

The Proceedings of 2006 IEEE/RSJ International Conference on Robots and Intelligent Systems



[Foreword](#)

[Conference Information](#)

[Program at a Glance](#)

[Table of Contents](#)

[Author Index](#)

[Videos](#)

Beijing, China
October 9-15, 2006

IEEE Catalog Number: 06CH37780D
ISBN: 1-4244-0259-X
Library of Congress: 2006921529

© 2006 IEEE. Personal use of this material is permitted. However, permission to reprint/republish this material for advertising or promotional purposes or for creating new collective works for resale or redistribution to servers or lists, or to reuse any copyrighted component of this work in other works must be obtained from the IEEE.

An Entropy-Based Measurement of Certainty in Rao-Blackwellized Particle Filter Mapping

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

Abstract – In Bayesian based approaches to mobile robot simultaneous localization and mapping, Rao-Blackwellized particle filters (RBPF) enable the efficient estimation of the posterior belief over robot poses and the map. These particle filters have been recently adopted by many exploration approaches, to whom a central issue is measuring the certainty inherent to a given estimation in order to be able to select robot actions that increase it. In this paper we propose a new certainty measurement grounded in Information Theory that unifies the two kinds of uncertainty which are intrinsic to SLAM: in the robot pose and in the map content. Most previous works have considered only one of them or a weighted average. Our method combines them more appropriately by first building an *expected map* (EM) which condenses all the current map hypotheses and then computing its *mean information* (MI) – an entropy derived measurement that quantifies the inconsistencies in the EM. Experimental results comparing our method (EMMI) with others verify its correctness and its better behavior for detecting the decrease in certainty when the robot enters unexplored areas and its increase after closing a loop.

Index terms – Mobile robots, SLAM, particle filters, information theory, probabilistic mapping.

I. INTRODUCTION

Automatic acquisition of environment models while simultaneously performing self localization (SLAM) is one of the major challenges for autonomous mobile robots. Probabilistic approaches based on Estimation Theory have received a great attention in the last years for addressing this problem. In particular, sequential Monte Carlo sampling methods (particle filters) are powerful enough to cope with any shape in the probability distributions, non-linear models, and a diversity of hypotheses in localization-only applications ([5],[7]). Estimating both the robot position and the map can be addressed through Rao-Blackwellized Particle Filters (RBPF), which reduces the dimensionality of the estimation problem by marginalizing out some variables, in particular, the map ([1],[6]). RBPF-based approaches maintain a posterior belief over the robot path consisting of a set of particles, i.e. each particle is a path hypothesis, as those ones plotted in Fig. 1(b). Attached to each particle it is also maintained an estimation of the map, which can be analytically determined from the associated path ([6],[9],[10],[14]).

This probabilistic framework underlies many recent works in automatic robot exploration, where the robot usually takes actions with a high expected information gain and a low cost

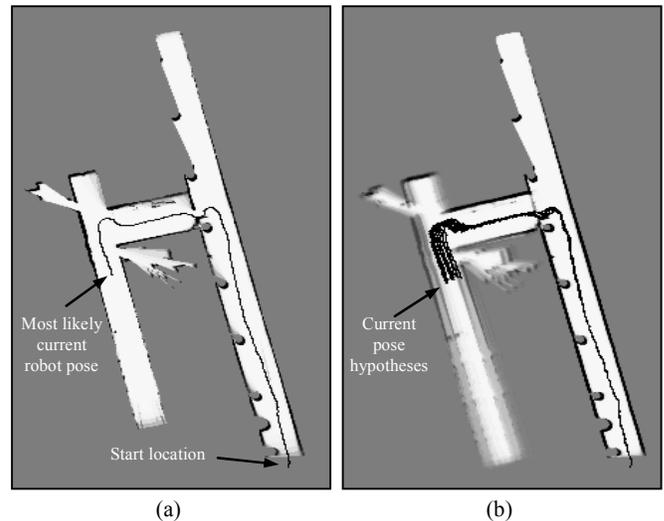


Fig. 1 (a) In RBPF mapping the map associated to the particle with a higher probability is considered the most likely map. (b) The expected map (EM), introduced in this paper, is a weighted map where all particles contribute. The uncertainty in the mapping process can be accurately determined by means of this map.

([3],[4],[14]). Making the robot to actively interact with its environment is of the greatest interest since the complexity of the mapping problem can be drastically reduced [4]. For example, think of a robot just before closing a large loop and having two possible paths: one of them which definitively closes the loop and the other entering a new unexplored area, as represented in Fig. 6(c). If the first option is taken before the second one, the map of the new area to be obtained later will possess a much higher certainty than if the loop is not closed first.

Therefore, a crucial point in the fields of exploration and in active localization is that of measuring the certainty in the mapping process. This can be performed by different means:

- The uncertainty in the robot pose can be estimated through its entropy. This method has been utilized in the contexts of localization and map building [12], localization only [18], and exploration [4]. Since the map is not taken into account this measurement reflects only part of the certainty present in RBPF mapping.
- The volume covered by the particles is used in [15] instead of the entropy. However it also directly reflects only the certainty in the pose, ignoring the map contents.

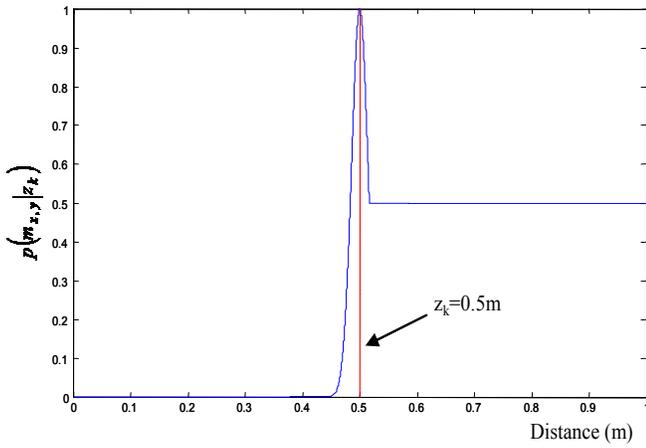


Fig. 2 The probabilistic inverse model of a sensor estimates the occupancy of a given cell, given that some reading z_k has been obtained. In the figure this density is plotted for a laser range finder (and for a measurement of $z_k=0.5m$) as a function of the distance from sensor to cells. The distribution is highly peaked due to the sensor high accuracy.

- The entropy values of both, paths and the map, are considered in [3] individually and then weighted together. A more elegant method is introduced in [14] by computing the joint entropy of both variables. However, they consider the individual maps separately without taking into account the relative consistency between them, as our method does.

In this paper we propose a new entropy-based approach to measure the *certainty* in the information of a RBPF-based map. The underlying intuition in our work is that contradicting hypotheses result in vague map estimations, whereas consistent ones lead to more certain (coherent) information. In this work we define the *expected map* (EM) as the weighted map which arises from taking into account all the particles (see Fig. 1(b)). We claim that this new map not only allows us to estimate the certainty of each original hypothesis but also reflects their relative consistency, therefore integrating better than other approaches both certainty sources: robot paths and contents of each proposed map. We compute the certainty in the resulting map with a novel entropy-based information measurement, namely the *mean information* (MI): our process takes all particles in a RBPF and computes their expected map mean information (EMMI) value, which provides a valuable estimation of the certainty that can be used, for example, in exploration or active localization methods.

It must be highlighted that this paper addresses maps based solely on occupancy grids, although the presented method can be extended to other map types provided that probabilistic representations are available for them. An interesting contribution of our certainty metric in the context of occupancy grids is its independence on the grid resolution and on its size, i.e. its rectangular limits. Thus the resulting EMMI values are contrastable across different experiments and configurations. This is not the case of previous approaches where entropy is applied directly.

In the next section we define the EM in the general framework of Bayesian paths and map tracking, and we also describe its practical implementation with particle filters. In section III we review the concept of the entropy of a map and introduce our certainty metric. Finally, section IV provides experimental results where our method is compared with other, including then some conclusions.

II. THE EXPECTED MAP (EM)

A. Problem formulation

Assume a robot moving in a planar scenario whose pose can therefore be described by $\mathbf{x}=[x \ y \ \phi]^T$, where (x,y) are 2D coordinates and ϕ is the robot heading. The actual pose is not accessible to us and at any given instant of time we are only provided with an estimation in the form of the probability distribution $p(\mathbf{x})$. Bayesian estimation has proved to be a powerful tool in recursively computing that distribution, assuring the convergence of that posterior belief towards the actual robot location, i.e. the *peak* in the distribution will ultimately coincide with the actual robot pose, even for the difficult case of the “robot awake” problem [17]. The Bayesian update rule for the robot pose can be written down in its iterative form as follows:

$$p(\mathbf{x}_n) \propto p(z_n | \mathbf{x}_n) \int p(\mathbf{x}_n | \mathbf{x}_{n-1}, a_{n-1}) p(\mathbf{x}_{n-1}) d\mathbf{x}_{n-1} \quad (1)$$

where a_n and z_n represent the actions of the robot (usually through odometry readings [17] or range scan matching [15]) and its observations at step n respectively. The density $p(\mathbf{x}_n)$ is estimated up to a proportionality constant, thus must be normalized at each iteration.

Apart from the robot pose, the environment must also be represented in some way if we are carrying out SLAM. In this work we consider occupancy grids as the world model (map). This representation of space is very popular in the robotics community, having been widely employed during the last twenty years ([8],[9],[11],[16]). An occupancy grid is a random field where we store the occupancy likelihood for each cell, which we will denote as:

$$p(m_{x,y}) \quad (2)$$

for any cell with indexes $\langle x,y \rangle$ in the map m . If no prior information is available about the obstacles present in a given environment, all cell occupancies can be set to 0.5: there is the same likelihood of any cell to be occupied or free. Subsequently sensors measurements update the map, and therefore that initial value. For that update, our method considers the inverse sensor model,

$$p(m_{x,y} | z_k) \quad (3)$$

that is, the cells occupancy likelihood conditioned to a given observation z_k . Fusing the current observation with the previous contents of the map is done on a Bayesian basis by using the following iterative expression:

$$p(m_{xy} | z_{1:n}) = \left(1 + \frac{1 - p(m_{xy} | z_{1:n-1})}{p(m_{xy} | z_{1:n-1})} \cdot \frac{1 - p(m_{x,y} | z_n)}{p(m_{x,y} | z_n)} \right)^{-1} \quad (4)$$

which can be easily derived from the log-odds representation of the update process [15] if we assume an initial occupancy

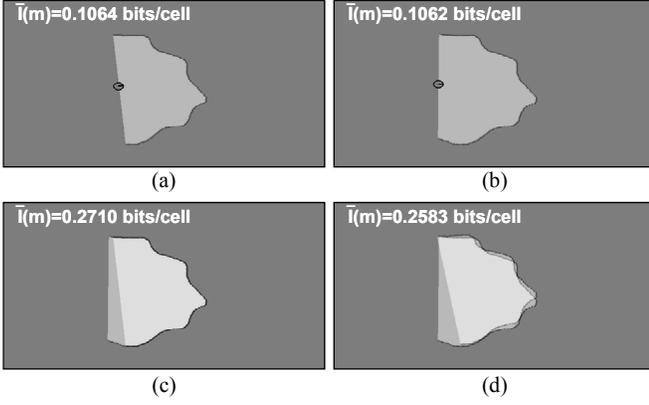


Fig. 3 (a)-(b) Scans and MI values for their associated maps. (c) When both scans are consistently fused in the same grid, the information value increases. (d) Inconsistencies due to poor robot localization decrease the quality of the fused map, which is confirmed by a lower MI value.

likelihood of 0.5 for all cells¹. Essentially, (4) increases the certainty in the occupancy/freeness of a given cell if subsequent observations confirms the earlier belief. The only density required to iterate (4) is the inverse sensor model of (3). For the case of a laser range scanner we consider the function depicted in Fig. 2.

At this point we have defined stochastic representations for both the robot pose and the map (in (1) and (4), respectively), thus we could compute their entropy values separately to measure the uncertainty in both estimations. However, we propose to first build a new map, the EM, with the intention of revealing the inconsistencies between maps associated to different particles in the RBPF.

B. Definition of the Expected Map

The expected map (EM) is the map expectation over all paths possibilities, given by the density $p(\mathbf{x}_{1:n})$, and for a set of associated observations $z_{1:n}$. In the case of occupancy grids, the EM is another grid defined as:

$$EM_{x,y}(p(\mathbf{x}_{1:n}), z_{1:n}) = E_{\mathbf{x}_{1:n}} [p(m_{x,y} | \mathbf{x}_{1:n}, z_{1:n})] \quad (5)$$

$$= \int p(m_{x,y} | \mathbf{x}_{1:n}, z_{1:n}) p(\mathbf{x}_{1:n}) d\mathbf{x}_{1:n}$$

for all cell indexes $\langle x,y \rangle$. Here the density $p(\mathbf{x}_{1:n})$ is supposed to be available in a continuous form, an impractical assumption. If this density is rather represented as a set of M particles the EM can be approximated with:

$$EM_{x,y}(p(\mathbf{x}_{1:n}), z_{1:n}) \approx \sum_{i=1}^M w^{[i]} p(m_{x,y}^{[i]} | \mathbf{x}_{1:n}^{[i]}, z_{1:n}) \quad (6)$$

The intuition behind (5)-(6) is that, by contrasting the occupancy values $p(m_{xy})$ of cells at the same location (the same $\langle x,y \rangle$ indexes) but from the different maps associated to each particle, we can prove the relative coherence between all the hypotheses. The result can be visualized as a “fuzzy” map, as in Fig. 1(b), where the sharper the image, the more certain is the map information for the associated area.

III. A CERTAINTY METRIC FOR THE EM

A. Definition of the Mean Information

In the following we derive an information measurement which stresses the certainty in the occupancy likelihood of cells in the grid in opposition to measuring the total *amount* of information in the map (which is related to the number of observed cells).

Information theory establishes that the amount of information associated with a random variable is related to its entropy [13]. Consider the entropy of a single cell in the grid,

$$H(p(m_{xy})) = -p(m_{xy}) \log p(m_{xy}) - (1-p(m_{xy})) \log(1-p(m_{xy})) \quad (7)$$

which is the entropy of a discrete random variable with two possible outcomes, i.e. a Bernoulli distribution. Notice that the maximum entropy is obtained for $p(m_{xy})=0.5$, that is, for unobserved cells. Recalling the independency assumption between cells, the entropy for the whole map turns into:

$$H(p(m)) = \sum_{x,y} H(p(m_{xy})) \quad (8)$$

This estimation of the entropy is widely used as a measurement of the information in the map ([3],[14]). However, it exhibits the following limitations:

1. Its absolute value depends on the grid size (the rectangular limits of the map) instead of the actually observed area. Note that from (7) it follows that unobserved cells will contribute to the global entropy with a maximum entropy value.
2. It depends also on the grid resolution, since it settles (together with the map limits) the total number of cells in the map. This means that the entropy of any map with unobserved areas (all maps in practice) increases without bound when resolution increases.

To overcome these drawbacks we employ the information (I) of a map instead of its entropy:

$$I(p(m_{xy})) = 1 - H(p(m_{xy})) \quad (\text{bits})$$

$$I(m) = \sum_{x,y} I(p(m_{xy})) \quad (\text{bits}) \quad (9)$$

where the entropy $H(\cdot)$ is computed by substituting the natural logarithms in (7) by base-2 ones. As a result we obtain a natural unit for information: *bits*. Noticeably the maximum information value (1 bit) is given to a certainly occupied/free cell while the minimum value (0 bits) is associated to any unobserved cell. The entropy dependency on the grid size is therefore avoided: the limits of the map become irrelevant since all unobserved cells now contribute with null information. Thus, the resulting measurement is a more practical quantifier of the amount of information carried by the map than the direct application of the entropy.

However, in this work we are not interested in the absolute amount of information in the map but in its *certainty*. To effectively reflect the certainty in a map m we introduce its *mean information* (MI), defined as

¹ Here the Markov assumption is not considered since all the past states are involved in the derivation of the current map.

$$\bar{I}(m) = \begin{cases} \frac{I(m)}{N_{obs}} & \text{if } N_{obs} > 0 \\ 0 & \text{otherwise} \end{cases} \quad (\text{bits/cell}) \quad (10)$$

where N_{obs} represents the number of observed cells in the map, i.e. cells with an occupancy likelihood different from 0.5. Notice that MI delivers bounded values in the range $[0,1]$. As an illustrative example of how the MI fits our purposes please consider the occupancy grids in Fig. 3. The first pair of maps represents two observations of the same environment from slightly different poses. In Fig. 3(c) both are merged into the same map (by means of (4)) with the correct alignment: each observation confirms the occupancy values of the overlapped cells (resultant occupancy values are closer to 0 and 1) and, as desired, the resultant MI value is greater than in maps where only one observation is inserted. This behavior follows from the properties of the information as defined in (9), whose maximum values are obtained for occupancy values of 0 and 1. When both observations are misaligned the consequent inconsistencies reduce the MI value, as shown in Fig. 3(d). Notice that in fact, the misaligned map contains more observed cells than the previous case (the total area of the misaligned mixture is bigger than that of the perfectly aligned maps), but the certainty of those cells is lower, in terms of the mean information.

We can highlight the following properties of the MI which set it apart from classical entropy-based measurements:

1. An empty map (containing only unobserved cells) has a null mean information value.
2. It is mostly independent on the grid resolution for practical cell sizes. This is shown in the maps of Fig. 4(a)-(c), whose MI increase as we consider higher resolutions. As the resolution increases, the MI asymptotically tends towards a maximum value, as appreciated in Fig. 4(d). This presents a great difference with the direct entropy behavior, which in that case tends to infinite. The asymptotical behavior of MI depends on the inverse sensor model, the specific environment being mapped, and other factors. In [2] we have derived a theoretical expression for this bound for the particular case of a circular synthetic environment.
3. The better the alignment between observations into the map, the higher the obtained MI values, as Fig. 3(c)-(d) illustrate. The intuitive idea behind this property is that well-aligned observations make the occupancy likelihood of cells to become almost 0 or 1, which correspond to maximum MI values. This property is crucial in understanding the highly distinctive behavior of MI after a loop closure, as shown later.

At this point we can now recover the problem of measuring the certainty in RBPF mapping since we have already defined new tools, namely the EM of a RBPF and the MI of a map. We denote the proposed certainty metric as the EMMI (*expected map mean information*). Formally,

$$EMMI(p(\mathbf{x}_{1:n}), z_{1:n}) = \bar{I}(EM(p(\mathbf{x}_{1:n}), z_{1:n})) \quad (11)$$

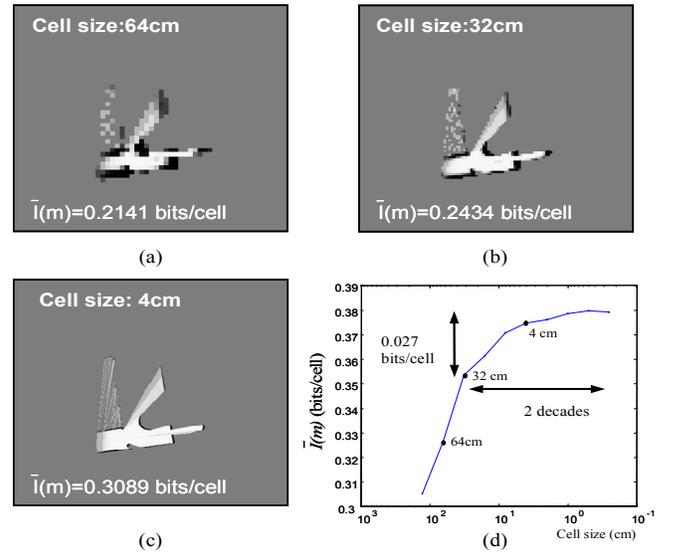


Fig. 4 (a)-(c) The occupancy grids obtained from mapping the same environment with different resolutions. (d) The mean information (MI) is plotted for different grid resolutions, from 1m to 2.5mm. It is remarkable the small increase in information produced by increasing the resolution of a map in the highlighted range, which coincides with commonly used resolutions. Therefore, in practice the MI can be regarded as resolution independent.

that is, the value resulting of computing the MI of the EM for a given RBPF at a given instant of time n .

B. Comparison with other approaches

Our proposed method can be compared to the entropy of the poses ([4],[12],[18]) and to the joint entropy [14]. Next we provide a theoretical discussion about the differences between those methods while in the next section experimental values are obtained for each of them.

Regarding their computational time complexities the most simple method is the entropy of poses since it implies $O(M)$ operations only, where M is the size of the particle population. On the other hand, the joint entropy has a complexity of $O(M(N+T))$, where T is the length of the paths and N the cells count in the grid. Clearly the cell count will be the dominant quantity in most situations, therefore the complexity tends toward $O(MN)$ which coincides with that of our method. Therefore, the entropy of poses is many orders of magnitude less complex than the others, at the cost of not considering the map content. It remains being a useful indicator of the quality of the pose estimation, anyway.

Concerning the space complexities, our method requires an additional map (the EM) to be maintained apart of the M maps attached to each particle in the RBPF, whereas the joint entropy does not require this supplementary storage.

The results of the joint entropy and the EMMI have quite different interpretations. We can realize of that by expanding the expression for the first estimator [14]:

$$H(p(\mathbf{x}_{1:n}, m_{xy} | z_{1:n})) = H(p(\mathbf{x}_{1:n})) + \int_{\mathbf{x}_{1:n}} p(\mathbf{x}_{1:n} | z_{1:n}) H(p(m_{xy} | \mathbf{x}_{1:n}, z_{1:n})) d\mathbf{x}_{1:n} \quad (12)$$

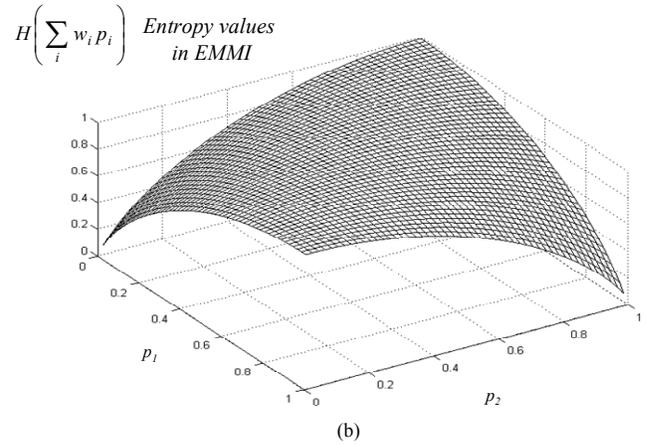
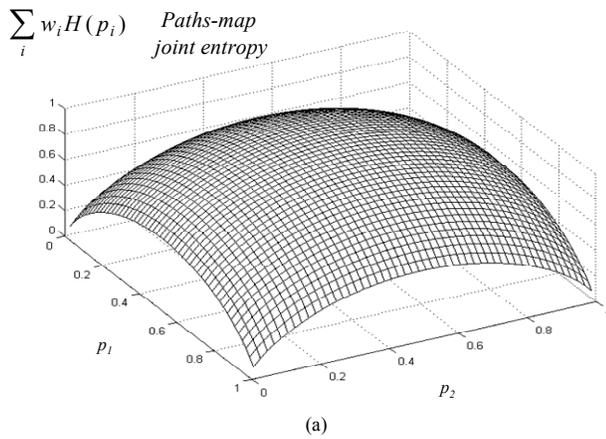


Fig. 5 Factors involved in the joint entropy (a) and in the EMMI (b) are plotted for the case of only two particles as an example of how both approaches deal with different occupancy values for the same cell in maps from different hypotheses. It is clear that the first one considers as certain (low entropy) contradictory values, like $p_1=0$ (the cell is free) and $p_2=1$ (the cell is occupied), whereas EMMI only gives low entropy values when both likelihoods are alike (consistent).

where it can be seen how the entropy of the map contents (m_{xy}) contributes to the total value weighted by the paths likelihood $p(\mathbf{x}_{1:n}|\mathbf{z}_{1:n})$. On the other hand, the EMMI considers the entropy of cells in the EM (by means of the MI as defined in (10)) which, for its part, is a weighted average of maps from each particle. Consequently, EMMI computes the entropy of the average of maps, whereas the joint entropy performs these operations in the opposite order. To graphically see the important implications of this difference, Fig. 5 shows the entropy factors regarding the map contents in each of these methods for the case of only two equally probable particles (this simplification is to allow values to be represented in 3D). The key distinctive point of EMMI is that inconsistencies like one hypothesis stating a cell is free (a value near 0) and the other stating it is occupied (a value near 1) are detected by EMMI as uncertainty, i.e. higher entropy values, see Fig. 5(b). On the other hand, the joint entropy, in Fig. 5(a), considers these cases instead as certainty (low entropy values). Obviously the EMMI is, by design, better suited to detect the relative consistency between several proposed maps.

IV. RESULTS AND CONCLUSIONS

To experimentally compare our method to previous ones we have considered the mapping of part of our building by means of a RBPF. Our algorithm is based on ideas from [9] to allow the mapping of large areas through a small number of particles. We use only 15 in this experiment. A video file of this experiment is available in <http://www.isa.uma.es/C16/research/>.

Fig. 6(a)-(c) show the EMs until the robot has almost completed a loop. At that point we consider two possible actions: to definitively complete the loop or to enter an unexplored area, marked as A and B , respectively. If action A is performed as in Fig. 6(d), the return to the start location makes the likelihood of some particles to become much higher than others, resulting in a resampling where incorrect hypotheses vanish: the estimations of robot path and map become much better than before. To measure this change in

the certainty we use the EMMI, plotted in Fig. 6(f). The most remarkable feature is the gradual degradation of the certainty (low EMMI values) until step 85 approximately, where the loop is completely closed and both the robot pose and the map estimations become closer to the truth (compare Fig. 6(c)-(d)). The abrupt falls in the EMMI at steps 12 and 27 correspond to increases in the uncertainty of the map while the robot turns corners to enter new corridors.

On the other hand, if the robot takes the other option (action B), a new area will be mapped but with a high degree of uncertainty. This is clear from the “fuzzy” appearance of the EM in Fig. 6(e). The EMMI for this case is the dashed plot in Fig. 6(f), where the difference with the other choice is patent: the certainty keeps decreasing whereas in the loop closure it raises.

This experiment shows how EMMI is well-suited to measure the overall certainty in a RBPF, but it must be contrasted with other estimators. The entropy of the poses, as defined in [4], is plotted for both discussed situations in Fig. 6(g). This measurement correctly captures the increase in uncertainty, but in the case of particles being resampled (as occurs in both situations) it restores its maximum value. Notice that this entropy only does not suffice to determine if the particles resampling is due to a loop closure or to their degeneration. The other uncertainty estimator to be compared with our method is the joint entropy [14], which evolves as shown in Fig. 6(h). As discussed in this work, the factors involved in this estimator do not consider the consistency between different hypotheses, but only the individual map contents. A particular effect of considering separately the entropy of each map is that the joint entropy decreases when exploring new areas (there are less unobserved cells in the individual maps) independently of the certainty in the robot pose estimation. Thus, no relevant information can be extracted from Fig. 6(h) indicating the drastic difference in the certainty of the RBPF after taken actions A and B . We can therefore conclude that EMMI stands out as a better certainty estimator than others. In addition, it exhibits other interesting

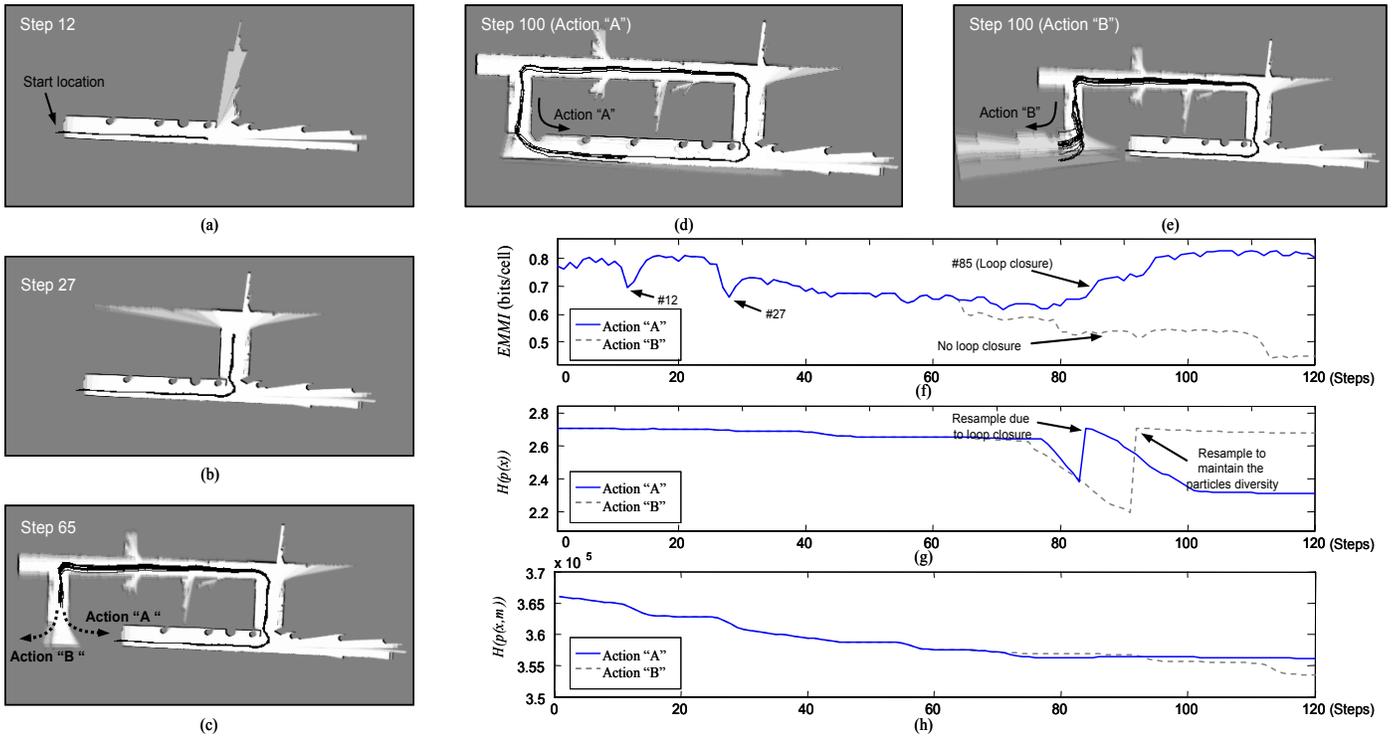


Fig. 6 The EM for a RBPF based mapping are shown in (a)-(c) for three instants of time. We consider two possible paths for the robot at step 65. One of them (d) definitively closes the loop and reduces the uncertainty, as clearly detected by the EMMI (f). The other path (e) makes the robot to enter an unexplored area, increasing the uncertainty as the EMMI also reflects. Other indicators are also plotted in (g)-(h): the poses entropy and the joint paths-map entropy, respectively. We claim that EMMI reflects more effectively the real uncertainty in the mapping process than the other measurements.

distinctive features, as being almost insensitive to map size and resolution and providing values in a bounded range.

To summarize, in this paper we have reviewed how RBPF can efficiently solve the SLAM problem and why measuring the certainty is an important issue in active localization and exploration methods. A novel measurement has been introduced, the EMMI, which reveals valuable information about the certainty in a RBPF mapping system at a given instant of time. This method represents a computational complexity similar to others but with a highly distinctive behavior derived from its novel consideration of the relative consistency between all the hypotheses. Future research is aimed to integrate our certainty estimator with existing active exploration approaches.

REFERENCES

- [1] C. Andrieu, N. Freitas, A. Doucet, "Rao-Blackwellised Particle Filtering via Data Augmentation", in *Advances in Neural Information Processing Systems 14*, MIT Press, 2002.
- [2] J.L. Blanco, "Derivation of an Expression for the Mean Information Maximum Value", *TR*, Dept. of System Engineering And Automation, Univ. of Malaga, 2006.
- [3] F. Bourgault, A.A. Makarenko, S.B. Williams, "Information based adaptive robotic exploration", in *IROS02*, pp. 540-545, 2002.
- [4] W. Burgard, D. Fox, S. Thrun, "Active Mobile Robot Localization by Entropy Minimization", in *Proc. of Second Euromicro Workshop on Advanced Mobile Robots*, 1997.
- [5] F. Dellaert, D. Fox, W. Burgard, S. Thrun, "Monte Carlo localization for mobile robots", in *ICRA99*, v2. pp. 1322-1328, 1999.
- [6] A. Doucet, N. Freitas, K. Murphy, S. Russel, "Rao-Blackwellised Particle Filtering for Dynamic Bayesian Networks", in *Proc. of Conf. on Uncertainty in Artificial Intelligence*, pp. 176-183, 2000.
- [7] A. Doucet, S. Godsill, C. Andrieu, "On sequential Monte Carlo sampling methods for Bayesian filtering", in *Statistics and Computing*, v. 10, pp. 197-208, 2000.
- [8] A. Elfes, "Using occupancy grids for mobile robot perception and navigation", in *Computer*, v. 22, pp. 46-57, 1989.
- [9] G. Grisetti, C. Stachniss, W. Burgard, "Improving Grid-based SLAM with Rao-Blackwellized Particle Filters by Adaptive Proposals and Selective Resampling", in *ICRA05*, pp. 2443-2448, 2005.
- [10] M. Montemerlo, S. Thrun, "Simultaneous Localization and Mapping with Unknown Data Association Using FastSLAM", in *ICRA03*, v. 2, pp. 1985-1991, 2003.
- [11] H.P. Moravec, A. Elfes, "High Resolution Maps from Wide Angle Sonar", in *ICRA85*, v.2, pp. 116-121, 1985.
- [12] N. Roy, W. Burgard, D. Fox, S. Thrun, "Coastal navigation-mobile robot navigation with uncertainty in dynamic environments", in *ICRA99*, v.1, pp. 35-40, 1999.
- [13] C.E. Shannon, "A mathematical theory of communication", in *Bell System Technical Journal*, v. 27, pp.379-423 and 623-656, Jul. and October, 1948.
- [14] C. Stachniss, G. Grisetti, W. Burgard, "Information Gain-based Exploration using Rao-Blackwellized Particle Filters", in *Proc. of Robotics: Science and Systems (RSS)*, 2005.
- [15] C. Stachniss, D. Hähnel, W. Burgard, "Exploration with Active Loop-Closing for FastSLAM", in *IROS04*, v.2, pp. 1502-1510, 2004.
- [16] S. Thrun, "Learning Occupancy Grid Maps With Forward Sensor Models", in *Autonomous Robots*, v. 15, no.2, pp. 111-127, 2003.
- [17] S. Thrun, D. Fox, W. Burgard, F. Dellaert, "Robust Monte Carlo Localization for Mobile Robots", in *Artificial Intelligence*, v.128, pp.99-141, 2000.
- [18] N. Vlassis, Y. Motomura, B. Kröse, "An information-theoretic localization criterion for robot map building", in *Proc. of Int. Conf. on Machine Learning and Applications*, 1999.