

# LIMITS TO THE CONSISTENCY OF EKF-BASED SLAM<sup>1</sup>

José A. Castellanos   José Neira   Juan D. Tardós

*Dept. Informática e Ingeniería de Sistemas  
Universidad de Zaragoza  
María de Luna 1, 50018 Zaragoza, Spain  
Email: {jacaste, jneira, tardos}@unizar.es*

**Abstract:** This paper analyzes the consistency of the classical extended Kalman filter (EKF) solution to the simultaneous localization and map building (SLAM) problem. Our results show that in large environments the map quickly becomes inconsistent due to linearization errors. We propose a new EKF-based SLAM algorithm, *robocentric mapping*, that greatly reduces linearization errors, improving map consistency. We also present results showing that large-scale mapping methods based on building local maps with a local uncertainty representation (Tardós *et al.*, 2002) have better consistency than methods that work with global uncertainties.

**Keywords:** SLAM, map building, extended Kalman filter, linearization errors, consistency.

## 1. INTRODUCTION

During the last five years, the robotics literature has been populated with scientific work on the problem of simultaneous localization and mapping (SLAM), also referred to as concurrent mapping and localization (CML). Basically, the problem relates to the autonomous exploration of an unknown environment using the relative observations obtained by onboard sensors mounted on a computerized vehicle. The most popular approach to SLAM dates back to the seminal work of Smith *et al.* (1988) where the idea of representing the structure of the navigation area in a discrete-time state-space framework was originally presented. They introduced the concept of *stochastic map* and they developed a rigorous solution to the SLAM problem using an EKF perspective.

Many successful implementations of this approach have been reported in indoors (Castellanos *et al.*, 1999), outdoors (Guivant and Nebot, 2001), subsea (Leonard *et al.*, 2001) and air-borne (Kim and Sukkarieh, 2003) applications.

Dissanayake *et al.* (2001) proved three important convergency properties of the EKF solution to SLAM, namely that: (1) the determinant of any submatrix of the map covariance matrix decreases monotonically as observations are successively made; (2) in the limit as the number of observations increases, the landmark estimates become fully correlated; and (3) in the limit, the covariance associated with any single landmark location estimate reaches a lower bound determined only by the initial covariance in the vehicle location estimate at the time of the first sighting of the first landmark.

However, it is important to note that these theoretical results only refer to the evolution of the covariance matrices computed by the EKF in the

---

<sup>1</sup> This research has been funded in part by the Dirección General de Investigación of Spain under projects DPI2000-1265 and DPI2003-07986.

ideal linear case. They overlook the fact that, given that SLAM is a nonlinear problem, there is no guarantee that the computed covariances will match the actual estimation errors, which is the true SLAM consistency issue. In fact, it is well known that the inherent approximations due to linearization of the system and measurement equations can lead to divergence of the EKF (Bar-Shalom *et al.*, 2001). In the case of a stationary vehicle equipped with a range-bearing sensor with nonzero angular uncertainty, Julier and Uhlmann (2001) showed that the EKF algorithm always produces an inconsistent map, with a null covariance matrix in steady-state. They also showed through simulations that the same inconsistency problem appears in SLAM with a moving vehicle, although the problem seems to arise after several hundred time steps.

Many recent efforts have concentrated on reducing the computational complexity of SLAM in large environments (Leonard and Feder, 2000; Guivant and Nebot, 2001; Tardós *et al.*, 2002; Leonard and Newman, 2003). In this paper we show that linearization errors lead to inconsistent estimates well before the computational problems arise. We present simulation results showing that methods like *local map joining* (Tardós *et al.*, 2002), based on building independent local maps, effectively reduce linearization errors. We also propose a new method called *robocentric mapping*, that can be used as a replacement for the standard EKF-SLAM algorithm for building medium-size maps, improving the map consistency.

The rest of the paper is structured as follows. Section 2 summarizes the classical EKF solution to SLAM and analyzes the importance of linearization errors in medium-size maps through a simulated example. The robocentric mapping technique is introduced in section 4. Section 5 presents simulation results that confirm the improved consistency properties of robocentric mapping and local map joining. Conclusions are drawn in section 6.

## 2. SLAM USING THE EXTENDED KALMAN FILTER

In this section we summarize the basic EKF SLAM equations to point out the steps in which linearized approximations are introduced. We subsequently analyze their influence in the consistency of the solution obtained.

In the standard EKF-based approach to SLAM, the environment information related to a set of elements  $\{R, F_1, \dots, F_n\}$  is represented by a stochastic state vector  $\mathbf{x}^W$  with estimated mean  $\hat{\mathbf{x}}^W$  and estimated error covariance  $\mathbf{P}^W$ :

$$\hat{\mathbf{x}}^W = \begin{bmatrix} \hat{\mathbf{x}}_R^W \\ \vdots \\ \hat{\mathbf{x}}_{F_n}^W \end{bmatrix}; \mathbf{P}^W = \begin{bmatrix} \mathbf{P}_R^W & \cdots & \mathbf{P}_{RF_n}^W \\ \vdots & \ddots & \vdots \\ \mathbf{P}_{F_n R}^W & \cdots & \mathbf{P}_{F_n}^W \end{bmatrix} \quad (1)$$

Vector  $\hat{\mathbf{x}}^W$  contains the estimated location of the vehicle  $R$  and the environment features  $F_1 \dots F_n$ , all with respect to a base reference  $W$ .

If a reference external to the vehicle is used as base reference, the vehicle location must be initialized with the corresponding nonzero uncertainty. It is a common misconception is that this nonzero initial level of uncertainty in the vehicle location may improve map consistency. In contrast, our results will show that this quickly results in optimistic covariance values due to linearization errors. For this reason, we use the vehicle location before the first motion (at step  $k = 0$ ) as the base reference ( $W = R_0$ ). Thus, the map can be initialized with zero covariance for the vehicle location:  $\hat{\mathbf{x}}_0^W = (0, 0, 0)^T$ ,  $\mathbf{P}_0^W = \mathbf{0}$ . Our results show that this improves the consistency of the EKF SLAM algorithm.

### 2.1 The prediction step

When the vehicle moves from position at step  $k-1$  to position at step  $k$ , its motion is estimated by odometry:

$$\mathbf{x}_{R_k}^{R_{k-1}} = \hat{\mathbf{x}}_{R_k}^{R_{k-1}} + \mathbf{v}_k \quad (2)$$

where  $\mathbf{v}_k$  (process noise) is assumed to be additive, zero-mean and white. After this motion, the robot location is predicted as follows:

$$\mathbf{x}_{R_k|k-1}^W = \mathbf{x}_{R_{k-1}}^W \oplus \mathbf{x}_{R_k}^{R_{k-1}} \quad (3)$$

where  $\oplus$  represents the composition of transformations. Note that through the composition operation  $\oplus$ , the prediction step introduces the *first* linearization of the problem: the covariance of the predicted value  $\mathbf{x}_{R_k|k-1}^W$  is obtained by linearizing the composition operation around the estimated values  $\hat{\mathbf{x}}_{R_{k-1}}^W$  and  $\hat{\mathbf{x}}_{R_k}^{R_{k-1}}$  using the appropriate Jacobians (see appendix in (Tardós *et al.*, 2002)).

### 2.2 The estimation step

At step  $k$  an onboard sensor obtains a partial measurement  $\mathbf{z}_k$  of feature  $F$ , related to the state by a nonlinear function  $\mathbf{h}_k$ :

$$\mathbf{z}_k = \mathbf{h}_k(\mathbf{x}_k^W) + \mathbf{w}_k \quad (4)$$

where  $\mathbf{w}_k$  (measurement noise) is assumed to be additive, zero-mean, white and independent of the process noise, with covariance  $\mathbf{R}_k$ .

A *second* linearization, this time around the current map prediction, yields:

$$\begin{aligned} \mathbf{z}_k &\simeq \mathbf{h}_k(\hat{\mathbf{x}}_{k|k-1}^W) + \mathbf{H}_k(\mathbf{x}_k^W - \hat{\mathbf{x}}_{k|k-1}^W) \\ \mathbf{H}_k &= \left. \frac{\partial \mathbf{h}_k}{\partial \mathbf{x}_k^W} \right|_{(\hat{\mathbf{x}}_{k|k-1}^W)} \\ &= (\mathbf{H}_{R_k} \mathbf{0} \cdots \mathbf{H}_{F_k} \cdots \mathbf{0}) \\ \mathbf{H}_{R_k} &= \left. \frac{\partial \mathbf{h}_k}{\partial \mathbf{x}_{R_k}^W} \right|_{(\hat{\mathbf{x}}_{k|k-1}^W)} ; \mathbf{H}_{F_k} = \left. \frac{\partial \mathbf{h}_k}{\partial \mathbf{x}_{F_k}^W} \right|_{(\hat{\mathbf{x}}_{k|k-1}^W)} \end{aligned} \quad (5)$$

Measurement  $\mathbf{z}_k$  is used to obtain a new estimation of the state using the standard EKF update equations:

$$\begin{aligned} \hat{\mathbf{x}}_k^W &= \hat{\mathbf{x}}_{k|k-1}^W + \mathbf{K}_k(\mathbf{z}_k - \mathbf{h}_k(\hat{\mathbf{x}}_{k|k-1}^W)) \\ \mathbf{P}_k^W &= (\mathbf{I} - \mathbf{K}_k \mathbf{H}_k) \mathbf{P}_{k|k-1}^W \\ \mathbf{K}_k &= \mathbf{P}_{k|k-1}^W \mathbf{H}_k^T (\mathbf{H}_k \mathbf{P}_{k|k-1}^W \mathbf{H}_k^T + \mathbf{R}_k)^{-1} \end{aligned} \quad (6)$$

### 3. THE INCONSISTENCY OF EKF SLAM

A state estimator is called *consistent* if its state estimation error  $\mathbf{x}_k^W - \hat{\mathbf{x}}_k^W$  satisfies (Bar-Shalom *et al.*, 2001):

$$\begin{aligned} E[\mathbf{x}_k^W - \hat{\mathbf{x}}_k^W] &= \mathbf{0} \\ E[(\mathbf{x}_k^W - \hat{\mathbf{x}}_k^W)(\mathbf{x}_k^W - \hat{\mathbf{x}}_k^W)^T] &= \mathbf{P}_k^W \end{aligned} \quad (7)$$

This means that the estimator is *unbiased* and that the actual Mean Square Error matches the filter calculated covariances. Given that SLAM is a nonlinear problem, consistency checking is of paramount importance. When the ground true solution for the state variables is available, a statistical test for filter consistency can be carried out on the Normalized Estimation Error Squared (NEES), defined as:

$$D^2 = (\mathbf{x}_k^W - \hat{\mathbf{x}}_k^W)^T (\mathbf{P}_k^W)^{-1} (\mathbf{x}_k^W - \hat{\mathbf{x}}_k^W) \quad (8)$$

Consistency is checked using a chi-squared test:

$$D^2 \leq \chi_{r,1-\alpha}^2 \quad (9)$$

where  $r = \dim(\mathbf{x}_k^W)$  and  $\alpha$  is the desired significance level (usually 0.05).

To isolate the effects of linearization errors on the consistency of the EKF-based approach to SLAM, we have designed a simulated experiment where a vehicle travels along a rectangular-shaped trajectory of  $100 \times 20$  meters, i.e. a 240-meter loop trajectory, moving 1m per step. The map of the navigation environment is composed of 2-D point features, with a feature density of 0.5 feature/m.

The vehicle is equipped with a range-bearing sensor with a maximum range of 15 meters and a 180 degrees frontal field-of-view. Gaussian-distributed synthetic errors were generated for both the sensor measurements (standard deviation of 5 cm/m in range and 0.5deg in direction) and for the odometry model of the vehicle (displacement standard deviations of 0.2 m/m in position and 0.5deg in orientation). Additionally, known data association is considered.

Figure 1, left, shows the evolution of angular error and uncertainty ( $2\sigma$  bounds) in the vehicle location along the trajectory. For this SLAM simulation, the initial vehicle location is used as base reference, allowing to set the initial vehicle uncertainty to zero. The theoretical uncertainty level was obtained by simulating the same trajectory linearizing around the ground truth solution, so that there are no linearization errors. We can see that, while the theoretical angular uncertainty increases until loop closing, the uncertainty computed by the EKF reaches a maximum level (around 0.5deg in this case). This results in the vehicle location estimation failing the consistency check of eq. (9) after only 100 steps.

It is common practice to build a map relative to a fixed base reference, different from the initial vehicle location. This normally requires to assign an initial level of uncertainty to the vehicle estimated location. As it is argued in Dissanayake *et al.* (2001), the vehicle uncertainty should always be above this initial level. Surprisingly, our simulation shows that when a non-zero initial uncertainty is used (fig. 1, right), the estimated vehicle uncertainty rapidly drops below its initial value (1deg) making the estimation inconsistent after only 50 EKF update steps. This corroborates the results of Julier and Uhlmann (2001), but also shows that the problem arises in practice earlier than they suggested.

An alternative to maintain consistency by reducing the harmful effects of linearization errors (and also to reduce the computational cost of the process) is *Local Map Joining* (Tardós *et al.*, 2002). Instead of building one global map from the beginning, a sequence of independent local maps of limited size are built. Then, each local map is joined with the global map. Matchings between global and local features can then be found and fused to update the global map. The result is equivalent to building one global map from the beginning, except precisely for linearization errors. Most of the process consists in updating local maps where errors remain small, and their effect is local. In section 5 we show that Local Map Joining allows to attain better consistency in the stochastic map, for a fraction of the computational cost of standard EKF-based SLAM.




Fig. 1. Angular error and  $2\sigma$  uncertainty bounds of the vehicle estimated location for the cases of zero (left) and nonzero (right) initial uncertainty.

#### 4. ROBOCENTRIC MAPPING

In this section we formulate the EKF-based approach to SLAM using the reference frame attached to the vehicle  $R$  as base reference of the stochastic map. Thus, the environmental information  $\{R, F_1, \dots, F_n\}$  is represented by a stochastic state vector  $\mathbf{x}^R$  with estimated mean  $\hat{\mathbf{x}}^R$  and estimated error covariance  $\mathbf{P}^R$ :

$$\hat{\mathbf{x}}^R = \begin{bmatrix} \hat{\mathbf{x}}_W^R \\ \hat{\mathbf{x}}_{F_1}^R \\ \vdots \\ \hat{\mathbf{x}}_{F_n}^R \end{bmatrix}; \mathbf{P}^R = \begin{bmatrix} \mathbf{P}_W^R & \cdots & \mathbf{P}_{WF_n}^R \\ \vdots & \ddots & \vdots \\ \mathbf{P}_{F_n W}^R & \cdots & \mathbf{P}_{F_n}^R \end{bmatrix} \quad (10)$$

where the world reference frame  $W$  has been included as a non-observable feature in the stochastic state vector.

Again we take the initial vehicle location as base reference  $W = R_0$ , and thus the map is initialized with perfect knowledge of the world location:  $\hat{\mathbf{x}}_W^R = (0, 0, 0)^T$  and  $\mathbf{P}_W^R = \mathbf{0}$ .

##### 4.1 The prediction step

After the vehicle changes its location from step  $k-1$  to step  $k$ , the complete structure of the stochastic map should be affected by the process noise  $\mathbf{Q}_k$  associated with the displacement  $\mathbf{x}_{R_k}^{R_{k-1}}$  as estimated by odometry. Thus, the estimated location of a given map feature  $F$  should be updated as:

$$\mathbf{x}_{F_{k|k-1}}^{R_k} = \ominus \mathbf{x}_{R_k}^{R_{k-1}} \oplus \mathbf{x}_{F_{k-1}}^{R_{k-1}} \quad (11)$$

and therefore, its estimated covariance would be computed from the corresponding linearization around the estimated values  $\hat{\mathbf{x}}_{R_k}^{R_{k-1}}$  and  $\hat{\mathbf{x}}_{F_{k-1}}^{R_{k-1}}$ .

To reduce the effects of this linearization on the quality of the stochastic map, we first improve the

estimated relative displacement in the estimation step of the EKF algorithm. Therefore, in the prediction step the vehicle motion is included, as an independent feature, in the previously available stochastic map  $\mathbf{x}_{k-1}^{R_{k-1}}$ :

$$\hat{\mathbf{x}}_{k|k-1}^{R_{k-1}} = \begin{bmatrix} \hat{\mathbf{x}}_{k-1}^{R_{k-1}} \\ \hat{\mathbf{x}}_{R_k}^{R_{k-1}} \end{bmatrix}; \mathbf{P}_{k|k-1}^{R_{k-1}} = \begin{bmatrix} \mathbf{P}_{k-1}^{R_{k-1}} & \mathbf{0} \\ \mathbf{0} & \mathbf{Q}_k \end{bmatrix} \quad (12)$$

##### 4.2 The estimation step

Now, linearization of the measurement equation around the estimated values of both the stochastic state vector and the partial measurement  $\mathbf{z}_k$  yields:

$$\begin{aligned} \mathbf{z}_k &\simeq \mathbf{h}_k(\hat{\mathbf{x}}_{k|k-1}^{R_{k-1}}) + \mathbf{H}_k(\mathbf{x}_k^{R_{k-1}} - \hat{\mathbf{x}}_{k|k-1}^{R_{k-1}}) \\ \mathbf{H}_k &= \left. \frac{\partial \mathbf{h}_k}{\partial \mathbf{x}_k^{R_{k-1}}} \right|_{(\hat{\mathbf{x}}_{k|k-1}^{R_{k-1}})} \\ &= (\mathbf{0} \cdots \mathbf{0} \mathbf{H}_{F_k} \mathbf{0} \cdots \mathbf{0} \mathbf{H}_{R_k}) \end{aligned} \quad (13)$$

where:

$$\mathbf{H}_{F_k} = \left. \frac{\partial \mathbf{h}_k}{\partial \mathbf{x}_{F_k}^{R_{k-1}}} \right|_{(\hat{\mathbf{x}}_{k|k-1}^{R_{k-1}})}; \mathbf{H}_{R_k} = \left. \frac{\partial \mathbf{h}_k}{\partial \mathbf{x}_{R_k}^{R_{k-1}}} \right|_{(\hat{\mathbf{x}}_{k|k-1}^{R_{k-1}})}$$

Equations which are subsequently used to obtain a new estimation of the stochastic state vector  $\hat{\mathbf{x}}_{k|k}^{R_{k-1}}$  and its covariance matrix  $\mathbf{P}_{k|k}^{R_{k-1}}$ , using the previously described EKF update equations. Note that, because the relative displacement of the vehicle from time  $k-1$  to time  $k$  was included as a feature of the stochastic state vector it is also improved during the application of the update equations.

### 4.3 The composition step

Finally, the stochastic state vector of the robocentric map is obtained by affecting each estimated location by the improved vehicle motion:

$$\hat{\mathbf{x}}_k^{R_k} = \begin{bmatrix} \ominus \hat{\mathbf{x}}_{R_k}^{R_{k-1}} \oplus \hat{\mathbf{x}}_{R_0}^{R_{k-1}} \\ \ominus \hat{\mathbf{x}}_{R_k}^{R_{k-1}} \oplus \hat{\mathbf{x}}_{F_1}^{R_{k-1}} \\ \vdots \\ \ominus \hat{\mathbf{x}}_{R_k}^{R_{k-1}} \oplus \hat{\mathbf{x}}_{F_n}^{R_{k-1}} \end{bmatrix} \quad (14)$$

with corresponding covariance matrix:

$$\mathbf{P}_k^{R_k} = [\mathbf{J}_2 \ \mathbf{J}_1] \mathbf{P}_{k|k}^{R_{k-1}} \begin{bmatrix} \mathbf{J}_2^T \\ \mathbf{J}_1^T \end{bmatrix} \quad (15)$$

with,

$$\mathbf{J}_1 = \begin{bmatrix} \mathbf{J}_{1\oplus} \{ \ominus \hat{\mathbf{x}}_{R_k}^{R_{k-1}}, \hat{\mathbf{x}}_{R_0}^{R_{k-1}} \} \mathbf{J}_{\ominus} \{ \hat{\mathbf{x}}_{R_k}^{R_{k-1}} \} \\ \vdots \\ \mathbf{J}_{1\oplus} \{ \ominus \hat{\mathbf{x}}_{R_k}^{R_{k-1}}, \hat{\mathbf{x}}_{F_n}^{R_{k-1}} \} \mathbf{J}_{\ominus} \{ \hat{\mathbf{x}}_{R_k}^{R_{k-1}} \} \end{bmatrix}$$

and,

$$\mathbf{J}_2 = \begin{bmatrix} \mathbf{J}_{2\oplus} \{ \ominus \hat{\mathbf{x}}_{R_k}^{R_{k-1}}, \hat{\mathbf{x}}_{R_0}^{R_{k-1}} \} \cdots & \mathbf{0} \\ \vdots & \vdots \\ \mathbf{0} & \cdots \mathbf{J}_{2\oplus} \{ \ominus \hat{\mathbf{x}}_{R_k}^{R_{k-1}}, \hat{\mathbf{x}}_{F_n}^{R_{k-1}} \} \end{bmatrix}$$

where  $\mathbf{J}_{1\oplus}$  and  $\mathbf{J}_{2\oplus}$  are the Jacobians of transformation composition (see appendix in (Tardós *et al.*, 2002)).

## 5. EXPERIMENTAL RESULTS

In our controlled simulation environment we have compared the performance of the map joining and the robocentric mapping algorithms. Figure 2 shows that the robocentric mapping algorithm always obtains a non-optimistic estimation for the vehicle location uncertainty and that the map joining algorithm is slightly optimistic in the second-half of the trajectory.

The consistency of the solutions is determined using equations (8) and (9). We conclude that both algorithms perform similarly, for the consistency of the estimated vehicle locations along the trajectory (figure 3) and also for the consistency of the map feature locations (figure 4). Finally, figure 5 shows the estimated errors and their respective  $2\sigma$  bounds for both approaches. The solutions are mostly consistent with ground truth. Only some discrepancies are observed for the trajectory points furthest from the initial location of the vehicle.




Fig. 2. Angular uncertainty of the vehicle estimated location for the robocentric and the map joining approaches.




Fig. 3. Consistency of the estimated vehicle location along its trajectory.




Fig. 4. Consistency of the estimated map feature locations.

## 6. CONCLUSIONS

In this work we have shown that in the standard extended Kalman Filter for SLAM, linearization errors produce inconsistency problems that show up long before computational problems arise. We




Fig. 5. Comparison of position and orientation errors along the vehicle trajectory for the approaches discussed in the text: Robocentric (left) and Map Joining (right).

follow a precise definition of filter consistency that considers both the accuracy of the estimation and of its covariances. We show that methods based on local maps, such as local map joining, are effective in reducing linearization errors and thus attain better consistency. We also propose *robocentric mapping* as an alternative for the standard EKF-SLAM algorithm to improve map consistency. The combination of local map joining and robocentric mapping is likely to further improve map consistency and is the subject of immediate future work.

These techniques allow to apply the EKF-based solution to SLAM at a larger scale. It is however likely that nonlinearity problems will arise again as larger environments are tackled. We feel that to overcome these limitations it is important to investigate the use of alternative formulations to SLAM, nonlinear and non Gaussian methods. Making these methods computationally efficient to be used in real time is the next important challenge in SLAM.

## REFERENCES

Bar-Shalom, Y., X. Rong Li and T. Kirubarajan (2001). *Estimation with Applications to Tracking and Navigation*. Wiley InterScience.

Castellanos, J. A., J. M. M. Montiel, J. Neira and J.D. Tardós (1999). The SPmap: A probabilistic framework for simultaneous localization and map building. *IEEE Trans. on Robotics and Automation* **15**(5), 948–952.

Dissanayake, M. W. M. G., P. Newman, S. Clark, H. F. Durrant-Whyte and M. Csorba (2001). A solution to the simultaneous localization and map building (SLAM) problem. *IEEE Trans. on Robotics and Automation* **17**(3), 229–241.

Guivant, J. E. and E. M. Nebot (2001). Optimization of the simultaneous localization and map-building algorithm for real-time implementation. *IEEE Trans. on Robotics and Automation* **17**(3), 242–257.

Julier, S. J. and J. K. Uhlmann (2001). A counter example to the theory of simultaneous localization and map building. In: *Proc. 2001 IEEE Int. Conf. on Robotics and Automation*. Seoul, Korea. pp. 4238–4243.

Kim, J.H. and S. Sukkarieh (2003). Airborne simultaneous localisation and map building. In: *IEEE Int. Conf. on Robotics and Automation*. Taipei, Taiwan.

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

Leonard, J. J., R. Carpenter and H. J. S. Feder (2001). Stochastic mapping using forward look sonar. *Robotica* **19**, 467–480.

Leonard, J.J. and P.M. Newman (2003). Consistent, convergent and constant-time slam. In: *Int. Joint Conf. on Artificial Intelligence*. Acapulco, Mexico.

Smith, R., M. Self and P. Cheeseman (1988). A stochastic map for uncertain spatial relationships. In: *Robotics Research, The Fourth Int. Symposium* (O. Faugeras and G. Giralt, Eds.). pp. 467–474. The MIT Press.

Tardós, Juan D., José Neira, Paul M. Newman and John J. Leonard (2002). Robust mapping and localization in indoor environments using sonar data. *Int. J. Robotics Research* **21**(4), 311–330.