2007 IEEE International Conference on Robotics and Automation

Roma, Italy

10–14 April 2007

FrC5 Aula 3

Sensor Modelling and Uncertainty (Regular Sessions)

2007 IEEE International Conference on
Robotics and Automation
Roma, Italy, 10-14 April 2007

A Consensus-based Approach for Estimating the Observation Likelihood of Accurate Range Sensors

Jose-Luis Blanco, Javier Gonzalez, Juan-Antonio Fernández-Madrigal
Dept. of System Engineering and Automation
University of Malaga
Málaga, Spain
{jlblanco, jgonzalez,jafma}@ctima.uma.es

***Abstract* – One of the main elements of probabilistic localization and SLAM is the probabilistic sensor model (also known as the observation likelihood function). However, when dealing with very accurate sensors like laser range scanners, most approaches artificially inflate the uncertainty in the range measurements and assume conditional independence between the individual ranges of the scan to compute this likelihood function. In this paper we propose an alternative method where each sample in the scan can contribute an accurate estimation according to both its real uncertainty and its compatible correspondences with a given map. These likelihood values of individual measurements are fused via a Linear Opinion Pool (LOP), a method from Consensus Theory. Our approach results in a more precise likelihood function than others and excels in robustness in dynamic environments. To validate our research we provide systematic comparisons with other proposals in the context of localization with particle filters.**

Index Terms – Probabilistic robotics, scan matching, global localization, particle filters, Bayesian filtering.

I. INTRODUCTION

The problems of global localization and Simultaneous Localization and Mapping (SLAM) share the same probabilistic model, where unobserved variables are estimated from noisy sensor measurements ([8],[12]) through Bayesian filtering. Basically, the distribution of unknown variables is estimated through the Bayes rule, which states how to update a prior belief of the state of the system $p(x)$ if we are given a new observation z:

$$\underbrace{p(x \mid z)}_{\text{Posterior}} \propto \underbrace{p(x)}_{\text{Prior}} \underbrace{p(z \mid x)}_{\substack{\text{Observation} \\ \text{likelihood}}} \qquad (1)$$

A fundamental component of Bayesian filtering is the observation likelihood, stated by $p(z|x)$ in (1), since it contributes new information to the estimation. Unfortunately, the likelihood of an observation given a robot pose, namely the sensor model $p(z|x)$, can not be computed exactly. Strictly speaking, measurements depend on the sensor pose into the environment, but also on the environment itself. Formally, we could define the sensor model as $p(z|x,m^*)$, where m^* represents the real environment. Unfortunately, the ground truth m^* is unknown, thus the closest that we can be to this model is to consider $p(z|x,m)$ as the sensor model, where we are given an estimation of the map, $p(m)$. This approximation is assumed in all works on localization and SLAM, and it is unavoidable. Its effects can be ignored for non-accurate sensors (i.e. sonars), but they become a substantial problem for accurate ones: small discrepancies between the map and the real world lead to negligible and useless likelihood values. Previous works avoid this problem by artificially inflating the uncertainty in measurements to account for the uncertainty in the map (up to two orders of magnitude above the actual sensor uncertainty).

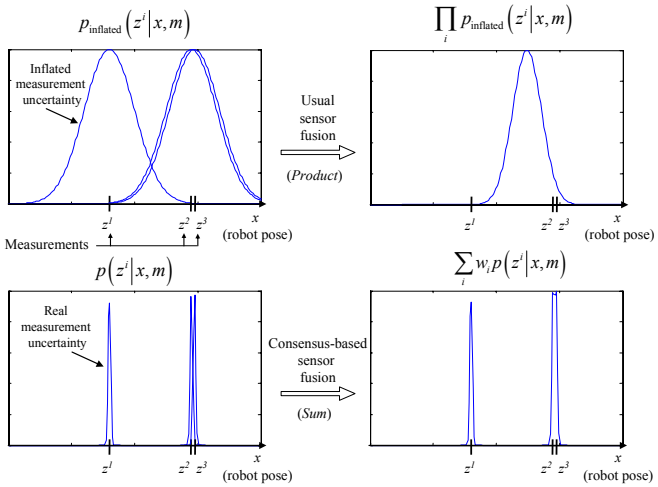

Fig. 1 The likelihood of individual range measurements z_i may be contradictory in dynamic or partially known environments. The usual method to fuse them is to inflate artificially the measurement uncertainty, to assume independence, and then compute their product (graphs on the top). The result may be peaked on robot poses actually inconsistent with observations. In this work we employ a method from Consensus Theory to fuse the high-precision likelihood functions of individual measurements (bottom graphs), which leads to an accurate and robust localization.

The common approach to integrate the likelihood values of individual range measurements of a sensor is assuming conditional independence between them, that is, computing the product of the likelihoods of each range (as illustrated in the upper graphs of Fig. 1). The problem with this technique becomes patent in dynamic environments, where there are significant differences between the expected and the actual measurements. In the example of Fig. 1 this is schematically represented by two close measures and a discrepant one on the left. The effect in the fused likelihood following the usual independence assumption is usually a high likelihood at robot poses that are actually inconsistent with all the measurements (refer to the upper graphs of Fig. 1).

The purpose of this work is to provide a more appropriate approximation of the likelihood function for highly accurate sensors, concretely for laser range-finders. We propose to consider individual likelihood values as

"*opinions*" about the final fused likelihood, which can then be calculated by means of fusion methods from Consensus Theory [9]. In this work we address this fusion via a Linear Opinion Pool (LOP), the most simple and intuitive method from consensus fusion techniques [1]. The resulting approximation of the likelihood function, that we name here *Range Scan Likelihood Consensus* (RSLC), allows considering the actual (very low) uncertainty of the sensor instead of some inflated, and thus distorted, version. We claim that our approximation leads to more accurate and dependable pose estimations than existing methods. The advantage of using *average* combination rules in the presence of outliers is well known in the field of robust sensor fusion ([3],[14]). To the best of our knowledge this work integrates for the first time these ideas into probabilistic robotics.

For simplicity, in this paper we focus on the problem of mobile robot localization only, although our approach can also be directly applied to SLAM without modifications. Since we formulate the problem in terms of raw range scans, the methods and ideas presented here are related to probabilistic scan matching techniques.

The rest of this paper is outlined as follows. In the next section we review related works in literature. Next, we set out the localization problem as a Bayesian filtering process, and in section IV we expose our approximation to computing the likelihood of observations. Finally, we present experimental results that show the performance of the RSLC method on a real robot and compare it to previous techniques.

II. Related Works

Providing an observation likelihood function for accurate range scan sensors has been a challenging issue for all probabilistic approaches to localization and map building. The most physically plausible likelihood function is the *beam model* (BM) ([21],[13]), where each range in the scan is assumed to be corrupted with zero-mean, independent identically distributed (iid) Gaussian noise. This assumption allows the following factorization:

$$p(z \mid x,m) = \prod_i p\left(z^i \mid x,m\right) \qquad (2)$$

where z represents the whole scan, z^i represents individual ranges, m is the estimated map, and the expected value of each range (the mean of the corresponding Gaussian distribution) is computed by performing ray-tracing in the grid-map. The BM has important drawbacks in practice [21]. Firstly, the resulting distribution is extremely peaked for accurate sensors, which indicates an extremely small uncertainty, but any tiny error in the map with respect to the real world makes the distribution to diverge from the ground truth. Also, and as a consequence of the previous drawback, if just one measurement out of the whole scan were affected by dynamic obstacles (those non-modelled in the map) the joint distribution would become practically zero. The following solutions have been proposed in the literature to overcome the problems with the BM: (i) to inflate artificially the uncertainty in the range measurements [22], and (ii) to preprocess ranges in order to remove those clearly caused by dynamic obstacles [6].

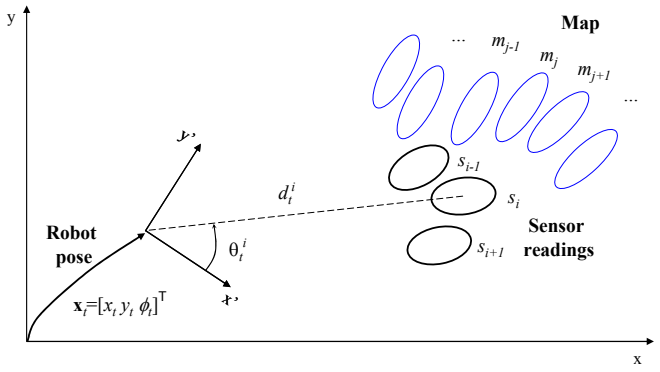

Fig. 2 A schematic representation of the variables involved in our problem. Both the map and the observations are given by a set of normally distributed 2D points m_j. The robot pose $\mathbf{x}_t$ is used to project the sensor readings into the fixed reference system <*x*,*y*>.

An alternative to the BM is the *likelihood field* (LF), an efficient approximation that avoids the costly ray-tracing operation by taking into account the 2D coordinates of the sensed points, and assigning likelihood values according to their nearness to correspondences in the map [20]. This model also inflates the sensor measurement uncertainty, typically up to values around one meter. In spite of its lack of a physical foundation, it has been successfully applied to localization and mapping ([10],[12]).

In both techniques, BM and LF, it is a common practice to use a small fraction of the ranges available in scans, achieving a speed-up in computation times and making methods more resistant to unmodelled obstacles at the cost of suboptimal solutions. Please refer to [21] for a more detailed discussion about these methods.

In the context of SLAM with Rao-Blackwellized Particle Filters ([4],[12]), an interesting alternative is proposed in [10]. There a Gaussian approximation of the likelihood function is computed by evaluating a matching function at random robot poses around a local maximum obtained from deterministic scan matching. Its dependence on the scan matching makes it prone to local minima problems (e.g. in long corridors without salient parts).

Our work contributes to the vast research field of scan matching (SM), since we manage both the observation and the map as "points" (more precisely, points distributed according to two-dimensional Gaussians). This contrasts with the most common employment of occupancy grids [7] in probabilistic localization and mapping [10]. Most of the best-known scan matching techniques, like the Iterative Closest Point (ICP) [2] or the IDC [16], aim to find the pose that achieves the optimal matching between scans. In general, these methods do no provide a measure of the uncertainty in the estimation, and hence they are not directly applicable to probabilistic SLAM. There are some exceptions, like the method proposed in [15], which considers the sensor measurement errors and the residuals from the least square error optimization. However, it does not take into account the uncertainty in the correspondence between points, which typically dominates the overall

uncertainty. The uncertainty in the correspondences is considered in the probabilistic Iterative Correspondence (pIC) method [17], but under the restrictive assumption of a normally-distributed prior of the robot pose. Other scan matching methods [13] do provide an estimation of the pose uncertainty, but they rely on the traditional assumption of independence and product-based fusion (refer to (2)), thus they are not as suited to dynamic environments as ours. Unlike iterative methods ([2],[16],[17]), our approach does neither require distance thresholds nor is iterative, since all the uncertainty in the pose is already represented by an arbitrarily distributed prior density.

III. Problem Statement

The problem of mobile robot localization, including the "robot awakening" case [5], consists of estimating the robot pose $\mathbf{x}_t$ from all the observations $z_{1:t}=\{z_1,\ldots z_t\}$ and actions $u_{1:t}=\{u_1,\ldots u_t\}$ up to the current instant t, given a known map distribution $p(m)$. The distribution to estimate is then:

$$p(\mathbf{x}_t \mid z_{1:t}, u_{1:t}, m) \tag{3}$$

We can apply the Bayes rule on the most recent observation z_t to obtain the known relation:

$$p(\mathbf{x}_t \mid z_{1:t}, u_{1:t}, m) \propto \underbrace{p(z_t \mid \mathbf{x}_t, m)}_{\text{Observation likelihood}} \underbrace{\int \overbrace{p(\mathbf{x}_t \mid \mathbf{x}_{t-1}, u_t)}^{\text{Motion model}} p(\mathbf{x}_{t-1} \mid z_{1:t-1}, u_{1:t-1}) d\mathbf{x}_{t-1}}_{\text{Prior}}$$

This expression indicates how to iteratively filter the robot pose distribution by means of the likelihood function of successive observations. In the next section we provide our proposal for the computation of this quantity for laser range scanners. But previously, we need to describe the notation and the meaning of the involved variables, which are graphically represented in Fig. 2 for clarity.

We assume a planar robot pose, represented by $\mathbf{x}_t=[x_t\ y_t\ \phi_t]^T$ for the time step t. Let the map m be a set of M points m_j, whose location uncertainty is assumed to be given by the Gaussian distributions:

$$m = \{m_j\}_{j=1..M} \qquad m_j \sim N(\boldsymbol{\mu}_m^j, \boldsymbol{\Sigma}_m^j)$$

Regarding the observations z_t, which represent points from a laser scan, they are described in robot-centric polar coordinates (d_t^i , θ_t^i):

$$z_t = \{z_t^i\}_{i=1..L} \quad z_t^i = \{d_t^i, \theta_t^i\}$$

Let s_t^i be the Cartesian coordinates of the sensed point z_t^i in the global reference system <x,y>, once transformed from the mobile system <x',y'> through:

$$s_t^i = f(\mathbf{x}_t, z_t^i) = \begin{pmatrix} f_x(\mathbf{x}_t, d_t^i, \theta_t^i) \\ f_y(\mathbf{x}_t, d_t^i, \theta_t^i) \end{pmatrix} = \begin{pmatrix} x_t + d_t^i \cos(\phi_t + \theta_t^i) \\ y_t + d_t^i \sin(\phi_t + \theta_t^i) \end{pmatrix} \tag{4}$$

Following a first-order approximation of the uncertainty of points s_t^i with normal distributions, that is, $s_t^i \sim N(\boldsymbol{\mu}_s^{t,i}, \boldsymbol{\Sigma}_s^{t,i})$, the mean and the covariance matrix for each point is obtained by propagating the sensor uncertainty through the linearization of the function $f(\cdot)$ in (4). Assuming that both errors in angle and in range are independent and normally distributed with standard deviations σ_θ and σ_d respectively, we obtain the following parameters for the distribution of s_t^i:

$$\boldsymbol{\mu}_s^{t,i} = f(\mathbf{x}_t, z_t^i) \qquad \boldsymbol{\Sigma}_s^{t,i} = \mathbf{J}_f^T\big|_{z=z_t^i} \begin{pmatrix} \sigma_\theta^2 & 0 \\ 0 & \sigma_d^2 \end{pmatrix} \mathbf{J}_f\big|_{z=z_t^i} \tag{5}$$

where **J** stands for the Jacobian of the function in (4).

IV. The Range Scan Likelihood Consensus (RSLC)

We define the RSLC as a consensus theoretic method for fusing the likelihood values of individual ranges of a scan. Consensus techniques have been employed in a variety of problems where probability values have to be fused, e.g. for combining different classification results. In this work we address data fusion by means of a particular consensus method, the Linear Opinion Pool (LOP) [1].

Let p be a probability density to be estimated from a set of L opinions $\{p_i\}$. Then, the general form for a LOP can be written as:

$$p = \sum_{i=1}^{L} w_i p_i$$

where w_i are weight factors for the individual opinions. If each p_i is a density function, we can assure that the result is also a density by imposing the condition:

$$\sum_{i=1}^{L} w_i = 1$$

For the problem we address in this paper, p is the likelihood of a whole range scan whereas $\{p_i\}$ is the set of likelihood values for individual ranges in the scan. Since we can not know in advance whether some likelihood values are more confident than others, we will simply assign an equal confidence factor to each one (a pessimistic assumption), leading to:

$$p(z_t \mid \mathbf{x}_t, m) \propto \sum_{i=1}^{L} \underbrace{p(z_t^i \mid \mathbf{x}_t, m)}_{\text{Individual likelihood values}}$$

which is a solution consistent with previous research on robust classification, where it is shown that this *average rule* for combination outperforms the classical *product rule* [14] (in the context of classifiers combination).

It remains to be described how to evaluate the individual likelihood values. As previously discussed, the problem of the BM method [21] is that inaccuracies in the map lead to non-smooth values of the likelihood function, with drastic variations for small displacements in the robot pose variable $\mathbf{x}_t$. As an alternative, we propose a novel approximation for this function that, like the LF method [20], avoids the costly ray-tracing operation by considering only the Cartesian coordinates of sensed points. Concretely, we propose to approximate the likelihood of a given range measurement z_t^i with the probability that the scanned point does correspond to any given map point. Put formally:

$$p(z_t^i \mid \mathbf{x}_t, m) \propto \sum_{j=1}^{M} \mathrm{P}(c_{ij} \mid \mathbf{x}_t, m_j) \tag{6}$$

where c_{ij} represents the correspondence between the map point m_j and the sensed point s_i, derived from z_t^i through (4)-(5). In the computation of the correspondence probabilities we also account for the possibility of a s_i not corresponding with any particular point of the map, which is represented by $c_{i\varnothing}$. To compute (6) we use:

$$\mathrm{P}\left(c_{ij}\middle|\mathbf{x}_t,m\right)=\eta_i C_{ij} \tag{7}$$

Here C_{ij} is the probability density of the pair of points s_i and m_j to coincide, normalized to the range [0,1]:

$$C_{ij}\propto\int \underbrace{p\left(s_t^i\middle|\mathbf{x}_t,z_t^i\right)}_{\sim N\left(\boldsymbol{\mu}_{st}^i,\boldsymbol{\Sigma}_{st}^i\right)}\underbrace{p\left(m_j\right)}_{\sim N\left(\boldsymbol{\mu}_m^j,\boldsymbol{\Sigma}_m^j\right)}ds_t^i$$

which has a closed form solution for Gaussian distributions [17]. The constants η_i in (7) are computed to satisfy the law of total probability:

$$\sum_{j=1}^{M}\mathrm{P}\left(c_{ij}\middle|\mathbf{x}_t,m\right)+\mathrm{P}\left(c_{i\varnothing}\middle|\mathbf{x}_t,m\right)=1$$

and the probability of no correspondence $c_{i\varnothing}$ is given by:

$$\mathrm{P}\left(c_{i\varnothing}\middle|\mathbf{x}_t,m\right)=\prod_{j=1}^{M}\left(1-C_{ij}\right)$$

To gain an insight into these expressions, consider the example in Fig. 3(a), where the probability of correspondence c_{ij} of a sensed point s_i is computed for a map with four points. Provided that the ellipses represent 95% confidence intervals, it is apparent that the sensed point s_i is probably a new point (it does not correspond to any point in the map). This fact is clearly reflected in Fig. 3(b), where this alternative receives the highest probability. Finally, according to (6), our method would assign a likelihood of $p\left(z^i\middle|\mathbf{x}_t,m\right)=0.2224$ to the sensed point of this example.

V. Experimental Results

In the following we provide systematic comparisons between the proposed method and other two well-known approximations: the BM [12] and the LF [20]. We have chosen robot localization with particles filter [5] as the framework for the tests. We also suggest the reader to view the videos available online in [18].

A. Synthetic Experiments

In the first part of the synthetic experiments, the robot pose has been estimated along a given trajectory in the environment, shown in Fig. 4(a), where the same reference map has been used both for generating 361 sensor readings within the 360º field-of-view and for computing the likelihood values. By doing so, we are reproducing the situation of a perfectly known static map. Under these conditions, we contrast the performance of particle filter-based localization for three likelihood methods: BM, LF, and our RSLC. The accuracy of the estimated pose is evaluated in two different ways: (i) by computing the mean error between the ground truth and the mean robot pose according to the particles, and (ii) by evaluating the robot pose distribution at the ground truth (since the pose

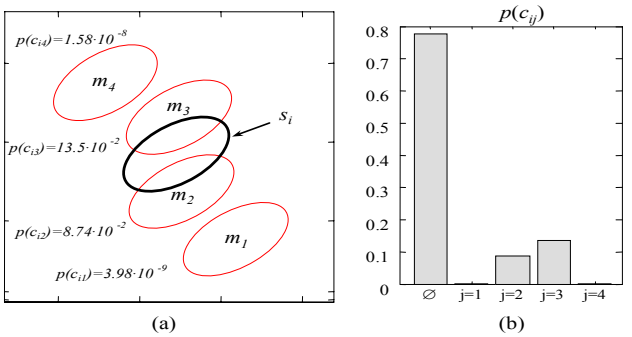

Fig. 3 An example of how our method computes the probability of the correspondence c_{ij} between a sensed point s_i and map points m_j. A sensed point and four map points are represented in (a) among with the corresponding probability values, which are graphically represented in (b), along with the probability of the point not corresponding to any map point (the symbol ∅).

distribution is given by a set of particles, a continuous version is reconstructed by a Parzen window with a Gaussian kernel [19]). The results are shown in Fig. 4(b)-(c), respectively. Due to the stochastic nature of the experiments we represent the mean values and 1σ confidence intervals for each chart after executing each experiment 10 times. The ordinates of the graphs stand for the ratio of range values employed from the whole scan: the experiments have been repeated for ratios starting at 2% (7 ranges) and up to the whole scan (361 ranges). These results reveal that RSLC provides the most accurate pose estimation (1cm mean error, approximately), even using only 2% of the ranges in the scan, while the LF method requires almost 100% of them to achieve the same accuracy. Regarding the probability assigned to the ground truth (see Fig. 4(c)), the RSLC method is surpassed by the LF and BM: RSLC is too pessimistic in assigning likelihood values. Thus, for perfectly known environments, the estimations from BM and LF are less uncertain than the one from RSLC (i.e. RSLC is excessively pessimistic for this ideal situation).

To emulate a dynamic scenario, sensor readings are simulated through the modified map shown in Fig. 4(d), whereas the robot uses the "reference" map in Fig. 4(a) to compute likelihood values. In the "dynamic" map some obstacles have been moved, removed or added to simulate typical problems. These experiments reveal a superior performance of RSLC: the mean error for our method, in Fig. 4(e), is the lowest from the three methods over the whole range of ratio values. Even more significant are the results for the probability assigned to the ground truth: by comparing Fig. 4(c) and Fig. 4(f) it can be seen how this indicator decreases slightly for the RSLC, whereas it abruptly falls for the other methods. Please, notice that particle filters perform a resampling process that discard particles at poses with a low likelihood value, thus RSLC is the most robust method in the sense that *it minimizes the probability of a correct particle to be removed*. The robustness of RSLC can be better visualized in Fig. 4(g), where the estimated path is shown according to each method. For ease of representation, the 99.7% confidence intervals are used instead of the original particles. It is clear

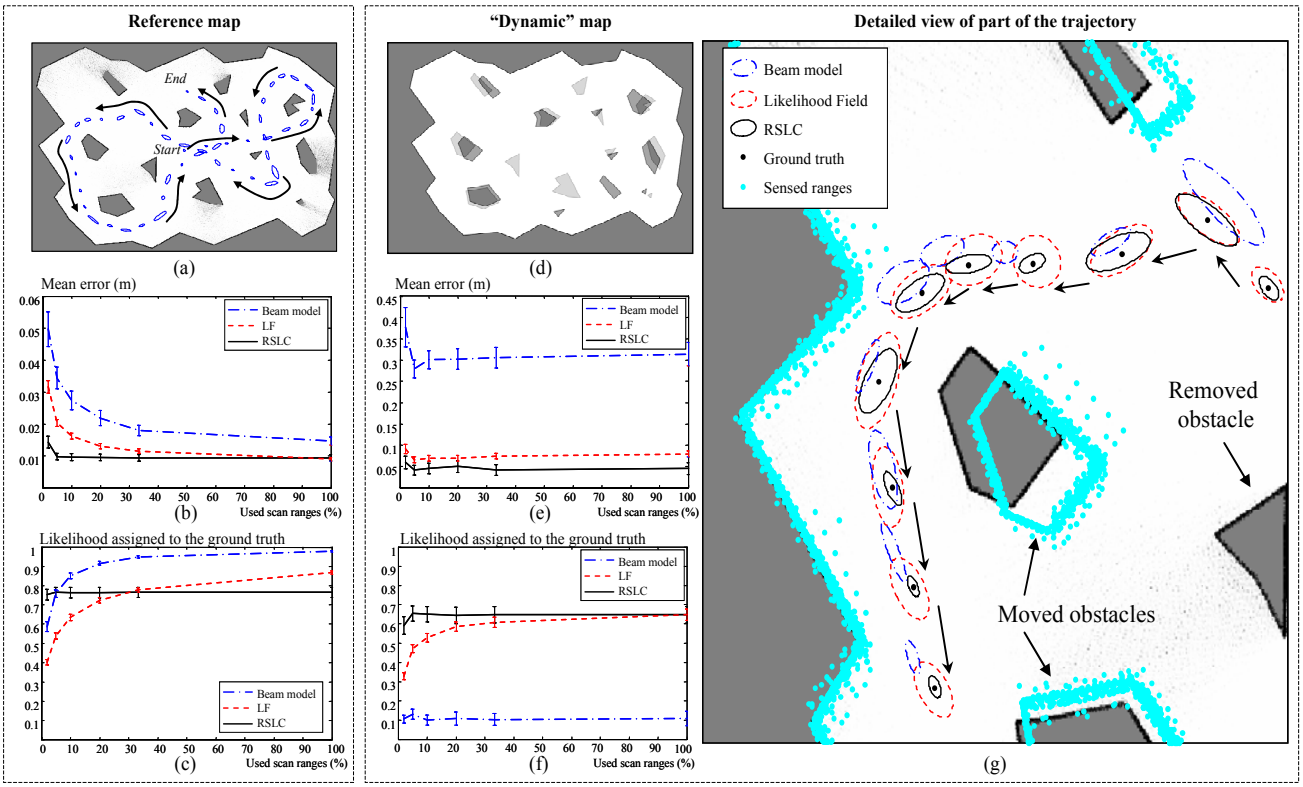

Fig. 4 Results for experiments in a synthetic environment. In a first set of simulations, the actual map (a) is available to the likelihood methods (a perfectly modelled environment). The resulting mean error from the ground truth and the likelihood assigned to the actual robot pose are shown in (b)-(c) respectively for the three methods (BM, LF, and RSLC). The graphs show the evolution of the results with the percentage of employed ranges from the scan. Confidence intervals of 68% are marked in all the graphs. The second set of experiments simulates a dynamic environment by using a slightly different map (d), whose results are plotted in (e)-(f), which reveal that RSLC outperforms the other methods. (g) A close look at part of the estimated trajectory in this case, according to each method.

that the estimation from RSLC is closer to the ground truth and less biased than the others.

B. Real Robot Experiment

To test the robustness of the different likelihood estimation methods against dynamic objects and discrepancies between the map and the real environment we have carried out an experiment where the robot moves throughout a dynamic cluttered room populated with people. A map of the environment was built, and then it was modified by moving furniture, removing objects, etc. The results are summarized in Fig. 5, where the particles are plotted for some instants of time together with the scanner readings projected from the weighted mean given by the particles (please, refer to the online videos for a better grasp of the experiment [18]). It can be visually verified that the estimation using RSLC provides a better alignment of the readings with the map, even in situations where most ranges do not have correspondences into the map. It is remarkable the poor performance of the BM, whose estimation is absolutely wrong from time step t=20sec on, approximately (please, observe that the sensed scans do not match the map at all). This is due to the inability of this model to cope with objects that appeared in the map but were removed afterwards. The opposite case (sensing new objects not present in the map) is typically solved by pre-processing the range scan [6]. By contrast, the RSLC does not require additional artifacts to deal with the problems derived from dynamic environments.

VI. Conclusions

In this paper we have addressed the problem of deriving a likelihood function for highly accurate range scanners. Instead of assuming an unrealistic measurement uncertainty for each range as previous works do, we have presented an accurate likelihood model for individual ranges, which are fused by means of a Consensus Theoretic method. Our main contributions are: the employment of the actual uncertainty in sensed points (according to the sensor precision), the consideration of uncertain in the correspondences with the map, and a sensor fusion method that confers RSLC an unprecedented degree of robustness for applications in dynamic scenarios.

Exhaustive simulations in static, synthetic environments reveal an excellent performance of the RSLC method even when considering only a small fraction of the whole range scan (e.g. 7 range values), while other methods require significantly more measurements to achieve similar results. Furthermore, results for dynamic environments demonstrate a qualitative improvement in the robustness of the robot pose estimation.

These promising results suggest that integrating our approach into probabilistic SLAM methods would improve the building of maps for dynamic, cluttered environments, a challenging issue that requires further research.

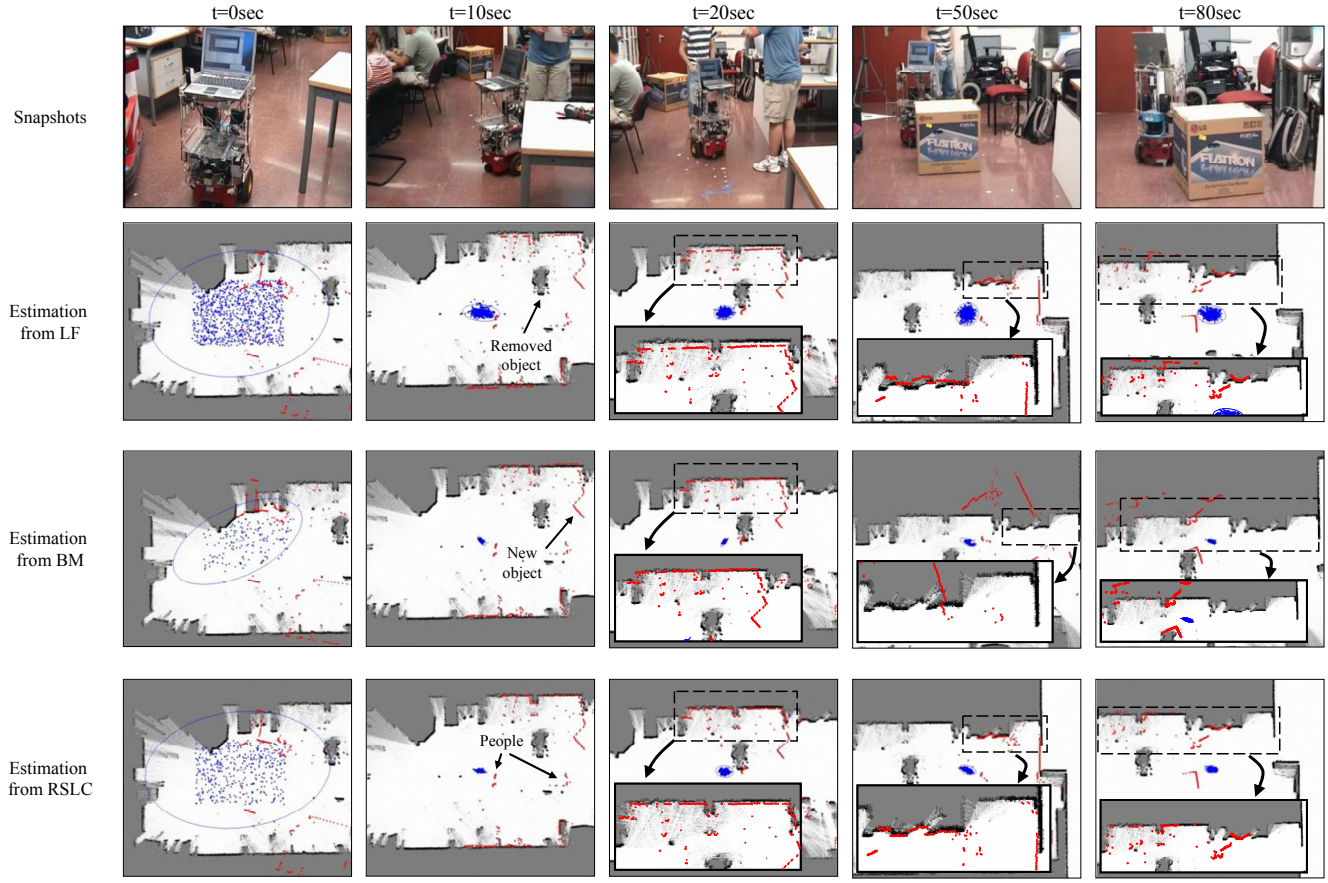

Fig. 5 Results for localization in a real dynamic scenario. Snapshots on the upper row show some instants of time along the robot navigation, while the rest illustrate the evolution of the particle filter for each likelihood computation method. See the text for further discussion.

REFERENCES

[1] J.A. Benediktsson, P.H. Swain, "Consensus Theoretic Classification Methods", in IEEE Trans. On Systems, Man, and Cybernetics, v.22, no.4, pp. 688-692, 1992.

[2] P.J. Besl, N.D. McKay, "A Method for Registration of 3-D Shapes", in *IEEE Tran. on Pattern Analysis and Mach. Int.*, pp. 239-256, 1992.

[3] H. Borotschnig, L. Paletta, A. Pinz, "A Comparison of Probabilistic, Possibilistic and Evidence Theoretic Fusion Schemes for Active Object Recognition", in *Computing*, v.62, no.4, pp.293-319, 1999.

[4] A. Doucet , N. de Freitas, K.P. Murphy, S.J. Russell, "Rao-Blackwellised Particle Filtering for Dynamic Bayesian Networks", in *Proc. of the 16th Conf. on Uncertainty in Art. Int.*, pp. 176-183, 2000.

[5] D. Fox, W. Burgard, F. Dellaert, S. Thrun, "Monte Carlo Localization: Efficient Position Estimation for Mobile Robots", in *Proc. of AAAI-99*, pp. 343-349, 1999.

[6] D. Fox, W. Burgard, S. Thrun, "Markov Localization for Mobile Robots in Dynamic Environments", in *Journal of Art. Int. Res.*, v.11, pp. 391-427, 1999.

[7] A. Elfes, "Using occupancy grids for mobile robot perception and navigation", in *Computer*, v. 22, pp. 46-57, 1989.

[8] M.W.M. Gamini Dissanayake, P. Newman, S. Clark, H.F. Durrant-Whyte, M. Csorba, "A Solution to the Simultaneous Localization and Map Building (SLAM) Problem", in *IEEE Transactions on Robotics And Automation*, v.17, no. 3, pp.229-241 , 2001.

[9] C. Genest, J.V. Zidek, "Combining Probability Distributions: A Critique and an Annotated Bibliography", in *Stat. Science*, v.1, pp.114-135, 1986.

[10] G. Grisetti, C. Stachniss, W. Burgard, "Improving Grid-based SLAM with Rao-Blackwellized Particle Filters by Adaptive Proposals and Selective Resampling", in *Proc. of the IEEE Int. Conf. on Robotics and Automation*, pp. 2432-2437, 2005.

[11] J.S. Gutmann, K. Konolige, "Incremental Mapping of Large Cyclic Environments", in Proc. of IEEE Int. Symp. On Comp. Int. in Robotics and Automation, v.318-325, 1999.

[12] D. Hähnel, W. Burgard, D. Fox, S. Thrun, "An Efficient FastSLAM Algorithm for Generating Maps of Large-Scale Ciclic Environments from Raw Laser Range Measurements", in *IEEE/RSJ Int. Conf. on Int. Robots and Systems*, v.3, pp.206-211, 2003.

[13] D. Hähnel, D. Schulz, W. Burgard, "Map Building with Mobile Robots in populated Environments", in *Proc. of the IEEE/RSJ Int. Conf. on Int. Robots and Systems*, v.1, pp. 496-501, 2002.

[14] J. Kittler, M. Hatef, R.P.W. Duin, J. Matas, "On combining classifiers", in *IEEE Trans. Pattern Anal. Mach. Intell.*, v.20, pp.226–238, 1998.

[15] F. Lu, E. Milios, "Globally Consistent Range Scan Alignment for Environment Mapping", in *Autonomous Robots*, v.4, pp.333-349, 1997.

[16] F. Lu, E. Milios, "Robot pose estimation in unknown environments by matching 2D rangescans", in *Proc. of IEEE Conf. on Comp. Vision and Pattern Recognition*, pp. 935-938, 1994.

[17] L. Montesano, J. Minguez, L. Montano, "Probabilistic Scan Matching for Motion Estimation in Unstructured Environments", in Proc. of the *IEEE/RSJ Int. Conf. on Int. Robots and Systems*, pp.3499-3504, 2005.

[18] Multimedia content, experiment videos available online in `http://www.isa.uma.es/C13/jlblanco/default.aspx`.

[19] E. Parzen, "On Estimation of a Probability Density Function and Mode", in *Annals of Math. Stats*, v.33, no.3, pp.1065-1076, 1962.

[20] S. Thrun, "A Probabilistic Online Mapping Algorithm for Teams of Mobile Robots", in *The International Journal of Robotics Research*, v.20, no.5, pp. 335-363, 2001.

[21] S. Thrun, W. Burgard, D. Fox, *Probabilistic Robotics*, the MIT Press, pp. 153-183, 2006.

[22] S. Thrun, D. Fox, W. Burgard, F. Dellaert, "Robust Monte Carlo localization for mobile robots", in *Artificial Intelligence*, v.128, pp.99-141, 2001.