition is assumed to be a Markov process in which the next state \mathbf{x}_k depends only on the immediately preceding state \mathbf{x}_{k-1} and the applied control \mathbf{u}_k and is independent of both the observations and the map.

The SLAM algorithm is now implemented in a standard two-step recursive (sequential) prediction (time-update) correction (measurement-update) form:

Time-update

$$P(\mathbf{x}_k, \mathbf{m} | \mathbf{Z}_{0:k-1}, \mathbf{U}_{0:k}, \mathbf{x}_0) = \int P(\mathbf{x}_k | \mathbf{x}_{k-1}, \mathbf{u}_k) \times P(\mathbf{x}_{k-1}, \mathbf{m} | \mathbf{Z}_{0:k-1}, \mathbf{U}_{0:k-1}, \mathbf{x}_0) d\mathbf{x}_{k-1} \quad (4)$$

Measurement Update

$$P(\mathbf{x}_k, \mathbf{m} | \mathbf{Z}_{0:k}, \mathbf{U}_{0:k}, \mathbf{x}_0) = \frac{P(\mathbf{z}_k | \mathbf{x}_k, \mathbf{m}) P(\mathbf{x}_k, \mathbf{m} | \mathbf{Z}_{0:k-1}, \mathbf{U}_{0:k}, \mathbf{x}_0)}{P(\mathbf{z}_k | \mathbf{Z}_{0:k-1}, \mathbf{U}_{0:k})} \quad (5)$$

Equations (4) and (5) provide a recursive procedure for calculating the joint posterior $P(\mathbf{x}_k, \mathbf{m} | \mathbf{Z}_{0:k}, \mathbf{U}_{0:k}, \mathbf{x}_0)$ for the robot state \mathbf{x}_k and map \mathbf{m} at a time k based on all observations $\mathbf{Z}_{0:k}$ and all control inputs $\mathbf{U}_{0:k}$ up to and including time k . The recursion is a function of a vehicle model $P(\mathbf{x}_k | \mathbf{x}_{k-1}, \mathbf{u}_k)$ and an observation model $P(\mathbf{z}_k | \mathbf{x}_k, \mathbf{m})$.

It is worth noting that the map building problem may be formulated as computing the conditional density $P(\mathbf{m} | \mathbf{X}_{0:k}, \mathbf{Z}_{0:k}, \mathbf{U}_{0:k})$. This assumes that the location of the vehicle \mathbf{x}_k is known (or at least deterministic) at all times, subject to knowledge of initial location. A map \mathbf{m} is then constructed by fusing observations from different locations. Conversely, the localization problem may be formulated as computing the probability distribution $P(\mathbf{x}_k | \mathbf{Z}_{0:k}, \mathbf{U}_{0:k}, \mathbf{m})$. This assumes that the landmark locations are known with certainty, and the objective is to compute an estimate of vehicle location with respect to these landmarks.

Structure of Probabilistic SLAM

To simplify the discussion in this section, we will drop the conditioning on historical variables in (1) and write the required joint posterior on map and vehicle location as $P(\mathbf{x}_k, \mathbf{m} | \mathbf{z}_k)$ and even $P(\mathbf{x}_k, \mathbf{m})$ as the context permits.

The observation model $P(\mathbf{z}_k | \mathbf{x}_k, \mathbf{m})$ makes explicit the dependence of observations on both the vehicle and landmark locations. It follows that the joint posterior cannot be partitioned in the obvious manner

$$P(\mathbf{x}_k, \mathbf{m} | \mathbf{z}_k) \neq P(\mathbf{x}_k | \mathbf{z}_k) P(\mathbf{m} | \mathbf{z}_k),$$

and indeed it is well known from the early papers on consistent mapping [17], [39] that a partition such as this leads to inconsistent estimates. However, the SLAM problem has more structure than is immediately obvious from these equations.

Referring again to Figure 1, it can be seen that much of the error between estimated and true landmark locations is common between landmarks and is in fact due to a single

source; errors in knowledge of where the robot is when landmark observations are made. In turn, this implies that the errors in landmark location estimates are highly correlated. Practically, this means that the relative location between any two landmarks, $\mathbf{m}_i - \mathbf{m}_j$, may be known with high accuracy, even when the absolute location of a landmark \mathbf{m}_i is quite uncertain. In probabilistic form, this means that the joint probability density for the pair of landmarks $P(\mathbf{m}_i, \mathbf{m}_j)$ is highly peaked even when the marginal densities $P(\mathbf{m}_i)$ may be quite dispersed.

The most important insight in SLAM was to realize that the correlations between landmark estimates increase *monotonically* as more and more observations are made. (These results have only been proved for the linear Gaussian case [14]. Formal proof for the more general probabilistic case remains an open problem.) Practically, this means that knowledge of the relative location of landmarks always improves and never diverges, regardless of robot motion. In probabilistic terms, this means that the joint probability density on all landmarks $P(\mathbf{m})$ becomes monotonically more peaked as more observations are made.

This convergence occurs because the observations made by the robot can be considered as “nearly independent” measurements of the *relative* location between landmarks. Referring again to Figure 1, consider the robot at location \mathbf{x}_k observing the two landmarks \mathbf{m}_i and \mathbf{m}_j . The relative location of observed landmarks is clearly independent of the coordinate frame of the vehicle, and successive observations from this fixed location would yield further independent measurements of the relative relationship between landmarks. Now, as

the robot moves to location \mathbf{x}_{k+1} , it again observes landmark \mathbf{m}_j this allows the estimated location of the robot and landmark to be updated relative to the previous location \mathbf{x}_k . In turn, this propagates back to update landmark \mathbf{m}_i —even though this landmark is not seen from the new location. This occurs because the two landmarks are highly correlated (their relative location is well known) from previous measurements. Further, the fact that the same measurement data is used to update these two landmarks makes them *more* correlated. The term *nearly independent measurement* is appropriate because the observation errors will be correlated through successive vehicle motions. Also note that in Figure 1 at location \mathbf{x}_{k+1} , the robot observes two new landmarks *relative* to \mathbf{m}_j . These new landmarks are thus immediately linked or correlated to the rest of the map. Later updates to these landmarks will also update landmark \mathbf{m}_j and through this landmark \mathbf{m}_i and so on. That is, all landmarks end up forming a network linked by relative location or correlations whose precision or value increases whenever an observation is made.

This process can be visualized (Figure 2) as a network of springs connecting all landmarks together or as a rubber sheet in which all landmarks are embedded. An observation in a neighborhood acts like a displacement to a spring system or rubber sheet such that its effect is great in the neighborhood and, dependent on local stiffness (correlation) properties, diminishes with distance to other landmarks. As the robot moves through this environment and takes observations of the landmarks, the springs become increasingly (and monotonically) stiffer. In the limit, a rigid map of landmarks or an accurate *relative* map of the environment is obtained. As the map is built, the location accuracy of the robot measured relative to the map is bounded only by the quality of the map and relative measurement sensor. In the theoretical limit, robot relative location accuracy becomes equal to the localization accuracy achievable with a given map.

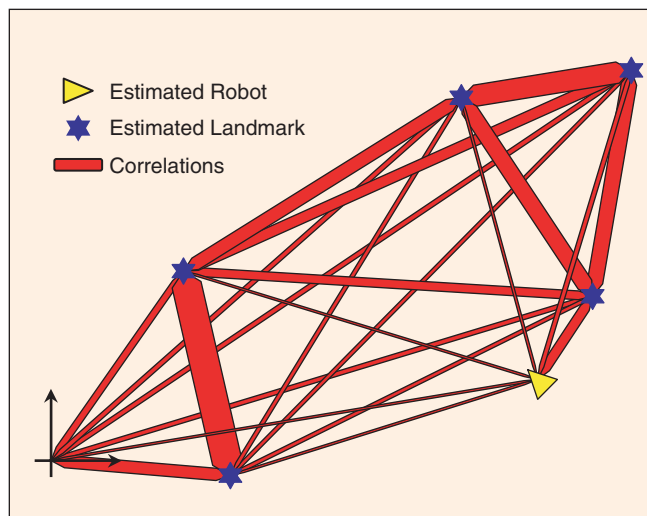


Figure 2. Spring network analogy. The landmarks are connected by springs describing correlations between landmarks. As the vehicle moves back and forth through the environment, spring stiffness or correlations increase (red links become thicker). As landmarks are observed and estimated locations are corrected, these changes are propagated through the spring network. Note, the robot itself is correlated to the map.

Solutions to the SLAM Problem

Solutions to the probabilistic SLAM problem involve finding an appropriate representation for both the observation model (2) and motion model (3) that allows efficient and consistent computation of the prior and posterior distributions in (4) and (5). By far, the most common representation is in the form of a state-space model with additive Gaussian noise, leading to the use of the extended Kalman filter (EKF) to solve the SLAM problem. One important alternative representation is to describe the vehicle motion model in (3) as a set of samples of a more general non-Gaussian probability distribution. This leads to the use of the Rao-Blackwellized particle filter, or FastSLAM algorithm, to solve the SLAM problem. While EKF-SLAM and FastSLAM are the two most important solution methods, newer alternatives, which offer much potential, have been proposed, including the use of the information-state form [43]. These are discussed further in Part II of this tutorial.

EKF-SLAM

The basis for the EKF-SLAM method is to describe the vehicle motion in the form

$$P(\mathbf{x}_k | \mathbf{x}_{k-1}, \mathbf{u}_k) \iff \mathbf{x}_k = \mathbf{f}(\mathbf{x}_{k-1}, \mathbf{u}_k) + \mathbf{w}_k, \quad (6)$$

where $\mathbf{f}(\cdot)$ models vehicle kinematics and where \mathbf{w}_k are additive, zero mean uncorrelated Gaussian motion disturbances with covariance \mathbf{Q}_k . The observation model is described in the form

$$P(\mathbf{z}_k | \mathbf{x}_k, \mathbf{m}) \iff \mathbf{z}_k = \mathbf{h}(\mathbf{x}_k, \mathbf{m}) + \mathbf{v}_k, \quad (7)$$

where $\mathbf{h}(\cdot)$ describes the geometry of the observation and where \mathbf{v}_k are additive, zero mean uncorrelated Gaussian observation errors with covariance \mathbf{R}_k .

With these definitions, the standard EKF method [14], [31] can be applied to compute the mean

$$\begin{bmatrix} \hat{\mathbf{x}}_{k|k} \\ \hat{\mathbf{m}}_k \end{bmatrix} = \mathbb{E} \begin{bmatrix} \mathbf{x}_k \\ \mathbf{m} \end{bmatrix} | \mathbf{Z}_{0:k},$$

and covariance

$$\begin{aligned} \mathbf{P}_{k|k} &= \begin{bmatrix} \mathbf{P}_{xx} & \mathbf{P}_{xm} \\ \mathbf{P}_{xm}^T & \mathbf{P}_{mm} \end{bmatrix}_{k|k} \\ &= \mathbb{E} \left[\begin{pmatrix} \mathbf{x}_k - \hat{\mathbf{x}}_k \\ \mathbf{m} - \hat{\mathbf{m}}_k \end{pmatrix} \begin{pmatrix} \mathbf{x}_k - \hat{\mathbf{x}}_k \\ \mathbf{m} - \hat{\mathbf{m}}_k \end{pmatrix}^T \middle| \mathbf{Z}_{0:k} \right] \end{aligned}$$

of the joint posterior distribution $P(\mathbf{x}_k, \mathbf{m} | \mathbf{Z}_{0:k}, \mathbf{U}_{0:k}, \mathbf{x}_0)$ from:

Time-update

$$\hat{\mathbf{x}}_{k|k-1} = \mathbf{f}(\hat{\mathbf{x}}_{k-1|k-1}, \mathbf{u}_k) \quad (8)$$

$$\mathbf{P}_{xx,k|k-1} = \nabla \mathbf{f} \mathbf{P}_{xx,k-1|k-1} \nabla \mathbf{f}^T + \mathbf{Q}_k, \quad (9)$$

where $\nabla \mathbf{f}$ is the Jacobian of \mathbf{f} evaluated at the estimate $\hat{\mathbf{x}}_{k-1|k-1}$. There is generally no need to perform a time-update for stationary landmarks. (A time-update, however, is necessary for landmarks that may move [44].)

Observation-update

$$\begin{bmatrix} \hat{\mathbf{x}}_{k|k} \\ \hat{\mathbf{m}}_k \end{bmatrix} = [\hat{\mathbf{x}}_{k|k-1} \hat{\mathbf{m}}_{k-1}] + \mathbf{W}_k [\mathbf{z}_k - \mathbf{h}(\hat{\mathbf{x}}_{k|k-1}, \hat{\mathbf{m}}_{k-1})] \quad (10)$$

$$\mathbf{P}_{k|k} = \mathbf{P}_{k|k-1} - \mathbf{W}_k \mathbf{S}_k \mathbf{W}_k^T, \quad (11)$$

where

$$\begin{aligned} \mathbf{S}_k &= \nabla \mathbf{h} \mathbf{P}_{k|k-1} \nabla \mathbf{h}^T + \mathbf{R}_k \\ \mathbf{W}_k &= \mathbf{P}_{k|k-1} \nabla \mathbf{h}^T \mathbf{S}_k^{-1} \end{aligned}$$

and where $\nabla \mathbf{h}$ is the Jacobian of \mathbf{h} evaluated at $\hat{\mathbf{x}}_{k|k-1}$ and $\hat{\mathbf{m}}_{k-1}$.

This EKF-SLAM solution is very well known and inherits many of the same benefits and problems as the standard EKF solutions to navigation or tracking problems. Four of the key issues are briefly discussed next.

Convergence

In the EKF-SLAM problem, convergence of the map is manifest in the monotonic convergence of the determinant of the map covariance matrix $\mathbf{P}_{mm,k}$ and all landmark pair submatrices, toward zero. The individual landmark variances converge toward a lower bound determined by initial uncertainties in robot position and observations. The typical convergence behavior of landmark location variances is shown in Figure 3 (from [14]).

Computational Effort

The observation update step requires that all landmarks and the joint covariance matrix be updated every time an observation is made. Naively, this means computation grows quadratically with the number of landmarks. There has been a great deal of work undertaken in developing efficient variants of the EKF-SLAM solution and real-time implementations with many thousands of landmarks have been demonstrated [21], [29]. Efficient variants of the EKF-SLAM algorithm are discussed in Part II of this tutorial.

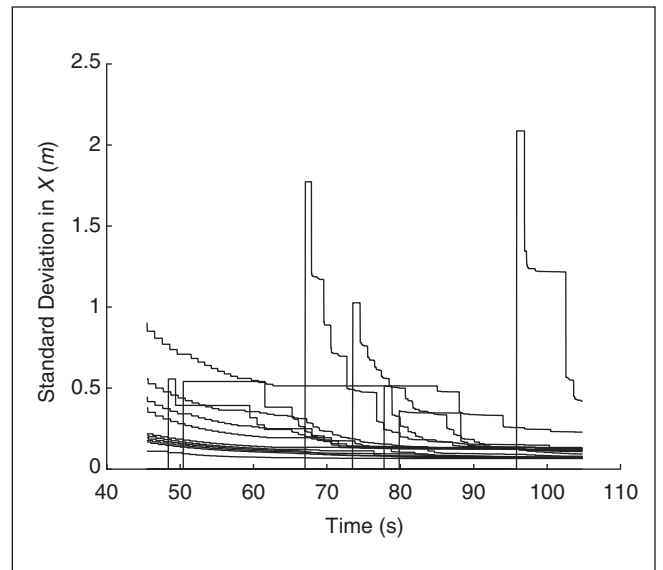


Figure 3. The convergence in landmark uncertainty. The plot shows a time history of standard deviations of a set of landmark locations. A landmark is initially observed with uncertainty inherited from the robot location and observation. Over time, the standard deviations reduce monotonically to a lower bound. New landmarks are acquired during motion (from [14]).

Data Association

The standard formulation of the EKF-SLAM solution is especially fragile to incorrect association of observations to landmarks [35]. The loop-closure problem, when a robot returns to reobserve landmarks after a large traverse, is especially difficult. The association problem is compounded in environments where landmarks are not simple points and indeed look different from different viewpoints. Current work in this area will be described in Part II of this tutorial.

Nonlinearity

EKF-SLAM employs linearized models of nonlinear motion and observation models and so inherits many caveats. Nonlinearity can be a significant problem in EKF-SLAM and leads to inevitable, and sometimes dramatic, inconsistency in solutions [24]. Convergence and consistency can only be guaranteed in the linear case.

Rao-Blackwellized Filter

The FastSLAM algorithm, introduced by Montemerlo et al. [32], marked a fundamental conceptual shift in the design of recursive probabilistic SLAM. Previous efforts focused on improving the performance of EKF-SLAM, while retaining its essential linear Gaussian assumptions. FastSLAM, with its basis in recursive Monte Carlo sampling, or particle filtering, was the first to directly represent the nonlinear process model and non-Gaussian pose distribution. (FastSLAM still linearizes the observation model, but this is typically a reasonable approximation for range-bearing measurements when the vehicle pose is known.) This approach was influenced by earlier probabilistic mapping experiments of Murphy [34] and Thrun [41].

The high dimensional state-space of the SLAM problem makes direct application of particle filters computationally infeasible. However, it is possible to reduce the sample-

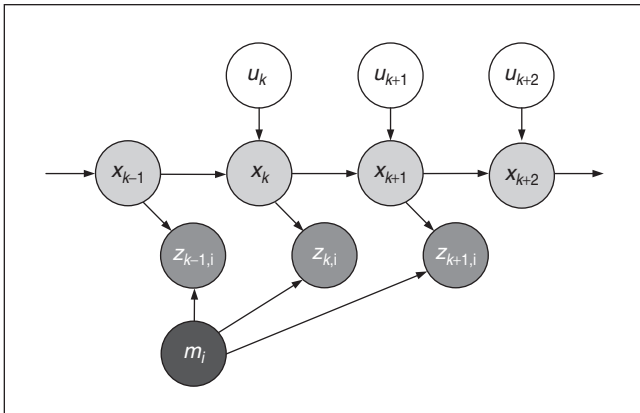


Figure 4. A graphical model of the SLAM algorithm. If the history of pose states are known exactly then, since the observations are conditionally independent, the map states are also independent. For FastSLAM, each particle defines a different vehicle trajectory hypothesis.

space by applying Rao-Blackwellization (R-B), whereby a joint state is partitioned according to the product rule $P(\mathbf{x}_1, \mathbf{x}_2) = P(\mathbf{x}_2|\mathbf{x}_1)P(\mathbf{x}_1)$ and, if $P(\mathbf{x}_2|\mathbf{x}_1)$ can be represented analytically, only $P(\mathbf{x}_1)$ need be sampled $\mathbf{x}_1^{(i)} \sim P(\mathbf{x}_1)$. The joint distribution, therefore, is represented by the set $\{\mathbf{x}_1^{(i)}, P(\mathbf{x}_2|\mathbf{x}_1^{(i)})\}_i^N$ and statistics such as the marginal

$$P(\mathbf{x}_2) \approx \frac{1}{N} \sum_i^N P(\mathbf{x}_2|\mathbf{x}_1^{(i)})$$

can be obtained with greater accuracy than is possible by sampling over the joint space.

The joint SLAM state may be factored into a vehicle component and a conditional map component:

$$\begin{aligned} P(\mathbf{X}_{0:k}, \mathbf{m}|\mathbf{Z}_{0:k}, \mathbf{U}_{0:k}, \mathbf{x}_0) \\ = P(\mathbf{m}|\mathbf{X}_{0:k}, \mathbf{Z}_{0:k})P(\mathbf{X}_{0:k}|\mathbf{Z}_{0:k}, \mathbf{U}_{0:k}, \mathbf{x}_0). \end{aligned} \quad (12)$$

Here the probability distribution is on the trajectory $\mathbf{X}_{0:k}$ rather than the single pose \mathbf{x}_k because, when conditioned on the trajectory, the map landmarks become independent (see Figure 4). This is a key property of FastSLAM and the reason for its speed; the map is represented as a set of independent Gaussians, with linear complexity, rather than a joint map covariance with quadratic complexity.

The essential structure of FastSLAM, then, is a Rao-Blackwellized state, where the trajectory is represented by weighted samples and the map is computed analytically. Thus, the joint distribution, at time k , is represented by the set $\{\mathbf{w}_k^{(i)}, \mathbf{X}_{0:k}^{(i)}, P(\mathbf{m}|\mathbf{X}_{0:k}^{(i)}, \mathbf{Z}_{0:k})\}_i^N$, where the map accompanying each particle is composed of independent Gaussian distributions

$$P(\mathbf{m}|\mathbf{X}_{0:k}^{(i)}, \mathbf{Z}_{0:k}) = \prod_j^M P(\mathbf{m}_j|\mathbf{X}_{0:k}^{(i)}, \mathbf{Z}_{0:k}).$$

Recursive estimation is performed by particle filtering for the pose states and the EKF for the map states.

Updating the map, for a given trajectory particle $\mathbf{X}_{0:k}^{(i)}$, is trivial. Each observed landmark is processed individually as an EKF measurement update from a known pose (see Figure 5). Unobserved landmarks are unchanged. Propagating the pose particles, on the other hand, is more complex, as we discuss below.

We forego giving the background on particle filters, except to say that the theory is derived from a recursive form of sampling known as *sequential important sampling* (SIS) [15], which actually samples from a joint *state history* but “telescopes” the joint into a recursion via the product rule.

$$\begin{aligned} P(\mathbf{x}_0, \mathbf{x}_1, \dots, \mathbf{x}_T|\mathbf{Z}_{0:T}) \\ = P(\mathbf{x}_0|\mathbf{Z}_{0:T})P(\mathbf{x}_1|\mathbf{x}_0, \mathbf{Z}_{0:T}) \dots P(\mathbf{x}_T|\mathbf{X}_{0:T-1}, \mathbf{Z}_{0:T}). \end{aligned}$$

At each time-step k , particles are drawn from a *proposal distribution* $\pi(\mathbf{x}_k|\mathbf{X}_{0:k-1}, \mathbf{Z}_{0:k})$, which approximates the true

distribution $P(\mathbf{x}_k | \mathbf{X}_{0:k-1}, \mathbf{Z}_{0:T})$, and the samples are given importance weights to compensate for any discrepancy. The approximation error grows with time (and inherent joint state space), increasing the variation in sample weights, degrading statistical accuracy. A *resampling* step reinstates uniform weighting, but causes loss of historical particle information. This leads to a crucial property: SIS with resampling can produce reasonable statistics only for systems that “exponentially forget” their past [8] (i.e., systems whose process noise cause the state at time k to become increasingly independent of preceding states).

The general form of a R-B particle filter for SLAM is as follows. We assume that, at time $k - 1$, the joint state is represented by $\{\mathbf{w}_{k-1}^{(i)}, \mathbf{X}_{0:k-1}^{(i)}, P(\mathbf{m} | \mathbf{X}_{0:k-1}^{(i)}, \mathbf{Z}_{0:k-1})\}_i^N$.

- 1) For each particle, compute a proposal distribution, conditioned on the specific particle history, and draw a sample from it

$$\mathbf{x}_k^{(i)} \sim \pi(\mathbf{x}_k | \mathbf{X}_{0:k-1}^{(i)}, \mathbf{Z}_{0:k}, \mathbf{u}_k). \quad (13)$$

This new sample is (implicitly) joined to the particle history $\mathbf{X}_{0:k}^{(i)} \triangleq \{\mathbf{X}_{0:k-1}^{(i)}, \mathbf{x}_k^{(i)}\}$.

- 2) Weight samples according to the importance function

$$w_k^{(i)} = w_{k-1}^{(i)} \frac{P(\mathbf{z}_k | \mathbf{X}_{0:k}^{(i)}, \mathbf{Z}_{0:k-1}) P(\mathbf{x}_k^{(i)} | \mathbf{x}_{k-1}^{(i)}, \mathbf{u}_k)}{\pi(\mathbf{x}_k^{(i)} | \mathbf{X}_{0:k-1}^{(i)}, \mathbf{Z}_{0:k}, \mathbf{u}_k)}. \quad (14)$$

The numerator terms of this equation are the observation model and the motion model, respectively. The former differs from (2) because R-B requires dependency on the map be marginalized away.

$$\begin{aligned} & P(\mathbf{z}_k | \mathbf{X}_{0:k}, \mathbf{Z}_{0:k-1}) \\ &= \int P(\mathbf{z}_k | \mathbf{x}_k, \mathbf{m}) P(\mathbf{m} | \mathbf{X}_{0:k-1}, \mathbf{Z}_{0:k-1}) d\mathbf{m} \end{aligned} \quad (15)$$

- 3) If necessary, perform resampling. (When best to instigate resampling is an open problem. Some implementations resample every time-step, others after a fixed number of time-steps, and others once the weight variance exceeds a threshold.) Resampling is accomplished by selecting particles, with replacement, from the set $\{\mathbf{X}_{0:k}^{(i)}\}_i^N$, including their associated maps, with probability of selection proportional to $w_k^{(i)}$. Selected particles are given uniform weight, $w_k^{(i)} = 1/N$.
- 4) For each particle, perform an EKF update on the observed landmarks as a simple mapping operation with known vehicle pose.

The two versions of FastSLAM in the literature, FastSLAM 1.0 [32] and FastSLAM 2.0 [33], differ only in terms of the form of their proposal distribution (Step 1) and, consequently, in their importance weight (Step 2). FastSLAM 2.0 is by far the more efficient solution.

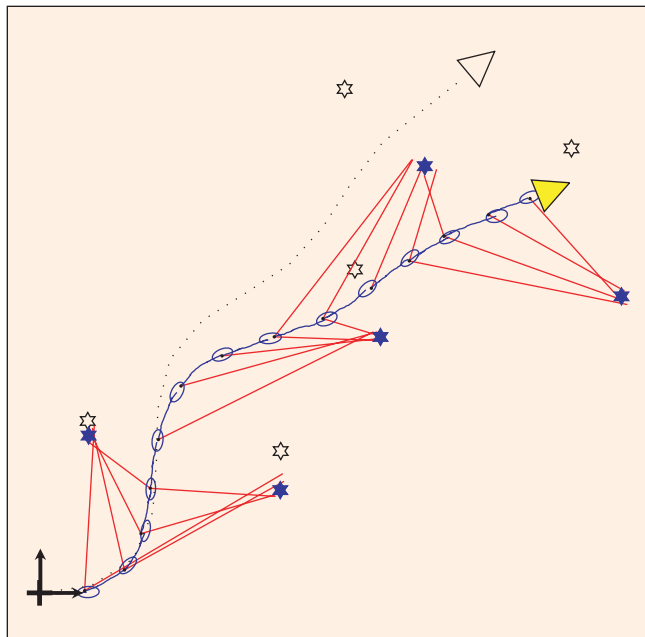


Figure 5. A single realization of robot trajectory in the FastSLAM algorithm. The ellipsoids show the proposal distribution for each update stage, from which a robot pose is sampled, and, assuming this pose is perfect, the observed landmarks are updated. Thus, the map for a single particle is governed by the accuracy of the trajectory. Many such trajectories provide a probabilistic model of robot location.

For FastSLAM 1.0, the proposal distribution is the motion model

$$\mathbf{x}_k^{(i)} \sim P(\mathbf{x}_k | \mathbf{x}_{k-1}^{(i)}, \mathbf{u}_k). \quad (16)$$

Therefore, from (14), the samples are weighted according to the marginalized observation model.

$$w_k^{(i)} = w_{k-1}^{(i)} P(\mathbf{z}_k | \mathbf{X}_{0:k}^{(i)}, \mathbf{Z}_{0:k-1}) \quad (17)$$

For FastSLAM 2.0, the proposal distribution includes the current observation

$$\mathbf{x}_k^{(i)} \sim P(\mathbf{x}_k | \mathbf{X}_{0:k-1}^{(i)}, \mathbf{Z}_{0:k}, \mathbf{u}_k), \quad (18)$$

where

$$\begin{aligned} & P(\mathbf{x}_k | \mathbf{X}_{0:k-1}^{(i)}, \mathbf{Z}_{0:k}, \mathbf{u}_k) \\ &= \frac{1}{C} P(\mathbf{z}_k | \mathbf{x}_k, \mathbf{X}_{0:k-1}^{(i)}, \mathbf{Z}_{0:k-1}) P(\mathbf{x}_k | \mathbf{x}_{k-1}^{(i)}, \mathbf{u}_k), \end{aligned}$$

and C is a normalizing constant. The importance weight according to (14) is $w_k^{(i)} = w_{k-1}^{(i)} C$. The advantage of FastSLAM 2.0 is that its proposal distribution is locally optimal [15]. That is, for each particle, it gives the smallest possible

variance in importance weight $w_k^{(i)}$ conditioned upon the available information, $\mathbf{X}_{0:k-1}^{(i)}$, $\mathbf{Z}_{0:k}$ and $\mathbf{U}_{0:k}$.

Statistically, FastSLAM (1.0 and 2.0) suffers degeneration due to its inability to forget the past. Marginalizing the map in (15) introduces dependence on the pose and measurement history, and so, when resampling depletes this history, statistical accuracy is lost [2]. Nevertheless, empirical results of FastSLAM 2.0 in real outdoor environments [33] show that the algorithm is capable of generating an accurate map in practice.

Implementation of SLAM

Practical realizations of probabilistic SLAM have become increasingly impressive in recent years, covering larger areas in more challenging environments. Here we discuss two representative implementations and mention other notable applications.

The “explore and return” experiment by Newman et al. [37] was a moderate-scale indoor implementation that validated the nondivergence properties of EKF-SLAM by returning to a precisely marked starting point. The experiment is remarkable because its return trip was fully autonomous. The robot was manually driven during the exploration phase, although without visual contact by the operator, who relied solely on a real-time rendering of the robot’s map (see Figure 6). For the return trip, the robot plans a path and returns without human intervention.

Guivant and Nebot [21] pioneered the application of SLAM in very large outdoor environments (see Figure 7). They addressed computational issues of real-time operation, while also dealing with high-speed vehicle motion, non-flat terrain, and dynamic clutter. Their results are particularly interesting because they are accompanied by accurate RTK-GPS ground truth, showing the practical veracity of the algorithm, which involved closing several large loops.

The logged data from their Victoria Park trials is available online and has become a popular benchmark within the SLAM research community.

SLAM applications now exist in a wide variety of domains. They include indoor [3], [4], [7], [12], outdoor [19], [21], aerial [25], and subsea [18], [36], [45]. There are different sensing modalities such as bearing only [13] and range only [30].

We also make honorable mention of consistent pose estimation (CPE) [22], [26], which is an entirely different SLAM paradigm based on topological mapping and data alignment, due to its exemplary results in large indoor environments.

Various researchers in the SLAM community have written software demonstrating SLAM, implemented in MATLAB, C++, and Java and available online (see Table 1). Collections of logged data are listed in Table 2. These datasets are from real sensors in real environments and are a valuable resource to assess and benchmark the various SLAM algorithms.

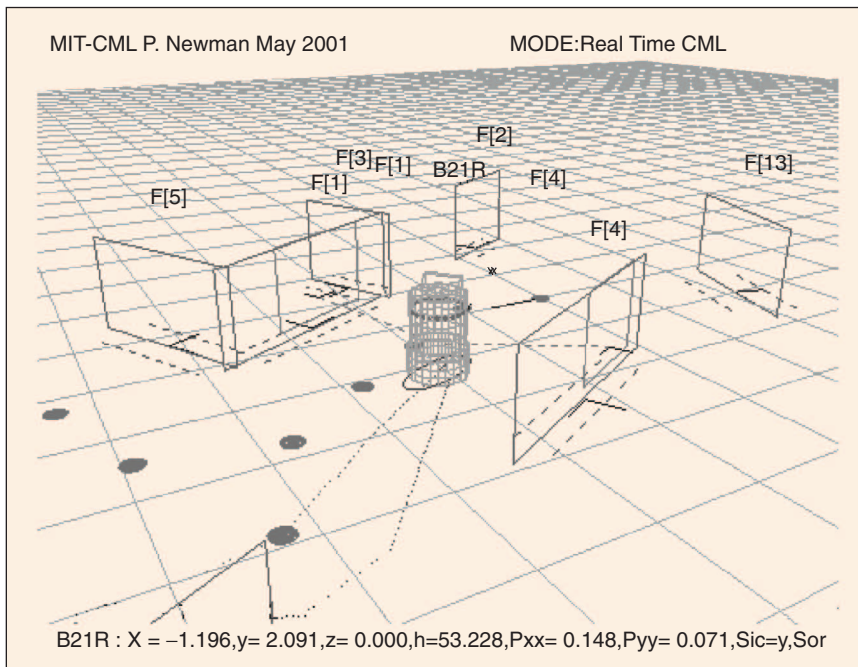


Figure 6. A real-time SLAM visualization by Newman et al. [37].

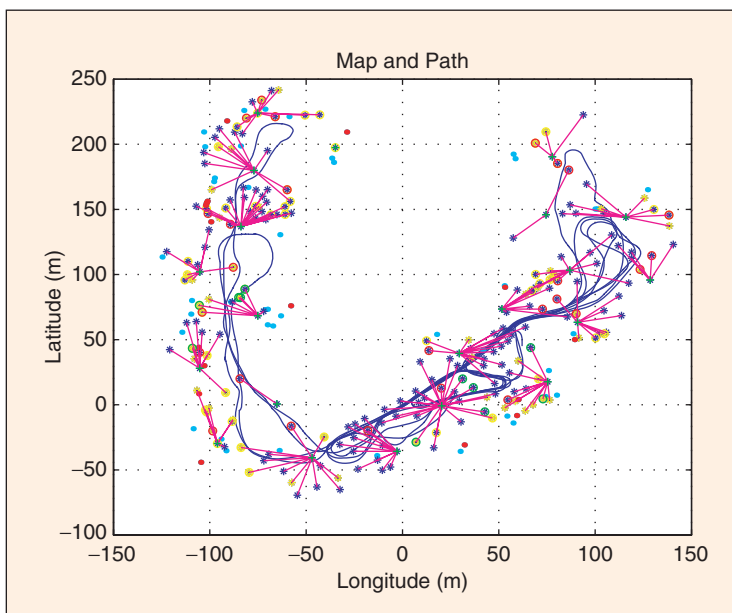


Figure 7. Large-scale outdoor SLAM by Guivant and Nebot [21].

Table 1. Open-source SLAM software.

Author	Description	Link
Kai Arras	The <i>CAS Robot Navigation Toolbox</i> , a MATLAB simulation toolbox for robot localization and mapping	http://www.cas.kth.se/toolbox/index.html
Tim Bailey	MATLAB simulators for EKF-SLAM, UKF-SLAM, and FastSLAM 1.0 and 2.0.	http://www.acfr.usyd.edu.au/homepages/academic/tbailey/software/index.html
Mark Paskin	Java library with several SLAM variants, including Kalman filter, information filter, and thin junction tree forms. Includes a MATLAB interface.	http://www.stanford.edu/~paskin/slam/
Andrew Davison	<i>Scene</i> , a C++ library for map-building and localization. Facilitates real-time single camera SLAM.	http://www.doc.ic.ac.uk/~ajd/Scene/index.html
José Neira	MATLAB EKF-SLAM simulator that demonstrates joint compatibility branch-and-bound data association.	http://webdiis.unizar.es/~neira/software/slam/slamsim.htm
Dirk Hähnel	C language grid-based version of FastSLAM.	http://www.informatik.uni-freiburg.de/~haehnel/old/download.html
Various	MATLAB code from the 2002 SLAM summer school.	http://www.cas.kth.se/slam/toc.html

Table 2. Online datasets.

Author	Description	Link
Eduardo Nebot	Numerous large-scale outdoor datasets, notably the popular Victoria Park data.	http://www.acfr.usyd.edu.au/homepages/academic/enebot/dataset.htm
Chieh-Chih Wang	Three large-scale outdoor datasets collected by the Navlab11 testbed.	http://www.cs.cmu.edu/~bobwang/datasets.html
Radish (The Robotics Data Set Repository)	Many and varied indoor datasets, including large-area data from the CSU Stanislaus Library, the Intel Research Lab in Seattle, the Edmonton Convention Centre, and more.	http://radish.sourceforge.net/
IJRR (The International Journal of Robotics Research)	IJRR maintains a Web page for each article, often containing data and video of results. A good paper example is by Bosse et al. [3], which has data from Killian Court at MIT.	http://www.ijrr.org/contents/23_12/abstract/1113.html

Conclusions

This article has described the SLAM problem and the essential methods for solving the SLAM problem and has summarized key implementations and demonstrations of the method. While there are still many practical issues to overcome, especially in more complex outdoor environments, the general SLAM method is now a well understood and established part of robotics. Part II of this tutorial will summarize more recent work in addressing some of the remaining issues in SLAM, including computation, feature representation, and data association.

References

- [1] N. Ayache and O. Faugeras, "Building, registering, and fusing noisy visual maps," *Int. J. Robot. Res.*, vol. 7, no. 6, pp. 45–65, 1988.
- [2] T. Bailey, J. Nieto, and E. Nebot, "Consistency of the FastSLAM algorithm," in *Proc. IEEE Int. Conf. Robot. Automat.*, 2006.
- [3] M. Bosse, P. Newman, J. Leonard, and S. Teller, "Simultaneous localization and map building in large-scale cyclic environments using the Atlas framework," *Int. J. Robot. Res.*, vol. 23, no. 12, pp. 1113–1140, 2004.
- [4] J.A. Castellanos, J.M. Martnez, J. Neira, and J.D. Tardós, "Experiments in multisensor mobile robot localization and map building," in *Proc. 3rd IFAC Sym. Intell. Auton. Vehicules*, 1998, pp. 173–178.
- [5] J.A. Castellanos, J.D. Tardós, and G. Schmidt, "Building a global map of the environment of a mobile robot: The importance of correlations," in *Proc. IEEE Int. Conf. Robot. Automat.*, 1997, pp. 1053–1059.
- [6] R. Chatila and J.P. Laumond, "Position referencing and consistent world modeling for mobile robots," in *Proc. IEEE Int. Conf. Robot. Automat.*, 1985, pp. 138–143.
- [7] K.S. Chong and L. Kleeman, "Feature-based mapping in real, large scale environments using an ultrasonic array," *Int. J. Robot. Res.*, vol. 18, no. 1, pp. 3–19, 1999.
- [8] D. Crisan and A. Doucet, "A survey of convergence results on particle filtering methods for practitioners," *IEEE Trans. Signal Processing*, vol. 50, no. 3, pp. 736–746, 2002.
- [9] J. Crowley, "World modeling and position estimation for a mobile robot using ultra-sonic ranging," in *Proc. IEEE Int. Conf. Robot. Automat.*, 1989, pp. 674–681.
- [10] M. Csorba, "Simultaneous Localisation and Map Building," Ph.D. dissertation, Univ. Oxford, 1997.
- [11] M. Csorba and H.F. Durrant-Whyte, "A new approach to simultaneous localisation and map building," in *Proc. SPIE Aerosense*, Orlando, FL, 1996.
- [12] A.J. Davison, Y.G. Cid, and N. Kita, "Real-time 3D SLAM with wide-angle vision," in *Proc. IFAC/EURON Symp. Intell. Auton. Vehicules*, 2004.
- [13] M. Deans and M. Hebert, "Experimental comparison of techniques for localization and mapping using a bearing-only sensor," in *Proc. Int. Symp. Experimental Robot.*, 2000, pp. 395–404.

- [14] G. Dissanayake, P. Newman, H.F. Durrant-Whyte, S. Clark, and M. Csobra, "A solution to the simultaneous localisation and mapping (SLAM) problem," *IEEE Trans. Robot. Automat.*, vol. 17, no. 3, pp. 229–241, 2001.
- [15] A. Doucet, "On sequential simulation-based methods for Bayesian filtering," Dept. Eng., Cambridge Univ., Tech. Rep., 1998.
- [16] H. Durrant-Whyte, D. Rye, and E. Nebot, "Localisation of automatic guided vehicles," in *Robotics Research: The 7th International Symposium (ISRR'95)*, G. Giralt and G. Hirzinger, Eds. New York: Springer Verlag, pp. 613–625, 1996.
- [17] H.F. Durrant-Whyte, "Uncertain geometry in robotics," *IEEE Trans. Robot. Automat.*, vol. 4, no. 1, pp. 23–31, 1988.
- [18] R. Eustice, H. Singh, J. Leonard, M. Walter, and R. Ballard, "Visually navigating the RMS Titanic with SLAM information filters," in *Robotics: Science and Systems*. Cambridge, MA: MIT Press, 2005.
- [19] J. Folkesson and H.I. Christensen, "Graphical SLAM—a self-correcting map," in *Proc. IEEE Int. Conf. Robot. Automat.*, 2004, pp. 791–798.
- [20] J. Guivant, E.M. Nebot, and S. Baiker, "Localization and map building using laser range sensors in outdoor applications," *J. Robot. Syst.*, vol. 17, no. 10, pp. 565–583, 2000.
- [21] J.E. Guivant and E.M. Nebot, "Optimization of the simultaneous localization and map-building algorithm for real-time implementation," *IEEE Trans. Robot. Automat.*, vol. 17, no. 3, pp. 242–257, 2001.
- [22] J.S. Gutmann and K. Konolige, "Incremental mapping of large cyclic environments," in *Proc. IEEE Int. Symp. Comput. Intell. Robot. Automat.*, 1999, pp. 318–325.
- [23] J. Hollerbach and D. Koditschek, Eds., *Robotics Research, The Ninth International Symposium (ISRR'99)*. New York: Springer-Verlag, 2000.
- [24] S.J. Julier and J.K. Uhlmann, "A counter example to the theory of simultaneous localization and map building," in *Proc. IEEE Int. Conf. Robot. Automat.*, pp. 4238–4243, 2001.
- [25] J. Kim and S. Sukkarieh, "Airborne simultaneous localisation and map building," in *Proc. IEEE Int. Conf. Robot. Automat.*, 2003, pp. 406–411.
- [26] K. Konolige, "Large-scale map-making," in *Proc. Nat. Conf. AI (AAAI)*, 2004, pp. 457–463.
- [27] J.J. Leonard and H.F. Durrant-Whyte, "Simultaneous map building and localisation for an autonomous mobile robot," in *Proc. IEEE Int. Workshop Intell. Robots Syst. (IROS)*, Osaka, Japan, 1991, pp. 1442–1447.
- [28] J.J. Leonard and H.F. Durrant-Whyte, *Directed Sonar Navigation*. Norwell, MA: Kluwer, 1992.
- [29] J.J. Leonard and H.J.S. Feder, "A computational efficient method for large-scale concurrent mapping and localisation," in *Robotics Research, The Ninth International Symposium (ISRR'99)*, J. Hollerbach and D. Koditschek, Eds. New York: Springer-Verlag, pp. 169–176, 2000.
- [30] J.J. Leonard, R.J. Rikoski, P.M. Newman, and M.C. Bosse, "Mapping partially observable features from multiple uncertain vantage points," *Int. J. Robot. Res.*, vol. 21, no. 10–11, pp. 943–975, 2002.
- [31] P.S. Maybeck, *Stochastic Models, Estimation and Control*, vol. 1. New York: Academic 1979.
- [32] M. Montemerlo, S. Thrun, D. Koller, and B. Wegbreit, "Fast-SLAM: A factored solution to the simultaneous localization and mapping problem," in *Proc. AAAI Nat. Conf. Artif. Intell.*, 2002, pp. 593–598.
- [33] M. Montemerlo, S. Thrun, D. Koller, and B. Wegbreit, "Fast-SLAM 2.0: An improved particle filtering algorithm for simultaneous localization and mapping that provably converges," in *Proc. Int. Joint Conf. Artif. Intell.*, 2003, pp. 1151–1156.
- [34] K. Murphy, "Bayesian map learning in dynamic environments," in *Advances in Neural Information Processing Systems*. San Mateo, CA: Morgan Kaufman, 1999.
- [35] J. Neira and J.D. Tardós, "Data association in stochastic mapping using the joint compatibility test," *IEEE Trans. Robot. Automat.*, vol. 17, no. 6, pp. 890–897, 2001.
- [36] P.M. Newman and J.J. Leonard, "Pure range-only subsea SLAM," in *Proc. IEEE Int. Conf. Robot. Automat.*, 2003, pp. 1921–1926.
- [37] P.M. Newman, J.J. Leonard, J. Neira, and J. Tardós, "Explore and return: Experimental validation of real time concurrent mapping and localization," in *Proc. IEEE Int. Conf. Robot. Automat.*, 2002, pp. 1802–1809.
- [38] W.D. Renken, "Concurrent localization and map building for mobile robots using ultrasonic sensors," in *Proc. IEEE Int. Workshop Intell. Robots Syst. (IROS)*, 1993.
- [39] R. Smith and P. Cheesman, "On the representation of spatial uncertainty," *Int. J. Robot. Res.*, vol. 5, no. 4, pp. 56–68, 1987.
- [40] R. Smith, M. Self, and P. Cheeseman, "Estimating uncertain spatial relationships in robotics," in *Autonomous Robot Vehicles*, I.J. Cox and G.T. Wilfon, Eds. New York: Springer-Verlag, pp. 167–193, 1990.
- [41] S. Thrun, W. Burgard, and D. Fox, "A real-time algorithm for mobile robot mapping with applications to multi-robot and 3D mapping," in *Proc. Int. Conf. Robot. Automat.*, 2000, pp. 321–328.
- [42] S. Thrun, D. Fox, and W. Burgard, "A probabilistic approach to concurrent mapping and localization for mobile robots," *Mach. Learning*, vol. 31, no. 1, pp. 29–53, 1998.
- [43] S. Thrun, Y. Liu, D. Koller, A. Ng, and H. Durrant-Whyte, "Simultaneous localisation and mapping with sparse extended information filters," *Int. J. Robot. Res.*, vol. 23, no. 7–8, pp. 693–716, 2004.
- [44] C.C. Wang, C. Thorpe, and S. Thrun, "On-line simultaneous localization and mapping with detection and tracking of moving objects," in *Proc. IEEE Int. Conf. Robot. Automat.*, 2003, pp. 2918–2924.
- [45] S.B. Williams, P. Newman, G. Dissanayake, and H.F. Durrant-Whyte, "Autonomous underwater simultaneous localisation and map building," in *Proc. IEEE Int. Conf. Robot. Automat. (ICRA)*, San Francisco, CA, Apr. 2000, pp. 1793–1798.

Hugh Durrant-Whyte received the B.Sc. in nuclear engineering from the University of London, United Kingdom, in 1983, and the M.S.E. and Ph.D. degrees, both in systems engineering, from the University of Pennsylvania, in 1985 and 1986, respectively. From 1987–1995, he was a senior lecturer in engineering science, the University of Oxford, United Kingdom, and a fellow of Oriel College Oxford. From 1995–2002 he was a professor of mechatronic engineering at the University of Sydney. In 2002, he was awarded an Australian Research Council (ARC) Federation Fellowship. He also now leads the ARC Centre of Excellence in Autonomous Systems. His research work focuses on autonomous vehicle navigation and decentralized data fusion methods. His work in applications includes automation in cargo handling, mining, defense, and marine systems. He has published over 300 technical papers and has won numerous awards and prizes for his work.

Tim Bailey received his B.E. in mechanical and mechatronic engineering at the University of Sydney in 1997 and his Ph.D. at the Australian Centre for Field Robots, University of Sydney, in 2003. He is currently a research fellow at the ACFR. His research interests include mobile robot localization and mapping, Bayesian estimation techniques, and probabilistic sensor modeling.

Address for Correspondence: Hugh Durrant-Whyte, Australian Centre for Field Robotics (ACFR) J04, The University of Sydney, Sydney NSW 2006, Australia. E-mail: hugh@acfr.usyd.edu.au.