J. L. Blanco
J. A. Fernández-Madrigal
J. Gonzalez
Department of System Engineering and Automation,
University of Málaga, 29071 Málaga, Spain
{jlblanco,jafma,jgonzalez}@ctima.uma.es

A Novel Measure of Uncertainty for Mobile Robot SLAM with Rao–Blackwellized Particle Filters

Abstract

Rao–Blackwellized particle filters (RBPFs) are an implementation of sequential Bayesian filtering that has been successfully applied to mobile robot simultaneous localization and mapping (SLAM) and exploration. Measuring the uncertainty of the distribution estimated by a RBPF is required for tasks such as information gain-guided exploration or detecting loop closures in nested loop environments. In this paper we propose a new measure that takes the uncertainty in both the robot path and the map into account. Our approach relies on the entropy of the expected map (EM) of the RBPF, a new variable built by integrating the map hypotheses from all of the particles. Unlike previous works that use the joint entropy of the RBPF for active exploration, our proposal is better suited to detect opportunities to close loops, a key aspect to reduce the robot path uncertainty and consequently to improve the quality of the maps being built. We provide a theoretical discussion and experimental results with real data that support our claims.

KEY WORDS—localization, mapping

1. Introduction

The automated mapping of unknown environments is one of the fundamental problems that need to be solved to achieve truly autonomous mobile robots. The difficulty of this task follows from the fact that a precise map can only be obtained from a well-localized robot, but in turn the quality of the robot pose estimation depends on the map accuracy: this is the simultaneous localization and mapping (SLAM) problem. In recent years, methods based on estimation theory have dominated the research in this field. In these approaches the robot pose and the map are represented in terms of probability densities which are tracked over time. Tracking is usually performed through a Bayesian filtering of either Gaussian (Dissanayake et al. 2001) or Gaussian mixture (Porta and Kröse 2006) distributions with Kalman filters, or sequential Monte Carlo (SMC) sampling schemes with particle filters (Doucet et al. 2000b). More concretely, the ability of SMC methods to solve the problems of global localization (Dellaert et al. 1999) and SLAM (Doucet et al. 2000a) efficiently has been demonstrated, because they can deal with non-linear models and any shape, multi-modal distributions. From this family of methods, Rao–Blackwellized particle filters (RBPFs) are widely employed to estimate both the robot path and the map simultaneously (Grisetti et al. 2007b; Montemerlo et al. 2002). In this scheme, probability densities are maintained by a set of weighted particles, which are hypotheses for the robot path. The Rao–Blackwellization consists of deriving maps analytically from these paths, which reduces the dimensional complexity of the SLAM problem (Doucet et al. 2000a). The most likely path (and, therefore, map) is usually considered to be that associated with the particle with the highest weight, as the example in Figure 1(a) shows.

In general, SLAM methods passively process incoming sensor data and update the map and path estimates iteratively. However, the advantages of allowing the robot to actively control its movements while building a map, that is, active exploration, are well known and have been reported in the literature (Stachniss et al. 2004; Sim and Roy 2005). Exploration methods aim at controlling a robot through an unknown scenario in such a way that the whole environment is mapped while minimizing some cost function, such as the total distance traveled (Yamauchi 1998; Burgard et al. 2000) or the uncertainty in localization (Stachniss et al. 2005b). To illustrate the potential impact of movement selection in the accuracy of the resulting map, consider the example in Figure 2(a) where a robot ex-

The International Journal of Robotics Research
Vol. 27, No. 1, January 2008, pp. 73–89
DOI: 10.1177/0278364907082610

Figures 3, 5, 8–10 appear in color online: http://ijr.sagepub.com

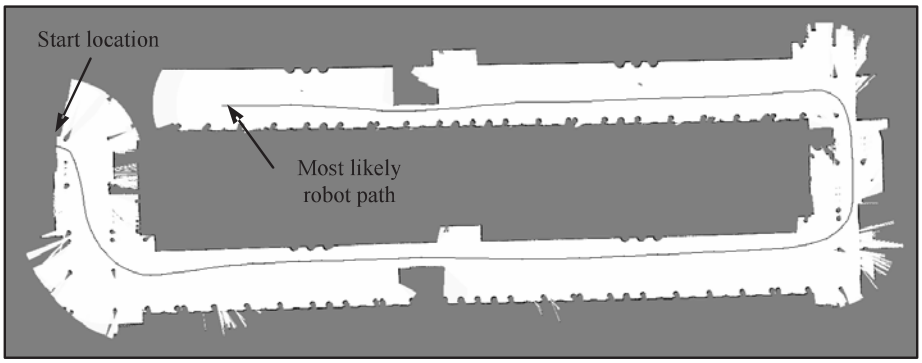

(a)

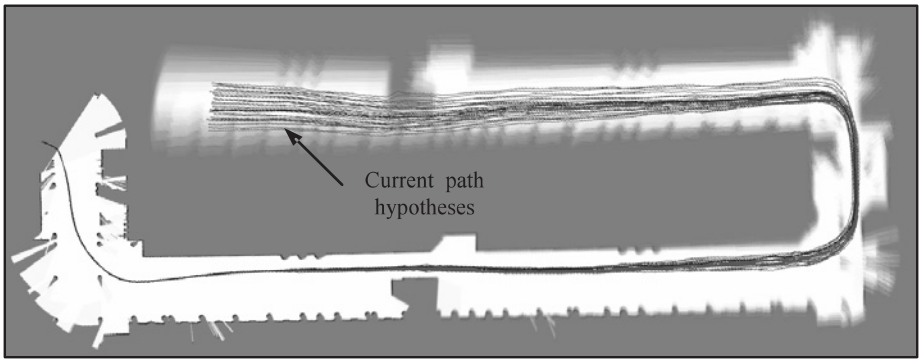

(b)

Fig. 1. (a) The most likely map in RBPF mapping is the map associated with the particle with the highest weight. (b) In this paper we introduce the expected map (EM), a weighted average of maps computed from the contribution of all of the particles in the RBPF. The uncertainty of the RBPF in both the robot path and the map content can be effectively determined by this new map.

plores an environment with a loop (an equivalent example from experimental data was reported by Stachniss et al. (2004)). In Figure 2(b) it has traveled almost all the way round. The uncertainty in the localization along the path increases as long as the robot does not revisit a known area, where the registration between recent and past observations would reset the accumulated errors (Lu and Milios 1997). This is schematically illustrated in Figure 2(b) with larger uncertainty ellipses as the robot moves further away from the origin. Next, the corridor on the right can be explored, but we should consider the convenience of closing the loop. In the first case (see Figure 2(c)), revisiting the known area drastically reduces the localization uncertainty, thus the final map will be more precise than if the loop is not closed, the situation illustrated in Figure 2(d). Note that higher uncertainty in the localization makes the detection of the closure of large loops more difficult. Therefore, the mapping process will be easier if the robot is prone to revisiting known places as soon as possible to reduce its localization uncertainty.

In addition to the uncertainty in the robot path, uncertainty also exists in the maps owing to the lack of knowledge about unexplored areas, the noisy nature of sensors, and the misalignment of observations from a poor localization. Any integrated approach considering SLAM and active exploration must take into account both uncertainty sources when choosing movement actions. On the one hand, it is desirable to take the robot towards unknown places in order to incorporate new information into the map, but, on the other hand, this will usually decrease accuracy in the localization (until revisit). Approaches in the literature quantify only one of these opposing factors, both of them, or propose some kind of combined measurements (Bourgault et al. 2002; Stachniss et al. 2005a). Measuring the uncertainty of a RBPF is also a fundamental part of the techniques for mapping environments with multiple nested loops (Stachniss et al. 2004, 2005b).

In this paper we introduce a new uncertainty measure that simultaneously considers the robot path and the map content. The key idea behind our paper is that of evaluating the consistency between individual maps from all of the hypotheses (particles) in a RBPF: because these maps are generated from each path hypothesis, we are also implicitly evaluating the consistency between the estimated paths. Note that in a RBPF the estimated paths are stored only because they determine the map hypotheses. To evaluate the consistency between hypotheses

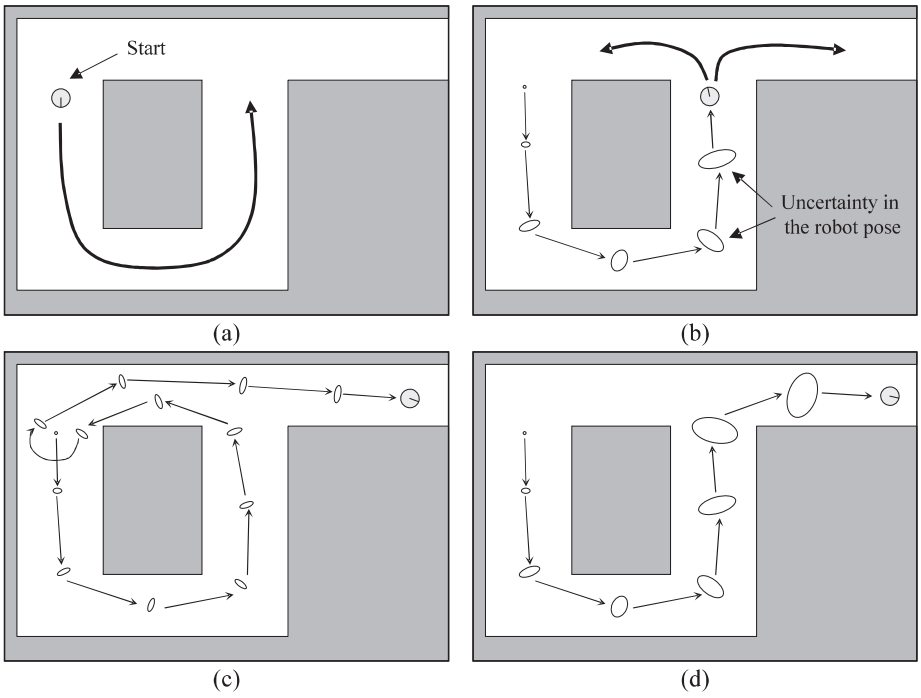

Fig. 2. (a) and (b) The uncertainty in the robot pose while exploring an unknown environment increases as long as it does not revisit a known area, as schematically illustrated with uncertainty ellipses. If the robot closes the loop (c) to reduce its uncertainty before exploring a new corridor (d), the resulting map will be much more accurate. Therefore, exploration methods must be favorable to performing loop closing.

we introduce the expected map (EM), a weighted average of maps computed from the contribution of all of the particles in the RBPF. As can be observed with the example in Figure 1(b), inconsistencies between maps appear in an EM as blurred areas. Instead of measuring the uncertainty in this map directly through its entropy, we define here the information (I) and the mean information (MI) of a grid map. Both information metrics are based on the entropy but avoid some of its undesirable properties when applied to grid maps. Concretely, both measures are independent of the grid map extent, while the MI is also practically independent of the map resolution for typical grid cell sizes. Although the information and the MI can be applied to any grid map in general, they are introduced in this paper to obtain a measure of the uncertainty in the EM of a given RBPF. In summary, in this paper we propose two different uncertainty measures:

(i) The MI of the EM (EMMI) provides an alternative to the entropy of the robot path in applications such as loop closure detection (Stachniss et al. 2004). As discussed in the following, the entropy of the path may not be well defined in the context of a RBPF owing to the sparsity of the particle representation, which does not affect EMMI.

(ii) The information of the EM (EMI) performs more appropriately than other measures for the case of exploration using occupancy grid maps, because actions aimed at closing loops are given more importance than the exploration of new areas when the current pose uncertainty is high. The benefits of closing loops while exploring have been reported previously (Stachniss et al. 2004, 2005). We discuss the theoretical ground of this behavior as well as providing real-world exploration experiments.

We must remark that our proposal is not a different approach for computing the joint entropy of a RBPF. Instead, we propose to measure the uncertainty in a RBPF by means of the entropy of a new variable (the EM of the RBPF) to avoid the undesirable property of the joint entropy of missing opportunities to close loops in grid map-based exploration.

The rest of this paper is organized as follows. In Section 2 we review previously proposed uncertainty measures. Then, in Section 3, the EM is defined in the general framework of SLAM and in the concrete case of a RBPF. We introduce our new information metrics for grid maps in Section 4 and in Section 5 we compare them to other uncertainty measures. We then show experimental results for active exploration and loop-closure detection using our proposed measures

in Section 6, and, finally, we provide some conclusions in Section 7.

2. Related Works

Probabilistic approaches have dominated the research in robot localization, map building, and active exploration for more than a decade. Classically, the most widely employed probabilistic representation for robot poses has been the multivariate Gaussian distribution, where a mean pose vector and a covariance matrix are maintained (Dissanayake et al. 2001). This approach can be used in the context of position tracking (where the initial pose distribution is known), global localization (there is no knowledge about the robot starting location), and landmarks-based SLAM (both the pose of the robot and a set of landmarks is to be estimated). This latter perspective has been dealt with through Kalman filtering (KF; see Kalman (1960)) or extended Kalman filtering (EKF; see Julier and Uhlmann (1997)). Uncertainty in these methods is related to the covariance matrices, thus the entropy of such matrices becomes a natural uncertainty measurement (Vlassis et al. 1999; Bourgault et al. 2002). However, the intractable growth in complexity of these filters ($O(N^2)$ for N landmarks), their restrictive assumption of Gaussian distributions, and their inability to deal with raw data (features must be extracted) have led to the proposal of more efficient approaches in recent years. In the scope of global localization, methods based on localization probability grids (Fox et al. 1999) avoid the Gaussian assumption by dividing the pose state space into a regular grid. The uncertainty of the robot pose can then be estimated by the entropy of the density discretized in a grid (Burgard et al. 1997; Roy et al. 1999). Nevertheless, key disadvantages of this space representation are its extensive computation and storage requirements.

More recently, Monte Carlo methods (particle filters) have gained huge popularity in the mainstream robotic research owing to their advantages: a limited number of particles ensure a bounded computation complexity and, because particles are focused on the regions of interest, the resources are employed efficiently. We must distinguish between the particle filters used for localization (Dellaert et al. 1999) and for SLAM (or active exploration; see Stachniss et al. (2005a)). The former assume that the map is known *a priori* and thus the robot pose is the only variable to estimate. However, the dimensionality of the SLAM problem (the problem addressed in this paper) is higher than in the case of localization because the content of the map must also be estimated. RBPFs represent a very efficient solution to these problems, allowing us to simultaneously estimate both the robot poses (path) and the map (Doucet et al. 2000a). There are some proposed measurements for the uncertainty of the robot pose only (ignoring the map), such as the volume covered by particles (Stachniss et al. 2004) or the entropy of the Gaussian distribution which approximates to the particles (Stachniss et al. 2005a). A particularly original proposal can be found in Fox (2003), where the Kullback–Leibler distance is used to measure (and bound) the error caused by the approximation of the pose distribution by a discrete set of particles.

Regarding the uncertainty in maps, entropy has also been employed as a measurement for different map representations: point-clouds, landmarks, or occupancy grids (Bourgault et al. 2002; Saez and Escolano 2005). However, in RBPF mapping we have multiple map hypotheses simultaneously. Thus, the entropy of maps needs a proper mechanism to explicitly consider the consistency between hypotheses, a role played by the EM in this paper and not found, to the best of our knowledge, in any previous work. We believe that although entropy is a founded measurement of the information in maps, its direct use on maps has some important drawbacks (as exposed in the next section). This is the reason why we propose only entropy-based measures, i.e. the information and the MI of a grid map, which solve these shortcomings. The methods mentioned above consider only the uncertainty either in the robot pose or the map. Others have proposed combined estimators which take into account both sources of uncertainty. We think that considering one uncertainty source only is inappropriate for active exploration approaches if we desire to obtain both a good localization and a consistent resulting map. Bourgault et al. (2002) computed and weighted individual entropies together, where weights are obtained experimentally. A more elegant method was reported by Stachniss et al. (2005a), which computes the joint entropy of the two variables. In spite of this latter method being mathematically founded, a number of problems discussed in this paper limit the utility of the method in practice. Interestingly, drawbacks to the use of joint entropy as the policy in exploration are also discussed by Sim and Roy (2005), where an alternative uncertainty measure is presented for EKF-based exploration.

In the context of multi-robot exploration, Ko et al. (2003) proposed separate heuristic functions for the cases of exploring a new area and meeting another robot (an event equivalent to a loop closure), while in (Burgard et al. 2000) other non-entropy based utility functions are proposed to prevent different robots from exploring the same areas.

From the methods discussed above, the joint entropy is used throughout this paper as a reference for comparisons because it is one of the most significant measures in the context of RBPF-based grid mapping.

3. The EM of a RBPF

In this section we first state the SLAM problem from a sequential filtering viewpoint, which includes RBPF. Next, we introduce the concept of an EM for such a particle filter and discuss its properties.

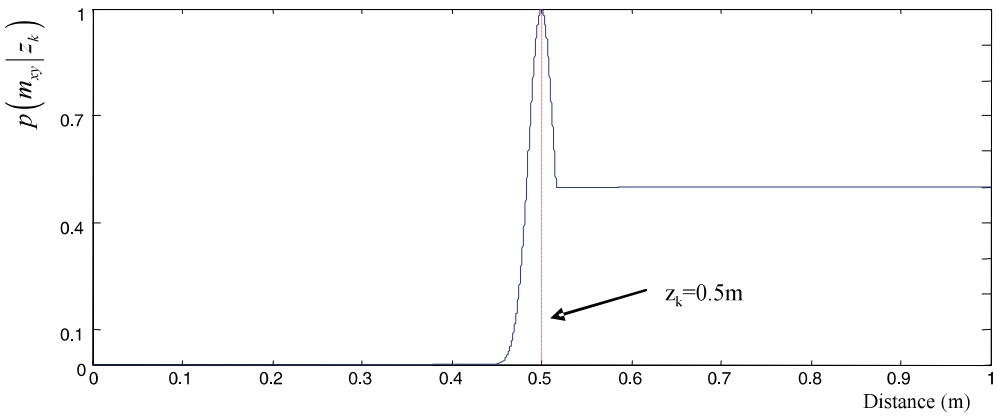

Fig. 3. The probabilistic inverse model of a sensor estimates the occupancy of map cells given that some reading z_k (distance to obstacle) has been obtained. This density is plotted for a laser range finder (and for a measurement of $z_k = 0.5m$) as a function of the distance from the sensor to cells.

3.1. Preliminary Definitions

Assume that we have a robot moving in a planar scenario whose pose can be described by $\mathbf{x} = [x \quad y \quad \phi]^\top$, where (x, y) are 2D coordinates and ϕ is the robot heading. The real pose at each instant of time k is unknown and can be estimated through Bayesian filtering (Dellaert et al. 1999). Within a RBPF (Doucet et al. 2000a) we estimate the probability density of the robot path through a particle filter that implements the recursive Bayesian formula:

$$\begin{aligned} & p(\mathbf{x}_k \mid u_{1:k}, z_{1:k}) \\ \propto \quad & p(z_k \mid \mathbf{x}_k, u_{1:k}, z_{1:k-1}) \int p(\mathbf{x}_k \mid \mathbf{x}_{k-1}, u_n) \\ \times \quad & p(\mathbf{x}_k \mid u_{1:k-1}, z_{1:k-1})\, d\mathbf{x}_{k-1}, \end{aligned} \tag{1}$$

where for any time step k, z_k are sensor observations and u_k are robot actions (typically incremental displacements given by odometry (Dellaert et al. 1999) or range scan matching (Hahnel et al. 2003; Stachniss et al. 2005b)). For clarity in the rest of the paper, we represent $z_{1:k}$ and $u_{1:k}$ by $d_{1:k}$.

In this paper we consider maps represented through occupancy grids, which have been a popular representation in the robotics community during the last 20 years (Moravec and Elfes 1985; Moravec 1988; Elfes 1989; Thrun 2003; Stachniss et al. 2004; Grisetti et al. 2007b). An occupancy grid map m is a discrete random field where we store the occupancy probability for each cell, which we denote by $p(m_{xy})$ for a cell with indexes $\langle x, y \rangle$. If no prior information about obstacles is available, the occupancy probability for all cells can be initially set to 0.5: each cell can be occupied or free with the same probability. A grid map is updated by integrating sensor measurements through the so-called *inverse sensor model* (Thrun et al. 2005),

$$p(m_{xy} \mid z_k), \tag{2}$$

that is, the likelihood of the cell m_{xy} being occupied, conditioned to a given observation z_k. Fusing the current observation with the previous contents of the map can be done on a Bayesian basis by using the following iterative expression:

$$\begin{aligned} & p(m_{xy} \mid z_{1:k}) \\ = \quad & \left(1 + \frac{1 - p(m_{xy} \mid z_{1:k-1})}{p(m_{xy} \mid z_{1:k-1})} \frac{1 - p(m_{xy} \mid z_k)}{p(m_{xy} \mid z_k)}\right)^{-1}, \end{aligned} \tag{3}$$

which can be easily derived from the log-odds representation of the update process (Moravec 1988) if we assume an initial occupancy likelihood of 0.5 for all cells. Essentially, (3) increases the certainty in the occupancy/freeness of a given cell if subsequent observations confirm the current belief. The only density required to iterate (3) is the inverse sensor model in (2). For the common case of a laser range scanner we can consider a function such as that depicted in Figure 3.

At this point we have defined stochastic representations for both the robot path and the map (in (1) and (3), respectively), thus we could compute their entropy values separately to measure their uncertainty. In turn, we propose here to build a new map, the EM, aimed at revealing inconsistencies between the maps hypotheses in a RBPF, and aimed at measuring the overall uncertainty of the filter.

3.2. Definition of the EM

The EM is defined as the mathematical expectation of the probability distribution of the map taken over all of the possible paths, given by the posterior $p(\mathbf{x}_{1:k} \mid d_{1:k})$, that is, the result of marginalizing out the robot path from the SLAM posterior:

$$p(\text{EM} \mid d_{1:k}) \doteq \text{E}_{\mathbf{x}_{1:k}}[p(m \mid \mathbf{x}_{1:k}, d_{1:k})]. \tag{4}$$

In the context of SLAM we know the posterior of the robot path, given by (1). Thus, the definition of the EM above can be rewritten as

$$p(\mathrm{EM} \mid d_{1:k}) = \int p(m \mid \mathbf{x}_{1:k}, d_{1:k})\, p(\mathbf{x}_{1:k} \mid d_{1:k})\, d\mathbf{x}_{1:k}. \quad (5)$$

In the concrete case of RBPF-based SLAM and occupancy grid maps, the EM is a grid map where the occupancy of each cell EM_{xy} can be obtained from the contributions of all of the map hypotheses in the RBPF, each one associated to a path hypothesis:

$$p(\mathrm{EM}_{xy} \mid d_{1:k}) \approx \sum_{i=1}^{M} \omega_k^{[i]}\, p(m_{xy} \mid x_{1:k}^{[i]}, d_{1:k}), \quad (6)$$

where M is the number of particles in the filter and the $\omega_k^{[i]}$ are their associated importance weights. The intuition behind (4)–(6) is that by contrasting the occupancy values of cells m_{xy} at the same location (the same $\langle x, y \rangle$ indexes) but from maps associated to different particles, we can test the coherence between hypotheses, which cannot be measured directly by other methods such as the joint entropy (as shown in Section 5.2). The resulting map can be visualized as an image where sharp areas indicate certain information whereas blurred areas involve uncertain regions, i.e. there are contradictory hypotheses about their content. This is illustrated in Figure 4(c), where we obtain an EM with uncertain occupancy values (close to 0.5) for cells with contradictory content in the individual maps. Another example for real data has been shown in Figure 1(b).

In the next section we define information metrics capable of measuring the information of an EM. As expected, those measures assign a higher amount of information to "sharp" maps than to "blurred" maps. If we think of the EM in an active exploration framework, it is clear that desirable actions are those that lead to information being gained in the EM, caused either by the exploration of new areas or by the closure of a long loop.

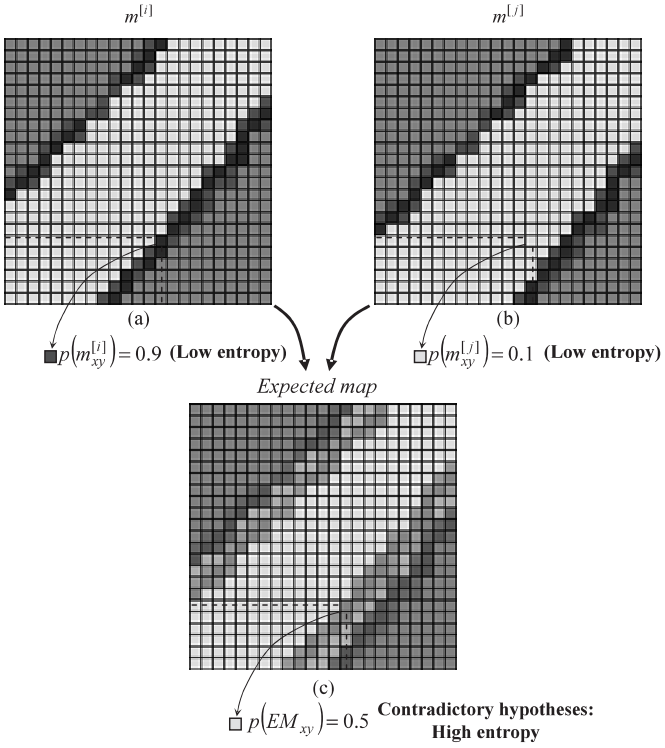

Fig. 4. (a) and (b) Map hypotheses according to a pair of particles i and j, respectively. (c) Their combination into an EM reveals contradictory values for some cells (such as those highlighted in the figures) by means of a corresponding cell in the EM with a highly uncertain occupancy likelihood.

4. Information Metrics for Grid Maps

In the following we define two metrics which measure the information in a given occupancy grid map (and, in particular, in an EM). Information theory establishes that the information associated with a random variable is related to its entropy (Cover and Thomas 1991). As each grid cell is a discrete random variable with two possible outcomes, i.e. a Bernoulli distribution, its entropy is given by

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

with $\bar{p}(m_{xy}) = 1 - p(m_{xy})$. Note that the maximum attainable entropy is given for $p(m_{xy}) = 0.5$, that is, for unobserved cells. Assuming statistical independency between cells, we can compute the entropy of the whole map m as

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

These equations for computing the entropy of a grid map have been widely used as a measure of the information in the maps (Bourgault et al. 2002; Stachniss et al. 2005a). However, in practice it exhibits the following drawbacks:

- Its absolute value depends on the grid extent (the rectangular limits of the map) instead of the actual observed area. Note that (8) implies that all of the unobserved cells in a map contribute to the entropy with their maximum value of uncertainty.
- It also depends on the grid resolution, since this parameter settles (along with the map extent) 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. We must remark that corrective scale factors have been proposed elsewhere to alleviate this problem (Stachniss 2006).

In order to overcome these drawbacks, we define here the information I of a map m as the following entropy-based measure:

$$\begin{aligned} I(m) &= \sum_{\forall x,y} I(m_{xy}) \quad \text{(bits)}, \\ I(m_{xy}) &= 1 - H(m_{xy}), \end{aligned} \tag{9}$$

where for convenience the entropy $H(\cdot)$ is computed using base-two logarithms. As a result we obtain a natural measure of information in units of bits. Notably, the maximum information value (1 bit) is given to a certain occupied or free cell, while the minimum value (0 bits) is associated with any unobserved cell. It is evident, therefore, that the dependency of the entropy on the grid extent is avoided using this definition of information: the limits of the map become irrelevant because all of the unobserved cells now contribute null information. Thus, the information of a map $I(m)$ is a more practical quantifier of the uncertainty of a map than the direct application of the entropy.

In this paper we are also interested in the *certainty* of the content of a map, not only in its absolute amount of information. To effectively reflect the certainty of a map m we introduce the MI, or $\bar{I}(m)$, which is defined as

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

where N_{obs} represents the number of observed cells in the map, i.e. all of those cells with occupancy likelihood different from 0.5. Note that this measure delivers bounded values in the range [0, 1], whereas entropy values can increase without bound.

Consider the following example which illustrates a key feature of the MI. Two observations are captured from different poses within the same environment, which separately give rise to the maps in Figures 5(a) and (b). Alternatively, Figure 5(c) shows an occupancy grid where both observations are fused (by means of (3)) with the correct alignment. Here each observation confirms the other, thus occupancy values of cells are closer to 0 and 1 than in the previous maps made from a single observation. Consequently, the MI of this map is greater than before because we have more certain information. This behavior follows from the properties of the information of a map as defined in (9), whose maximum value is obtained for occupancy values of 0 and 1. To illustrate an opposing situation, please note how both observations are misaligned in Figure 5(d), which is given a lower MI value.

To summarize, we enumerate next the properties of our information metrics that set them apart from the direct application of entropy to grid maps:

(i) An empty map (containing only unobserved cells) has null information and MI values.

(ii) Both measures are independent of the grid map rectangular limits, because unobserved cells do not contribute to the information in the map.

(iii) The MI is mostly independent of the grid resolution for practical cell sizes. This is shown in the maps of Figures 5(e) and (g), whose MI increase as we consider higher resolutions. However, the MI asymptotically tends towards a maximum value, as can be appreciated in Figure 5(h). This presents the remarkable difference in performance with the entropy, which in this case tends to infinity. The asymptotical behavior of MI depends on the inverse sensor model, the specific environment being mapped, and other factors. In Appendix A we show how a closed-form expression can be derived for a given synthetic environment.

(iv) The better the alignment between observations in the map, the higher the values obtained for the MI, as can be observed in Figures 5(c) and (d). The intuitive idea behind this property is that well-aligned observations make the occupancy of cells to tend towards either 0 or 1, which correspond to maximum information values. This property is the key for the distinctive behavior of the EMMI uncertainty estimator when closing a loop, as shown in the following.

We discuss next the experiment summarized in Figure 6, which is aimed at showing how we obtain high MI values for maps built from well-aligned observations (property (iv)). Simulation has been chosen here because in that case we know the real robot pose and we can change the sensor parameters freely. We have simulated several sets of observations (laser range scans) from the surroundings of the circled positions marked in Figure 6(a), according to some given errors in positioning (x, y) and orientation (θ). The observations are also corrupted with an additive Gaussian noise of standard deviation σ. Next we compute the MI of the grid map built from these observations following (3). The average results are plotted separately in Figures 6(b)–(e) for errors in orientation and position, and for different levels of sensor noise. As expected, the highest value of the MI is obtained for the correctly aligned observations. We can also remark how more precise observations (with lower error σ) lead to more selective MI values.

5. Comparison with Other Uncertainty Measurements

In the following we compare our proposed uncertainty measures, namely the EMI and EMMI, with other entropy-based methods, such as the entropy of the robot path (Burgard et al. 1997; Roy et al. 1999; Vlassis et al. 1999), the effective sample size, N_{eff} (Liu 1996), and the path–map joint entropy (Stachniss et al. 2005a). We provide here a theoretical discussion

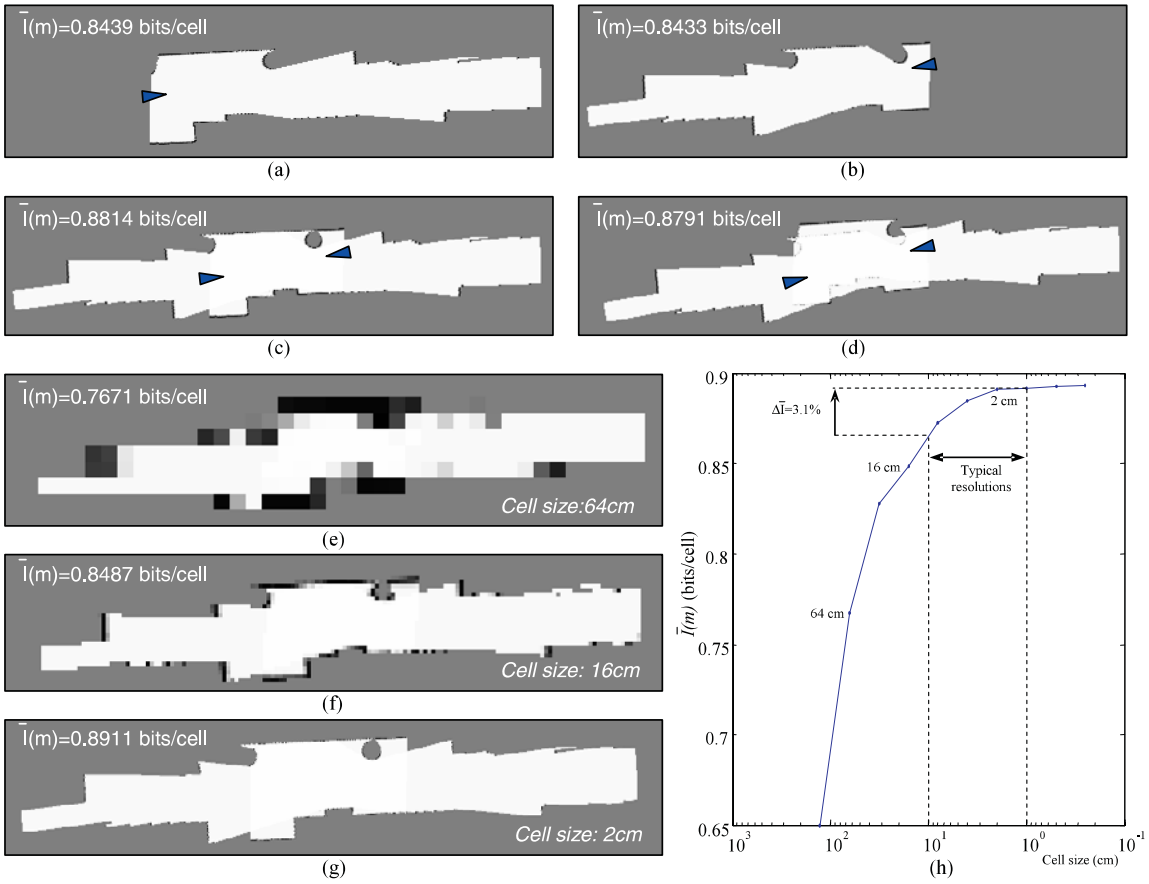

Fig. 5. Examples illustrating the main properties of the MI ($\bar{I}$). (a) and (b) Two observations and the MI value of the maps built from them separately. (c) When both are aligned and fused in the same grid, the information value increases. (d) Inconsistencies due to poor robot localization decrease the quality of the fused map, which is confirmed by a lower MI value. (e) and (g) The same map is shown for different cell resolutions, along with their MI values. (h) The MI values obtained for this environment are plotted for a wide range of resolutions, from 128 to 0.25 cm. Observe the small change in the MI value when varying the resolution of a map from 1 to 10 cm, which coincides with commonly used resolutions. Thus, in practice, the MI can be regarded as resolution independent for maps with cell size smaller than 10 cm.

about their complexities and their expected behaviors in typical mapping situations, while experimental results that support our reasonings are shown in the next section.

5.1. Expected Behaviors for the Uncertainty Measures

Before comparing the behavior of the different measures, it is convenient to consider the expression of the joint entropy (Stachniss et al. 2005a), the method most related to ours, to better understand its properties:

$$\begin{aligned} & H(\mathbf{x}_{1:k}, m \mid d_{1:k}) \\ = \; & H(\mathbf{x}_{1:k} \mid d_{1:k}) + \int p(\mathbf{x}_{1:k} \mid d_{1:k}) H(m \mid \mathbf{x}_{1:k}, d_{1:k}) \\ \approx \; & H(\mathbf{x}_{1:k} \mid d_{1:k}) + \sum_{i=1}^{M} \omega_k^{[i]} H(m \mid x_{1:k}^{[i]}, d_{1:k}). \end{aligned} \tag{11}$$

It is clear that this measure is a composition of the entropy of the robot path and the average entropy of individual maps $p(m \mid x_{1:k}^{[i]}, d_{1:k})$ weighted by $\omega_k^{[i]}$ for each particle i. Three fundamental drawbacks can be derived from (11):

(i) The term that considers the entropy of the maps (the integral in (11)) dominates the overall value, therefore hiding the contribution of the path uncertainty (the first term in that expression). The reason is that the former is typically many orders of magnitude higher than the entropy of the path. For instance, we have found in our experiments a typical ratio between both terms of approximately 10^5.

(ii) Only the content of individual maps determines the result, independently of their mutual consistency: this method is unable to detect inconsistencies between par-

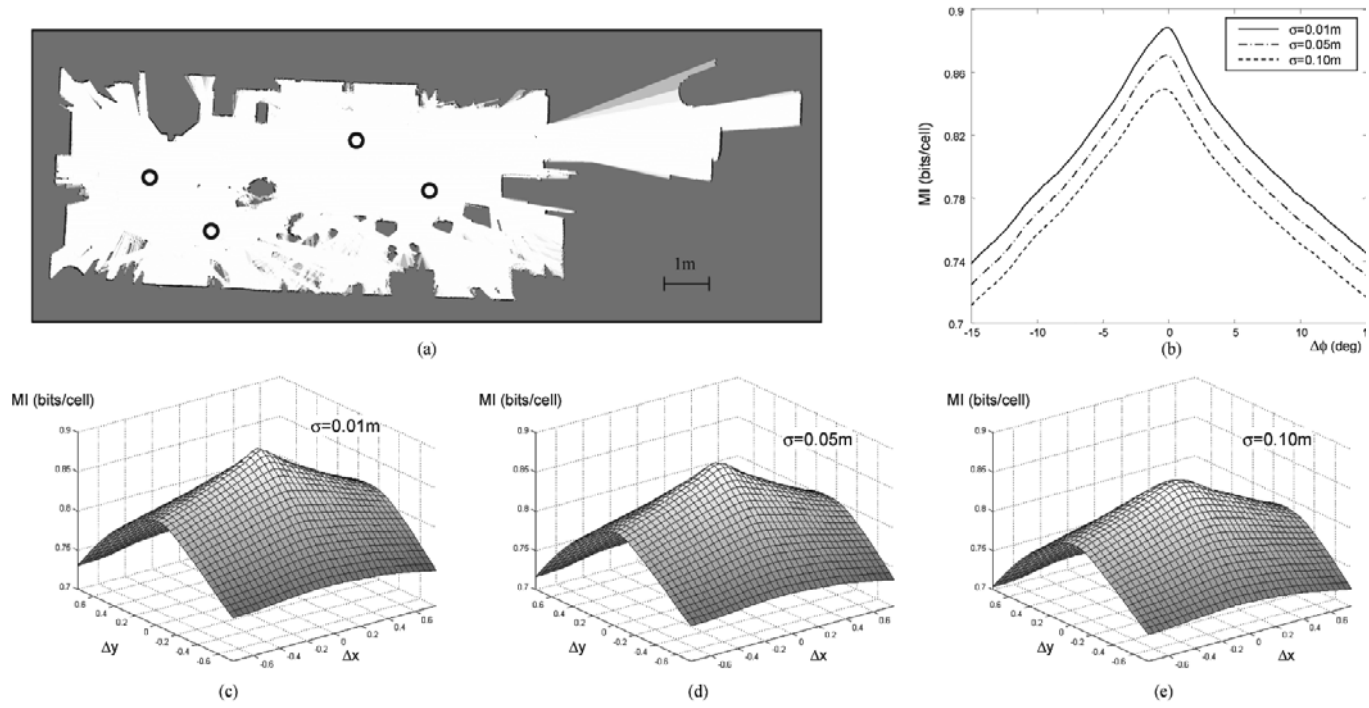

Fig. 6. (a) The simulated environment used to test the MI response against sensor noise and localization errors. Observations have been simulated from the circled positions, a map has been generated from them, and its MI has been computed. (b) The results for orientation errors only and different levels of sensor noise. (c)–(e) The response of the MI against errors in positioning (with the correct orientation). The higher selectivity for less noisy measurements is clearly visible. Anyway, the maximum value is obtained for the correct position and orientation, thus the results show the ability of the MI to measure the consistency between different observations.

ticles. In other words, the joint entropy is largely determined by the number of observed cells in maps, being mostly independent of whether individual maps contradict each other.

(iii) It is not clear how to compute the entropy of the robot path (the first term in (11)) in a RBPF. We consider here a good approximation of this value based on fitting Gaussians to the robot pose at each time step and averaging their entropy along the whole path, as proposed by Roy et al. (1999) and Stachniss et al. (2005a). However, we show in the following that this value can become undefined after closing a loop.

We examine next the theoretical behavior of the four previously mentioned uncertainty measures for two distinctive situations that can occur while performing mapping or exploration with RBPFs. The purpose of this comparison is to show how our estimators, EMI and EMMI, represent a more reliable measure than the others for selecting robot actions and detecting loop-closures, respectively.

5.1.1. Exploration of New Areas

In this case, particles tend to spread in space while their weights remain practically constant. As a result, the N_{eff} of the RBPF remains practically constant, while the entropy of the path increases owing to the higher uncertainty in localization. Note that this entropy is one of the components of the joint entropy, but in the case of exploring new areas the information added to grid maps more than compensates for the increase in uncertainty in the robot path, thus the overall value of the joint entropy increases while exploring. Regarding our uncertainty measures, here the EMI increases, but the number of observed cells grows as well. However, the increase in the robot pose uncertainty continuously introduces small incoherences between individual maps, which is reflected by decreasing EMMI values.

5.1.2. Loop Closure

In SLAM, coming back to an already known place following a previously unknown path is a special case of revisit called loop closure. If the loop is long enough, typically only a few

particles will be close to the actual path. This will result in: (i) only a few particles with non-negligible weights, (ii) particles are resampled, and (iii) the pose uncertainty is largely reduced (Arulampalam et al. 2002; Grisetti et al. 2007b). It is clear here that both N_{eff} and the robot path entropy fall drastically. Additionally, since this entropy is also a term of the joint entropy (and the other term for the maps remains practically constant in this situation), we can conclude that the joint entropy also falls drastically when closing a loop.

However, in practice both of the entropies above break down owing to the particle approximation of the SLAM posterior. The reason is that typically only one hypothesis survives the resampling after closing a long loop (for example, refer to Stachniss et al. (2004)), thus the covariance of the Gaussians fitted to the robot path collapses to zero. As the entropy of a d-dimensional Gaussian $x \sim \mathcal{N}(\mu, \Sigma)$ is given by

$$H(x) = \tfrac{1}{2} \log((2\pi e)^d |\Sigma|) \tag{12}$$

this would imply a negative infinity entropy. We believe that one solution to this numerical problem is to impose a limit on the minimum covariance of the robot pose (although this heuristic value then becomes a determinant of the behavior of the joint entropy for exploration). For instance, a standard deviation of 1 mm in x and y has an associated entropy of -10.98. However, we must highlight the small range of values for the path entropy in comparison with the second term of the joint entropy (the uncertainty in maps), which means that, in practice, the content of maps determines the overall value while disregarding the uncertainty in the path.

Consider the following numerical example which illustrates the problem of using the joint entropy for choosing actions in exploration. A robot faces two potential actions, which are evaluated trough the joint entropy. The first action closes a long loop and reduces the robot pose uncertainty in x and y from 1 m to 1 mm (considering these values as the standard deviations), whereas we predict that the second action will allow us to explore a new area of 1 m^2 (100 grid cells for a 0.10 m cell size). Straightforward calculations give us an expected reduction of the joint entropy of 13.82 and 69.3 for the first and second actions, respectively. Obviously, closing the long loop is a much more desirable action for the robot here than the exploration of a small area of new terrain (Stachniss et al. 2004; Sim and Roy 2005), thus the joint entropy would be not a good choice for active exploration in this situation.

In contrast, the EMI of the RBPF after closing a loop increases owing to the elimination of incoherent hypotheses (which generate more "blurred" EMs), the same reason why the EMMI increases. The particle depletion problem does not affect our measures because the different path hypotheses are taken into account through the EM, which is always well defined.

5.2. Consistency Between Map Hypotheses

We now discuss a key difference between how the EMMI and the joint path–map entropy detect inconsistencies between the individual maps from each particle in a RBPF. In the joint entropy, the entropy of the map contents $H(m)$ contributes to the total value averaged by the weights of the associated particles. In contrast, the EMMI considers the entropy of the cells in the EM, which, in turn, are the weighted average of the maps from each particle. Consequently, the 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, Figures 7(a) and (b) shows a simple example of how the joint entropy fails to capture inconsistencies. Here, the graphs show the entropy terms involved in each method for two equally probable particles (this simplification has been taken to enable the visualization of the results in three dimensions). The key observation is that inconsistencies such as 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, that is, high entropy values (see Figure 7(b)). In turn, similar situations are assigned a low entropy value (high certainty) by the joint entropy, as can be observed in Figure 7(a).

5.3. Computational and Storage Complexity

Regarding the computational time complexity of all of these uncertainty measures, the most efficient are the effective sample size N_{eff}, with $\mathrm{O}(M)$, and the entropy of the robot path, which involves $\mathrm{O}(ML)$ operations, where M and L are the number of particles and the length of the path, respectively. On the other hand, the joint entropy has a complexity of $\mathrm{O}(M(N + L))$, where N is the number of cells in the grid. Clearly, N will be the dominant quantity in practice, thus this complexity will tend towards $\mathrm{O}(MN)$. However, because each observation will modify only a limited area of the grid maps, we can reduce the computation complexity of the joint entropy to $\mathrm{O}(M)$. Unfortunately, the construction of the EM (required for the evaluation of the EMI and EMMI measures) depends on the weights of the particle filter, hence it cannot be simplified and the complexity of building the EM becomes $\mathrm{O}(MN)$. This is the price for testing the consistency between difference map hypotheses (refer to the previous section).

Concerning the space complexities, our methods require the computation of one additional map, the EM, in addition to the M maps associated with the particles in the RBPF, whereas the joint entropy does not require this supplementary storage. However, this cost is not significant because it implies keeping $M + 1$ maps instead of M, where M is the number of particles.

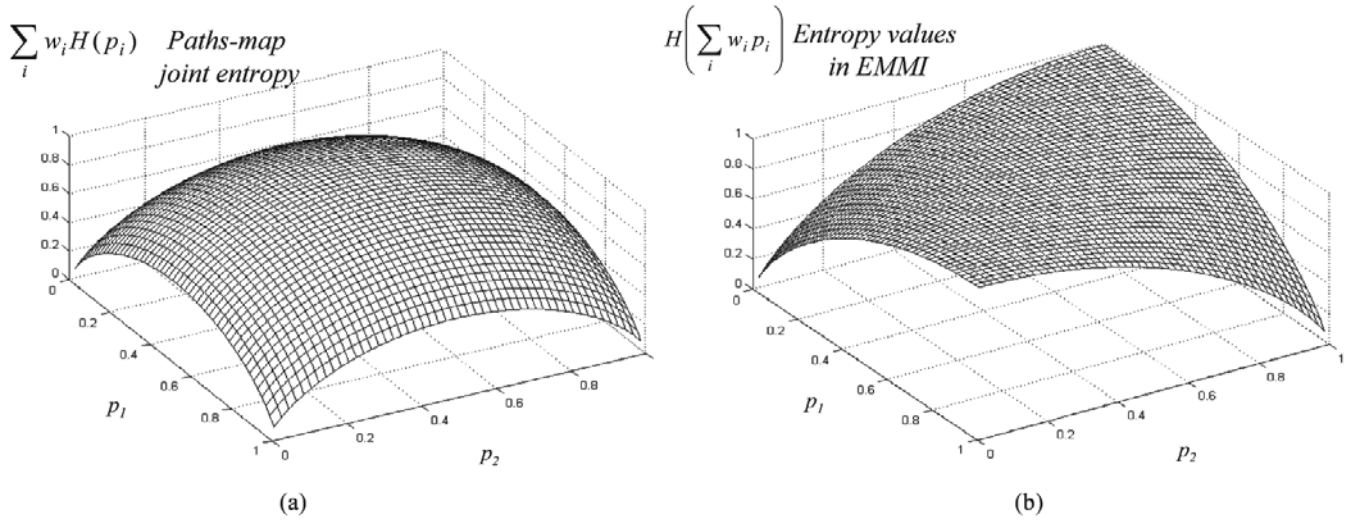

Fig. 7. Entropy terms involved in (a) the path–map joint entropy and (b) the EMMI are plotted for the case of only two particles, as an example of how each approach deals with different hypotheses. It is clear that the first approach considers contradictory values as certain (low entropy), such as $p_1 = 0$ (free cell) and $p_2 = 1$ (occupied cell), whereas EMMI only gives low entropy values when both likelihoods are similar (consistent).

6. Experimental Results

In this section we provide an experimental validation of our methods through comparisons with other common uncertainty estimators. Firstly, we perform a statistical analysis of the EMMI measure in order to check its ability to detect the closure of loops through the changes in the EM of the RBPF. Next we show how EMI can be applied to active exploration and compare its performance with that of the joint path–map entropy.

6.1. Statistical Characterization of EMMI

Owing to the stochastic nature of particle filters, it is desirable to carry out many realizations of each RBPF simulation to obtain statistically significant results. This is the reason why in this experiment we have repeated the mapping process 20 times (each realization takes 20 minutes with a 3.2 GHz Pentium 4), hence enabling a more convenient comparison by the mean values and variances of the different uncertainty measures. The experimental setup consists of an optimized RBPF (Grisetti et al. 2007a) which consumes data gathered by the robotic wheelchair SENA (Fernandez-Madrigal 2004; Gonzalez et al. 2006) in the streets surrounding a building in our campus (see Figure 8). Sensory data consist of odometry readings and scans from a SICK laser range finder.

In this experiment the robot completes two laps around the building, which is a total distance traveled of 550 m at an average speed of 1.5 m/s. The most critical point is when it closes the loop for the first time. Figure 8(a)–(d) represent the evolution of the EM at some instants of time during the mapping. The "blurred" aspect of the EM is more perceptible as the robot moves farther along the loop, until it definitively closes it. This closing starts approximately at time step 150, where the uncertainty owing to the dispersion of particles starts to decrease. After that, the path estimation and the map remain accurate while the robot travels over the same path again. This illustrates the fact that the uncertainty continuously increases while the robot explores unknown areas and that it rapidly decreases when the loop is closed. To measure these changes we have computed both the EMMI and the joint entropy (theoretically compared in Section 5). The entropy of the path, also discussed there, has not been considered in this statistical comparison because it does not reflect the uncertainty in maps. In addition, the effective sample size N_{eff} and the 3σ (99.7%) confidence area occupied by the particles in the space have also been computed for illustration purposes, although they are unfounded uncertainty estimators. Results are plotted in Figures 8(e)–(h), where mean values from the 20 experiments are shown as full curves and $\pm 2\sigma$ confidence intervals appear as dashed curves. Note that these intervals are for fitted Gaussian distributions, hence they can exceed the valid range of the variables, i.e. the upper confidence limit in Figure 8(g).

The evolution of the uncertainty estimators can be interpreted as follows. In the case of the EMMI, it decreases (less certainty) while the robot travels along the loop for the first time. After the loop closure, this estimator restores its high value (more certainty). This behavior can be observed in Figure 8(e) in the first gradual decrease and the posterior rise be-

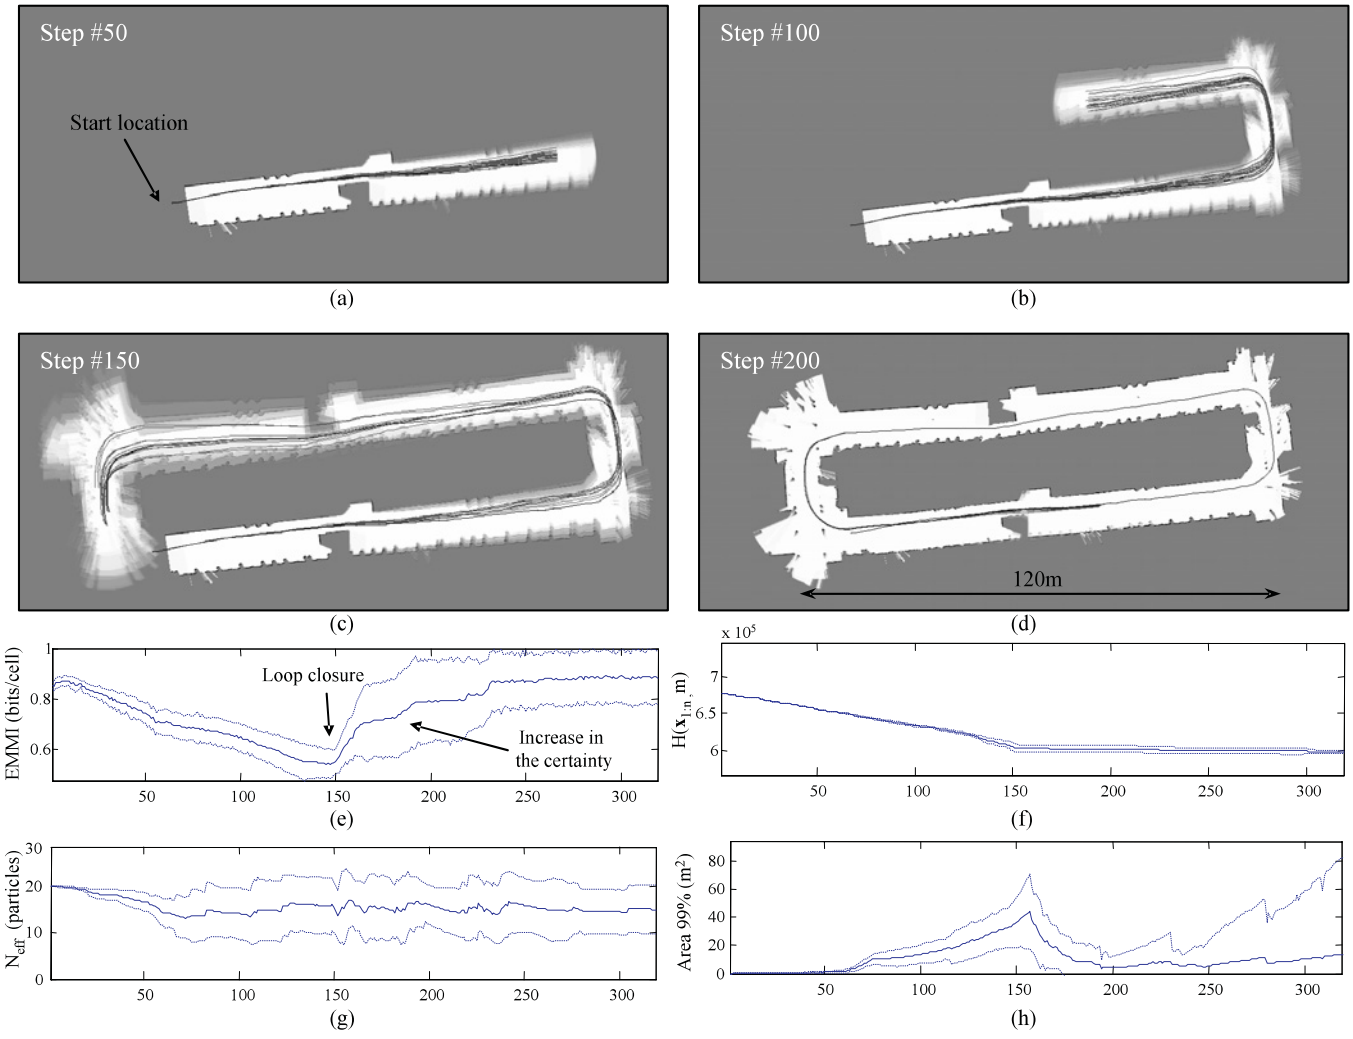

Fig. 8. (a)–(d) The EM for some steps of the mapping. The uncertainty estimators are (e) the EMMI, (f) the joint entropy, (g) the effective sample size, and (h) the area occupied by the particles. The ability of the EMMI to better reflect the actual uncertainty of the RBPF than the others is shown. The experiment has been repeated 20 times and the mean values for all of these estimators are plotted as full curves together with their 95.4% confidence intervals ($\pm 2\sigma$ using approximate Gaussian fits) as dashed curves.

yond time step 150. Note how the EMMI remains practically constant during the rest of the experiment. On the other hand, the joint entropy has been shown to be the estimator with the smallest variance, i.e. its values do not vary significantly between different runs. This is due to the dominance of the maps entropy in the estimator, which only depends on the number of observed cells, being (practically) independent of inconsistencies between hypotheses. Consequently, from the results given by this measure, shown in Figure 8(f), we can only know that, from step 150 to the end of the experiment, a small number of cells have been updated in the maps. Note that if we were to use this estimator in active exploration to decide robot actions, we would be hardly able to distinguish between closing a loop or any other movement with a similar outcome (in practice, any other action that does not involve exploration of a new area, e.g. going back along the traversed path, remaining still, etc.). This contrasts with the competent performance of the EMMI in this respect. Regarding the results for the effective sample size (N_{eff}), it can occasionally reflect a loop closure by means of restoring its maximum value owing to an associated particle resampling. However, the exact moment at which a resample occurs depends on many implementation parameters. Moreover, a resample can occur many times after a loop closure or even while exploring new areas, as previously discussed. All of these facts are reflected in the large variance observed in Figure 8(g). Moreover, its mean value is not correlated at all with the actual mapping uncertainty. The last computed estimator is the area occupied by particles in the space (Stachniss et al. 2005b), which we compute as the area of the three ellipsoids resulting from approximating particles with a Gaussian distri-

bution. No only is this estimator not mathematically grounded, it also does not take into account the map contents. This implies that, for example, it cannot distinguish between dispersed particles or a few separated modes where particles are concentrated. Furthermore, this estimator shows a high variance (see Figure 8(h)) while the robot traverses the loop for the second time, whereas the actual uncertainty is reduced following the first loop closure. In these experiments, the time required to compute the EMMI of the RBPF with a 3.2 GHz Pentium 4 is about 127 ms for a straightforward implementation in C++ and 49 ms for an optimized assembly version that exploits the highly parallel structure underlying the EM computation (averaging maps from all of the particles, the most time-consuming task in EMMI) by the Streaming SIMD Extension 2 (SSE2) instruction set (Intel 1999). Our implementation of the joint entropy takes 210 ms for the same particle filter, although more optimal algorithms could be used for its computation.

6.2. Exploration with EMI

To test the performance of EMI in selecting actions in active robot exploration, we have chosen the information gain-guided exploration framework proposed by Stachniss et al. (2005a). In short, the approach can be described as follows. At each time step we generate a set of potential targets around the robot, each a potential action. These points are generated using the most likely map from the RBPF and taking into account the feasibility of the path given the size of the robot and the potential existence of obstacles along the path. In our implementation we employ a simple path planner that assumes a circular robot and extracts collision-free paths according to a given occupancy grid map. Once we have established those potential paths (actions), we predict the observations along each of them by means of ray tracing over the grid. The observations for each path are integrated into a copy of the RBPF and the change in the uncertainty of the RBPF determines the information gain of each potential action. Finally, the robot takes the action with the highest utility value, which is computed by subtracting a cost proportional to the length of the path from the information gain. The navigation towards successive targets is performed by an obstacle avoidance method (Blanco et al. 2006) that runs on the robot in real time and concurrently to the update of the RBPF for mapping. In our implementation, the gain in information is computed as the increment of the EMI associated with the RBPF, although the decrease of the joint entropy has been also computed for comparison purposes. Finally, we must remark that the approximation of carrying out all of the computations just for the most likely particle is acceptable as long as the robot does not transverse loops that are too long. In those cases, more particles should be considered and the utility values averaged using the particle weights.

In this experiment the robot starts in one of the corridors of an environment containing a large loop, as shown in Figure 9. For each decision, the robot predicts the EM after the execution of all of the potential actions and chooses the action with the highest utility value, as defined above. Some of the predicted EMs in this way are shown in the middle graphs in Figure 9. It can be seen how it is typically predicted that unobserved areas contain free space, which means an increase in the map information and therefore provides the motivation for exploring them. As the robot moves farther from the starting point, as in the second row in Figure 9, the predicted EMs increasingly contain "blurred" areas, owing to the increasing uncertainty in the robot pose. It is notable that the utility values obtained from the EMI and the joint entropy for the first and second decisions in Figure 9 are very similar. This means that, for actions involving the exploration of new areas, both measures effectively detect which are the most advantageous actions: those leading to the exploration of large areas.

It is interesting to remark here that, as long as the predicted future observations of the robot are random variables, so are the predicted RBPF weights and the corresponding utility values. This helps to explain some apparently strange results in the experiment, such as the assignment of quite different utility values to targets close to each other (e.g. see targets #6 and #7 for the decision #12 or targets #7 and #8 for decision #18). Although we have verified that this is not a problem for exploring relatively small scenarios, it could be solved by predicting future observations a number of times and averaging the resulting utility values.

Consider now the third decision, the bottom row in Figure 9. Here there is a large amount of uncertainty in the robot path after traversing the whole loop, and some of the actions include definitively closing the loop (see target #7, for example), whereas others imply entering new areas to continue the exploration (such as target #12). Interestingly, in this case we obtain quite different values from the EMI and the joint entropy, as can be clearly observed in the bottom-right graphs in Figure 9. While the EMI assigns the highest utility to the action that definitively closes the loop (note the "sharp" EM which is predicted for action #7), the joint entropy assigns the highest values to those paths leading to unknown areas. This is because of the dominance of the uncertainty in maps in this measure, as discussed throughout this paper, which hides the potential reduction of uncertainty in the robot path derived from the loop closure. Therefore, we can conclude that EMI performs similarly to the joint entropy for exploration, with the additional advantage that it detects much more clearly the possibilities for closing loops, which reduces the overall uncertainty of the SLAM posterior in such a way that the rest of the exploration will be easier than if loops are not closed.

7. Conclusions

In this paper we have motivated the need for measuring the uncertainty of a RBPF solution to SLAM, which is required for tasks such as detecting loop closure or exploration. We have

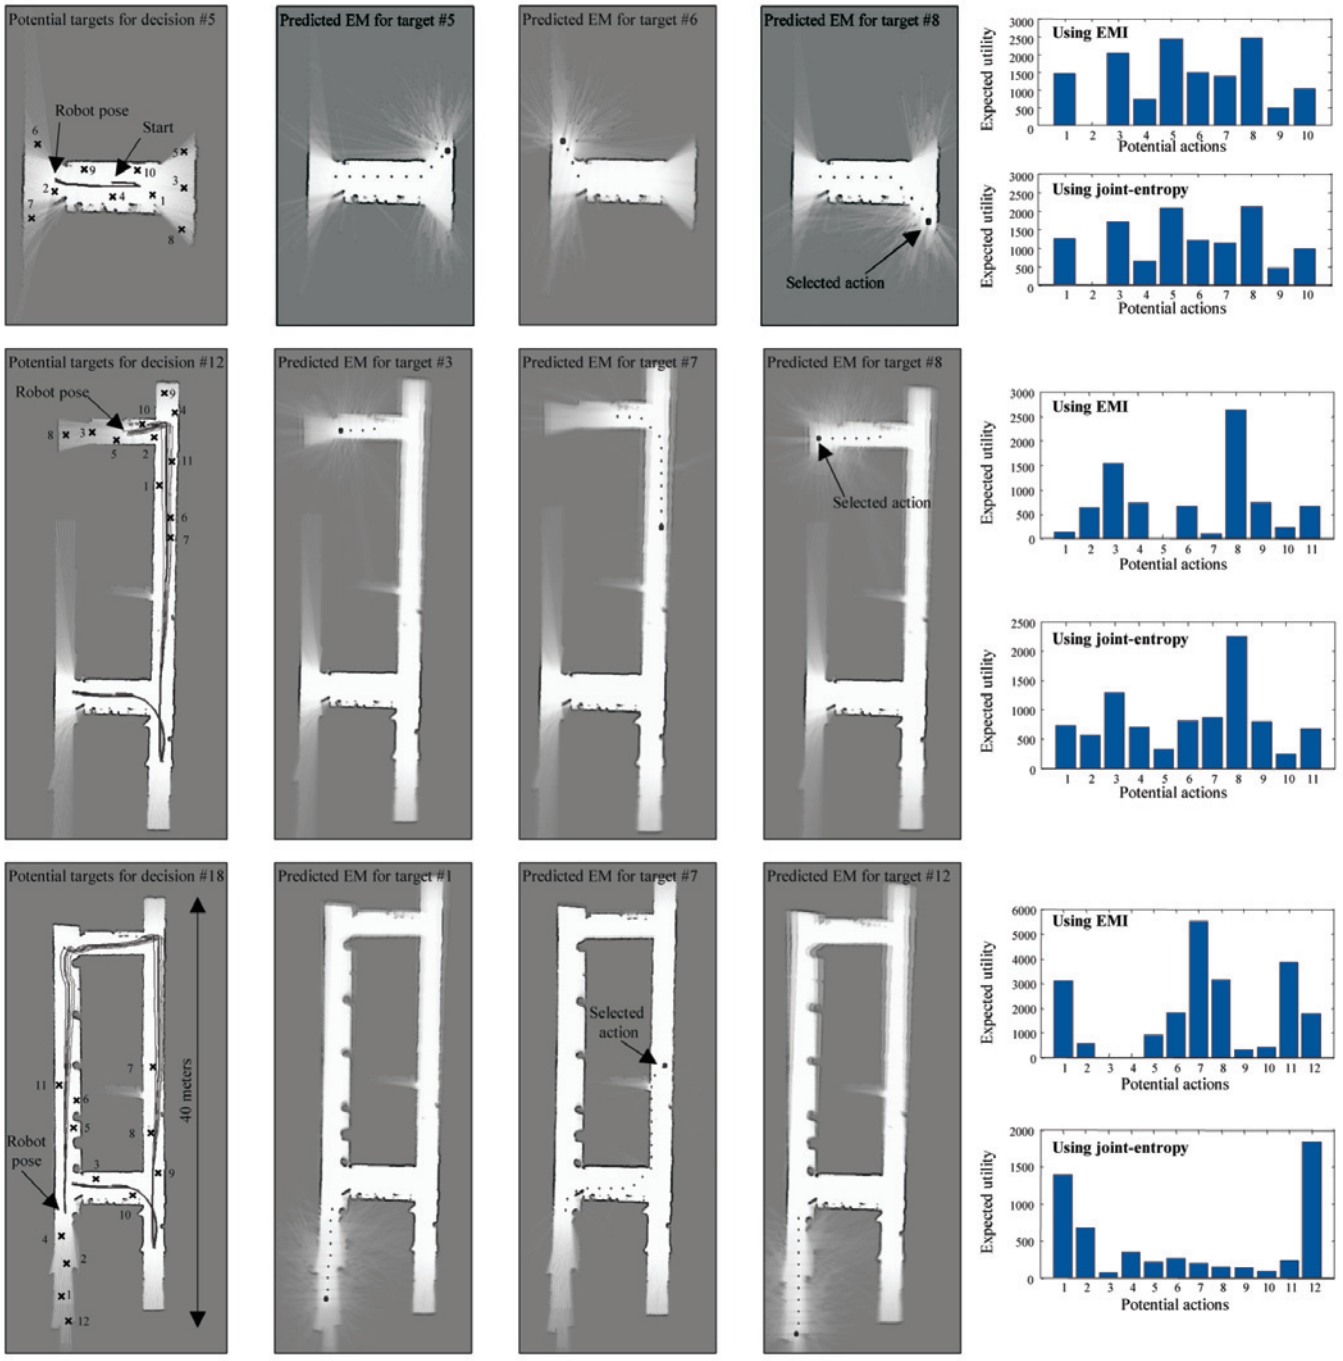

Fig. 9. Three snapshots of certain moments during an exploration when the robot has to decide between a set of potential movement actions, marked on the images on the left. The utility assigned to each of those actions is plotted both for our EMI measure and the joint entropy in the graphs on the right, whereas some of the EMs predicted while evaluating the actions are shown in the middle. As discussed in the text, the third decision clearly reflects that the actions selected by the EMI are more advantageous.

highlighted the drawbacks of previously employed measures such as the joint path–map entropy. Instead, we propose to construct the EM, a new map that condenses all of the uncertainty in the RBPF, both for the map and the robot path. By computing the information of this map we define the EMI of a RBPF, which has some advantages in comparison with other

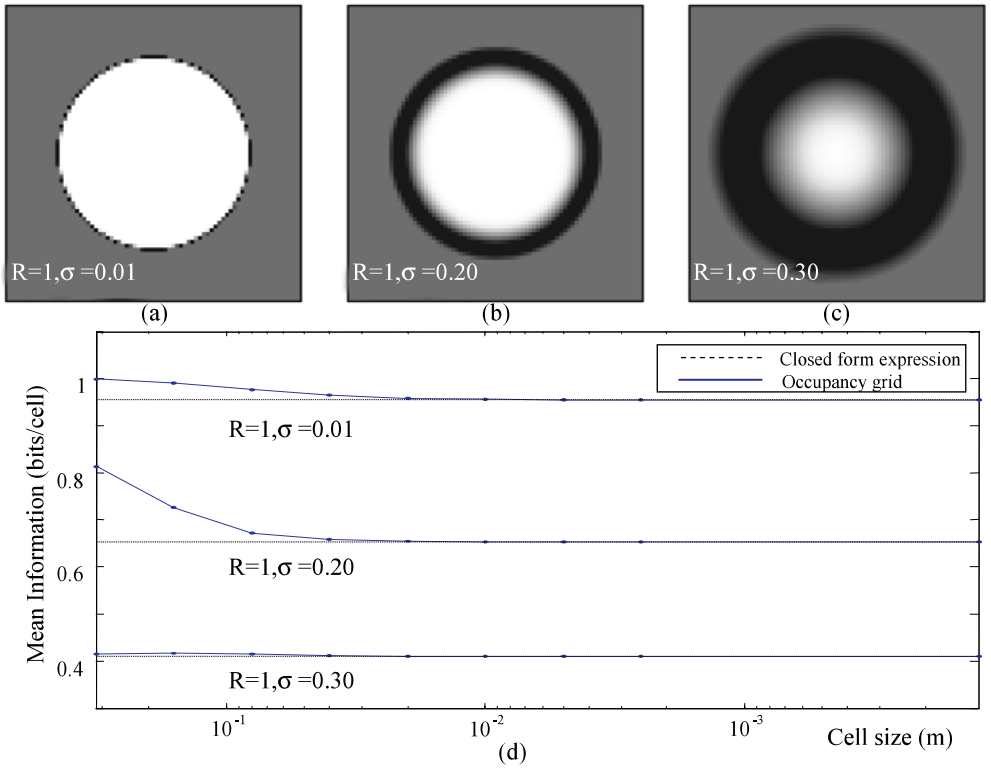

Fig. 10. The MI values for the synthetic scenario discussed in Appendix A are shown as continuous plots, together with the theoretical limit (derived in this paper) as a dashed line. It can be seen how the MI tends toward the computed limit as the resolution becomes more fine-grained.

measures for choosing actions in active exploration. Moreover, we have also defined the EMMI, which can be applied to the problem of detecting when to stop actively following an already traversed loop. *A priori* we do not know how long the robot must follow a known loop for the incorrect particles to be discarded. To detect when to stop the loop closing behavior, existing approaches either use the area covered by the particles or the effective sample size (Stachniss et al. 2004). None of these methods consider the map contents. We have presented experimental results that demonstrate that EMMI better fits this purpose than the other methods.

Acknowledgment

This work was partly supported by the Spanish Government under research contract DPI2005-01391.

Appendix A: Derivation of an Expression for the Maximum MI Value in a Synthetic Environment

As the resolution of an occupancy grid increases, its MI tends towards a maximum value. Unfortunately, a closed-form expression for this bound cannot be derived in the general case because it depends on the inverse sensor model, the number of observations that are merged into the map, and on the environment. For a particular case, however, we can derive such an expression to demonstrate the convergence of the MI and its high independence of the grid resolution.

Assume that we have a robot equipped with a 360° field-of-view radial range sensor, standing in the center of a circular environment of radius R, and making only one observation (z_1). Owing to the circular symmetry we consider the circular grid map $p(m_\rho)$ instead of the classical $p(m_{xy})$, where ρ is the distance from the origin (where the robot is initially) to a given point x, y. Thus, the map content is the same in all directions for the whole range of orientations $]-\pi, \pi]$ from the origin. Initially, there is no prior information about the map, thus $p(m_\rho) = 0.5$ for all of the values of ρ. If we consider a range sensor whose measurements are corrupted with additive Gaussian noise with standard deviation σ, its inverse sensor model can be defined as

$$p(m_\rho \mid z) = \begin{cases} e^{-((\rho - z)^2/2\sigma^2)} & \rho \leq z + \sigma\sqrt{\ln 4}, \\ 0.5 & \text{otherwise}, \end{cases} \quad (13)$$

where z represents the sensed range. Such a sensor model is plotted in Figure 3.

For the case of just one observation the Bayesian estimation of the map becomes simply $p(m_\rho \mid z_1) = p(m_\rho|z)$, where z is the actual measurement. If we make the cell size tend to zero, the expression for the MI in (10) becomes an integral over the map (which is no longer a discrete grid but a continuous surface)

$$\bar{I}(m) = \frac{\int\int (1 - H(m_\rho))\, d\rho\, d\theta}{\int_\theta \int_\rho \mathrm{Obs}(m_\rho)\rho\, d\rho\, d\theta}, \tag{14}$$

where the binary function $\mathrm{Obs}(m_\rho)$ takes the value of 1 for those positions which have been observed and 0 otherwise. The circular symmetry allows us to discard the integration over, because it leads to constant factors in both parts of the quotient. Taking this into account, replacing the integration limits, and solving the denominator, we end up with

$$\bar{I}(m) = \frac{\int_0^{z+\sigma\sqrt{\ln 4}} [1 - H(\exp(-((\rho - z)^2/2\sigma^2)))]\rho\, d\rho}{\frac{1}{2}(z + \sigma\sqrt{\ln 4})^2}. \tag{15}$$

The integral in the numerator has no analytical primitive, but precise values can still be obtained by numerical integration. To demonstrate the convergence of the grid-based MI measurement towards the continuous solution (15) we have generated grids for different R, σ, and resolution values. Some of the obtained grids are shown in Figures 10(a)–(c), whereas the computed MI values are plotted in Figure 10(d) together with the theoretical predictions from (15). It is clear from these graphs that MI effectively converges as the resolution increases and that this limit can be accurately predicted.

References

Arulampalam, M., Maskell, S., Gordon, N., Clapp, T., Sci, D., Organ, T. and Adelaide, S. (2002). A tutorial on particle filters for online nonlinear/non-GaussianBayesian tracking. *IEEE Transactions on Signal Processing*, **50**(2): 174–188.

Blanco, J., Gonzalez, J. and Fernandez-Madrigal, J. (2006). The Trajectory Parameter Space (TP-Space): a new space representation for non-holonomic mobile robot reactive navigation. *Proceedings of the IEEE/RSJ International Conference on Intelligent Robots and Systems*, pp. 1195–1200.

Bourgault, F., Makarenko, A., Williams, S., Grocholsky, B. and Durrant-Whyte, H. (2002). Information based adaptive robotic exploration. *Proceedings of IEEE International Symposium on Computational Intelligence in Robotics and Automation*, Vol. 1.

Burgard, W., Fox, D., and Thrun, S. (1997). Active mobile robot localization. *Proceedings of the Fourteenth International Joint Conference on Artificial Intelligence (IJCAI)*. San Mateo, CA, Morgan Kaufmann.

Burgard, W., Moors, M., Fox, D., Simmons, R. and Thrun, S. (2000). Collaborative multi-robot exploration. *Proceedings of the IEEE International Conference on Robotics and Automation*, Vol. 1, pp. 476–481.

Cover, T. and Thomas, J. (1991). *Elements of Information Theory*. New York, John Wiley & Sons.

Dellaert, F., Fox, D., Burgard, W., and Thrun, S. (1999). Monte Carlo localization for mobile robots. *Proceedings of the IEEE International Conference on Robotics and Automation*, Vol. 2.

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

Doucet, A., de Freitas, N., Murphy, K. and Russell, S. (2000a). Rao–Blackwellised particle filtering for dynamic Bayesian networks. *Proceedings of the Sixteenth Conference on Uncertainty in Artificial Intelligence*, pp. 176–183.

Doucet, A., Godsill, S. and Andrieu, C. (2000b). On sequential Monte Carlo sampling methods for Bayesian filtering. *Statistics and Computing*, **10**(3): 197–208.

Elfes, A. (1989). Using occupancy grids for mobile robot perception and navigation. *Computer*, **22**(6): 46–57.

Fernandez-Madrigal, J. (2004). Assistive navigation of a robotic wheelchair using a multihierarchical model of the environment. *Integrated Computer-Aided Engineering*, **11**(4): 309–322.

Fox, D. (2003). Adapting the sample size in particle filters through KLD-sampling. *International Journal of Robotics Research*, **22**(12): 985–1003.

Fox, D., Burgard, W. and Thrun, S. (1999). Markov localization for mobile robots in dynamic environments. *Journal of Artificial Intelligence Research*, **11**(3):391–427.

Gonzalez, J., Muñoz, A., Galindo, C., Fernandez-Madrigal, J. and Blanco, J. (2006). A description of the SENA robotic wheelchair. *Proceedings of the IEEE Mediterranean Electrotechnical Conference*, pp. 437–440.

Grisetti, G., Stachniss, C. and Burgard, W. (2007a). Improved techniques for grid mapping with Rao–Blackwellized particle filters. *IEEE Transactions on Robotics*, **23**: 34–46.

Grisetti, G., Tipaldi, G., Stachniss, C., Burgard, W. and Nardi, D. (2007b). Fast and accurate slam with Rao–Blackwellized particle filters. *Robotics and Autonomous Systems*, **55**(1): 30–38.

Hahnel, D., Burgard, W., Fox, D. and Thrun, S. (2003). An efficient FastSLAM algorithm for generating maps of large-scale cyclic environments from raw laser range measurements. *Proceedings of the IEEE/RSJ International Conference on Intelligent Robots and Systems*, Vol. 1.

Intel, Inc. (1999). Intel Architecture Software Developer's Manual. Volume 2: Instruction Set Reference. Available at: http://developer.intel.com/design/pentiumii/ manuals/243191.htm.

Julier, S. and Uhlmann, J. (1997). A new extension of the Kalman filter to nonlinear systems. *Proceedings of the International Symposium on Aerospace/Defense Sensing, Simulation and Controls*, Vol. 3.

Kalman, R. (1960). A new approach to linear filtering and prediction problems. *Journal of Basic Engineering*, **82**(1): 35–45.

Ko, J., Stewart, B., Fox, D., Konolige, K. and Limketkai, B. (2003). A practical, decision-theoretic approach to multi-robot mapping and exploration. *Proceedings of the IEEE/RSJ International Conference on Intelligent Robots and Systems*, Vol. 4, pp. 3232–3238.

Liu, J. (1996). Metropolized independent sampling with comparisons to rejection sampling and importance sampling. *Statistics and Computing*, **6**(2): 113–119.

Lu, F. and Milios, E. (1997). Globally consistent range scan alignment for environment mapping. *Autonomous Robots*, **4**(4): 333–349.

Montemerlo, M., Thrun, S., Koller, D. and Wegbreit, B. (2002). FastSLAM: A factored solution to the simultaneous localization and mapping problem. *Proceedings of the AAAI National Conference on Artificial Intelligence*, pp. 593–598.

Moravec, H. (1988). Sensor fusion in certainty grids for mobile robots. *AI Magazine*, **9**(2): 61–74.

Moravec, H. and Elfes, A. (1985). High resolution maps from wide angle sonar. *Proceedings of the IEEE International Conference on Robotics and Automation*, Vol. 2.

Porta, J. and Kröse, B. (2006). Appearance-based concurrent map building and localization. *Robotics and Autonomous Systems*, **54**(2): 159–164.

Roy, N., Burgard, W., Fox, D. and Thrun, S. (1999). Coastal navigation—mobile robot navigation with uncertainty in dynamic environments. *Proceedings of the IEEE/RSJ International Conference on Intelligent Robots and Systems*, Vol. 1.

Saez, J. and Escolano, F. (2005). Entropy minimization SLAM using stereo vision. *Proceedings of the IEEE/RSJ International Conference on Intelligent Robots and Systems*, pp. 36–43.

Sim, R. and Roy, N. (2005). Global A-optimal robot exploration in SLAM. *Proceedings of the IEEE International Conference on Robotics and Automation*, pp. 661–666.

Stachniss, C. (2006). Exploration and Mapping with Mobile Robots. *Ph.D. Thesis*, Universität Freiburg.

Stachniss, C., Grisetti, G. and Burgard, W. (2005a). Information gain-based exploration using Rao–Blackwellized particle filters. *Proceedings of Robotics: Science and Systems (RSS)*.

Stachniss, C., Grisetti, G. and Burgard, W. (2005b). Recovering particle diversity in a Rao–Blackwellized particle filter for SLAM after actively closing loops. *Proceedings of the IEEE International Conference on Robotics and Automation*.

Stachniss, C., Hahnel, D. and Burgard, W. (2004). Exploration with active loop-closing for FastSLAM. *Proceedings of the IEEE/RSJ International Conference on Intelligent Robots and Systems*, Vol. 2.

Thrun, S. (2003). Learning occupancy grid maps with forward sensor models. *Autonomous Robots*, **15**(2): 111–127.

Thrun, S., Burgard, W. and Fox, D. (2005). *Probabilistic Robotics*. Cambrideg, MA, MIT Press.

Vlassis, N., Motomura, Y., and Krose, B. (1999). An information-theoretic localization criterion for robot map building. *Proceedings of International Conference on Machine Learning and Applications*, pp. 1–6.

Yamauchi, B. (1998). Frontier-based exploration using multiple robots. *Proceedings of the Second International Conference on Autonomous Agents*. New York, ACM Press, pp. 47–53.