Probabilistic 2D Mapping of Unstructured Environments

Jeff Leal, Steve Scheding, Gamini Dissanayake

Australian Centre for Field Robotics.
The University of Sydney.
Sydney, NSW, 2006
Australia

Abstract

In this paper, a novel approach for online 2D multi-resolution mapping of unknown, unstructured and expandable environments is presented. The main contribution of this paper is a surface representation that captures the uncertainties in the sensed data, providing a compact and efficient map. The algorithm presented fits a set of linked line segments, in a least mean-square sense, to the statistical representation of the environment. Furthermore, the algorithm is driven by application specific utilities that enable it to adapt the set of line segments to the surface such that the resulting representation is appropriate for the desired application. Experimental results are presented to demonstrate the effectiveness of the proposed algorithms.

1 Introduction

This paper is concerned with the problem of building an on-line multi-resolution map of an environment, by information gathered from moving sensors. The goal is to generate a compact representation of the environment that can be modified based on the requirements of a given application. The main focus is on strategies to incorporate information observed from sensors into the existing map and to expand the map as sensors explore the environment.

Most multi-resolution mapping strategies are based on storing the vertices of a geometric structure (line segments, triangles, etc.) at different resolutions to obtain various levels of detail (LOD) representations. Examples of such an approach can be found in Klein *et. al*, [Klein *et .al*, 1997]. The main alternatives to such strategies are those that require maintaining an occupancy grid in order to store up to date information about the environment. Occupancy grids were pioneered by Moravec and Elfes, [Moravec and Elfes, 1985; Elfes, 1989]. Occupancy grids give the flexibility to accomadate various spatial sensors and can also be used to represent dynamic environments. The weakness of this approach is that it is necessary to maintain a grid that is discretised at predetermined intervals. This is usually computationally expensive. In this paper, a particle-based representation that avoids the need for a grid-based discretisation is presented. The proposed strategy maintains sufficient information about the environment in order to perform on-line multi-resolution mapping and have a compact representation of the environment.

Many researchers have proposed the idea of using a population of points (particles) to represent a probability distribution function in numerous applications. This idea was first suggested by Gordon [Gordon *et al.*, 1993] as the bootstrap filter, also denoted as the particle filter. Research on this topic can be found in [Carpenter *et al.*, 1999], [Pitt and Shephard, 1997]. In addition, some applications of this approach can be found in [Fox *et al.*, 1998].

The following section describes an overview of the presented approach. Section 3 discusses the main issues that influence the particle-based representation. Section 4 discusses the decision-making mechanisms used in order to maintain a good approximation of the map. The paper also presents some experimental results using data gathered from an unstructured environment.

2 Overview

The goal of the map building algorithm is to construct a two dimensional map that will approximate the true surface, S , of the environment. In many outdoor and in most indoor environments it is acceptable to assume that the walls or obstacles present span vertically from the ground upwards. The map representation chosen, therefore, consists of a linked set of line segments $S^* = \{s_1, s_2, \ldots, s_n\}$. Each line segment s_i locally approximates the true surface S in a least mean-square sense. The restriction on S^* to a linked set is mainly to maintain map compactness, although there is no constraint on the number of sets S_k^* used to represent a given environment.

During on-line map building, it is necessary to make fast and accurate decisions on the appropriateness of the map. Information contained in the linked set of segments is not sufficient for this purpose. For instance, a large segment may represent a large planar region or it may represent a region that has not been sensed. If only line segments are stored, these two situations are indistinguishable. Therefore, in order to update S^* to incorporate new observations or changes in the environment additional information is required.

In the algorithm presented here, the extra amount of information necessary to perform map updating is stored as sets of particles that provide a global approximation of the sensed environment S. Every line segment s_i has an associated set of particles $P_i = \{p_1, p_2, ..., p_{N_i}\}$. Use of a lower-level representation based on particles encapsulates all the uncertainties in the map building process (sensor location uncertainties and sensor measurement uncertainties), together with a higher-level representation based on line segments provide an excellent framework for decision-making.

The sensing issues and map representation issues are further developed in the following sections.

3 Sensing and Representation

The main issues that arise in on-line map building reside in how good is an observation, will it improve the map, and how to incorporate it into the map. In off-line map building, the latter is separated from the sensing issues, and the result is a map that is limited to a set of applications. The particle-based representation is an approach towards managing these issues simultaneously.

3.1 Sensing

Sensor observations introduce uncertainties in their measurements. In addition, there are uncertainties introduced by not knowing the exact location of the sensors. Observations are usually modeled by some process $z(t)$,

$$z(t) = \mathrm{h}[\mathbf{x}(t), t] + \mathrm{w}(t) \qquad (1)$$

where **x** is a state vector and **w** is the process noise. In most cases, **h** is nonlinear (e.g. range-bearing sensors). Furthermore, **w** is usually modeled as a Gaussian, which may not represent the true characteristics of the sensor uncertainties. Performing linearisation from sensor space to Cartesian space also introduces errors [Julier and Uhlmann, 1997]. These limitations can be overcome by using a particle representation. With the particle representation an explicit form for the observation model is not needed, instead sensor uncertainties are represented by a particle cloud.

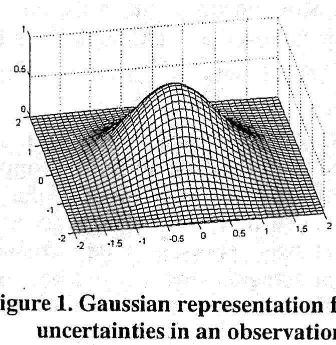

Figure 1. Gaussian representation for the uncertainties in an observation

Figure 1 shows the uncertainty, in sensor space, for an observation where **w** is zero-mean, uncorrelated, Gaussian noise.

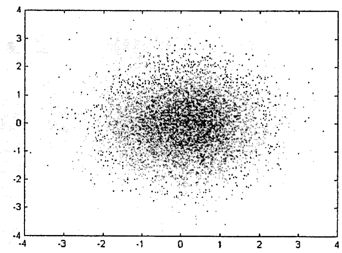

Figure 2. Particle-based representation of the uncertainties in an observation

Figure 2 shows the same observation uncertainty as a particle distribution. As can be seen, the likelihood of the observation is now represented by the density of the particles. Transforming the latter representation into Cartesian space is trivial (being simply the projection of the particles through the sensor model h), whilst in the former case one must deal with linearisation issues. Sensor location uncertainties simply add to the sensor measurement uncertainty.

3.2 Surface Representation

The previous subsection demonstrated how the particle distribution is useful for representing a sensor observation and its uncertainties. This subsection, illustrates how this representation can also describe an entire surface and its uncertainties.

Because the observations represent the amount of information in an infinitely small region of the true surface, may be easily understood that the total available information about the entire surface is simply a cumulative distribution of the sensed information.

In order to perform only local updates of the approximated surface the global particle representation is divided into local point clouds, each belonging to a line segment. Due to computation and storage issues, the algorithm for building this collective distribution relies on sampling/resampling and data association methods. These issues further discussed in Sections 4.3 and 4.4.

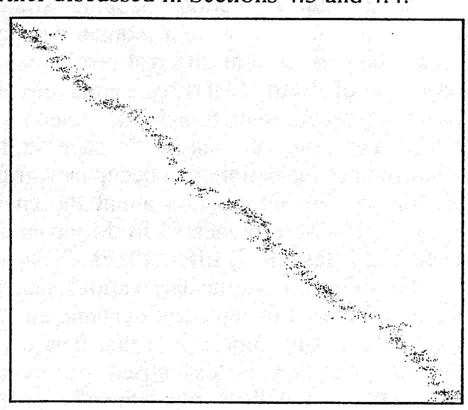

Figure 3. Particle representation of surface uncertainty

An example of a particle representation for a region of the environment is shown in Figure 3. Regions of higher density are the most likely to belong to the true surface S (in a probabilistic sense). The regions with low particle density correspond to parts of the surface that were unobserved or observed with poor accuracy. The particle representation, therefore, incorporates sufficient information that it can be used to formulate a decision-making mechanism for the entire surface, in terms of optimal line segment fitting.

Another advantage of this representation is that there is no need to store a grid for representing observation uncertainties or, for the entire distribution of the approximated surface. In a grid-based approach one would simply need to store mean and variance (assuming Gaussian uncertainty models) for all observations, in order to reconstruct a measurement of the uncertainties in the entire map. In the particle-based case, no assumptions are made on sensor models and the result is a global distribution that represents the true surface uncertainties. In addition, this representation preserves all the information necessary to perform isometric transformations on the segment set S^*. In other words, the statistical foundations of the underlying particle-based representation, allows modification of S^* without the loss of global information.

In the following section, a decision-making mechanism for fitting a set of line segments, S^* to the particles distribution is developed.

4 Decision-Making

The goal now is to fit a set of line segments to the particle representation. This is performed by creating the segment set S^* and fitting it to the particle distribution in a least mean-square sense. Section 4.1 describes the algorithm used to perform the least mean-square fit.

Once the line segment set is created it is modified based on subsequent information gathered from the sensors. It is also necessary to modify the structure of the line segment set in order to minimize the error in the between S^* and the true surface S, and to maintain map compactness. This can be achieved by splitting and merging the lines in S^*. Splitting consists of breaking a segment into finer detail so that S^* results in an improved local fit to the particle-based distribution. Merging consists of joining more than one segment to eliminate unnecessary (or unwanted) detail.

The decisions to perform certain operations locally in the map are based on utilities. Utilities are expressions that provide a voting mechanism on which operation is best suited for a particular region given a set of goals. The functionality of this approach enables the map to be used in various applications (e.g. navigation, active sensing, terrain-following, localization, etc.). These are discussed in Section 4.2.

In order not to fit segments to the uncertainties in the data, but to a local approximation of the true surface, a notion of maximum resolution or minimum segment length($l_{\min}$) is introduced. $l_{\min}$ correspond to the sensor resolution or to the finest detail required by the intended application for the map.

Additional choices that need to be made in order to maintain the representation compactness and computational efficiency are presented in sections 4.3 and 4.4.

4.1 PCA Algorithm

To be able to perform a least mean-square fit of a segment to a two-dimensional point cloud it is necessary to perform more than just a one axis-oriented fit. Principal Component Analysis (PCA) [Kendall, 1975], allows the decomposition of any number of particles into N maximum variation directions in N-dimensions.

In this application, PCA consists of decomposing the particle-based density distribution into two major directions. The basic idea is to find the principal components λ_1, and λ_2 so that they describe the maximum amount of variance possible in two orthogonal directions of the particle-based representation.

The computation of the principal components of a set of particles can be simply accomplished using the particle covariance matrix, C_x:

$$C_x = E\left\{(\mathbf{x}-\hat{\mathbf{x}}).(\mathbf{x}-\hat{\mathbf{x}})^T\right\} \qquad (2)$$

where $\mathbf{x}$ is the particles x-y locations, and $\hat{\mathbf{x}}$ is the mean of $\mathbf{x}$. The principal components are consequently given by the eigenvalues of C_x, and their directions are given by the corresponding eigenvectors.

The largest principal component is the projection on the direction in which the variance of the particles is maximized. It can be proven that the representation given by PCA is an optimal linear dimension decomposition technique in the least mean-square sense [Jolliffe, 1986].

The next step is to fit multiple segments to the particle representation using the PCA algorithm iteratively. In order to perform this, an iterative decision-making mechanism is used to perform merging and splitting operation on the line segments based on the utilities of each of these actions.

4.2 Utility Expressions

The structure of a utility expression is usually in the form of a measurement of value or *"usefulness"* for making a decision; in this case, splitting, merging or "no action". The general utility expressions can be stated, for any segment s_i, by

$$\begin{aligned} U_{split}(s_i) &= f(\mathbf{O}_1(s_i), \mathbf{O}_2(s_i), \ldots, \mathbf{O}_l(s_i), l_{\min}) \\ U_{merge}(s_i) &= g(\mathbf{O}_1(s_i), \mathbf{O}_2(s_i), \ldots, \mathbf{O}_l(s_i), l_{\min}) \\ U_{nothing}(s_i) &= h(\mathbf{O}_1(s_i), \mathbf{O}_2(s_i), \ldots, \mathbf{O}_l(s_i), l_{\min}) \end{aligned} \qquad (3)$$

where O_l is a goal function, dependant on the local particle representation for segment s_i. The functions f and g do not necessarily need to be the inverse of each other. Also h does not necessarily need to exist. The highest of the utilities decides the operation that will be performed on segment s_i.

The goal of the map-building algorithm is to minimize the mean-square error in the normal direction of a segment. In order to do this, a goal function is defined as the normal variance of the particles, σ_n to a segment. To enforce S^* to preserve at most the maximum resolution, a second goal function is defined. This is quickly calculated by the PCA algorithm. Consequently, the utility system for the splitting and merging during map building can be expressed as:

$$\begin{aligned} O_1(s_i) &= \sigma_n(P_i) \\ O_2(s_i) &= length(s_i) \end{aligned} \qquad (4)$$

$$U^1_{split}(s_i) = f(O_1(s_i), O_2(s_i), l_{\min}) = O_1(s_i) \geq \sigma_{th} \text{ and } O_2(s_i) \geq l_{\min}$$

$$U^1_{merge}(s_i) = g(O_1(s_i), O_2(s_i), l_{\min}) = O_1(s_i \mid s_{i+1}) \leq \alpha\sigma_{th}$$

$$U^1_{nothing}(s_i) = h(O_1(s_i), O_2(s_i), l_{\min}) = \text{not}(U_{split}(s_i)) \text{ and } \text{not}(U_{merge}(s_i))$$

where $0 \leq \alpha < 1$, and σ_{th} is the minimum variance threshold that describes the maximum map quality. The utilities described by Equations 4 are independent of the location of the observer. Segment sets obtained for part of the environment obtained using the above utility expression are shown in Figure 4.

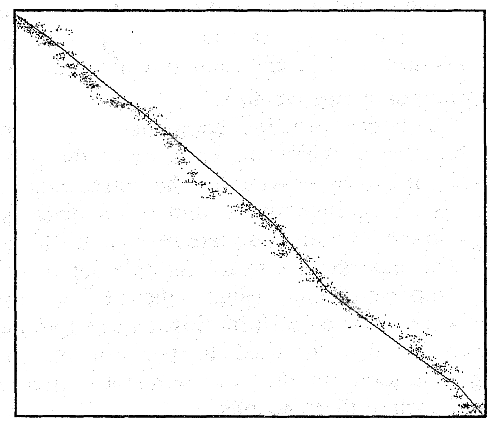

Figure 4. Particle distribution over five segments. Area of map: 7m by 8m.

Utilities can be selected in order to obtain a map that is appropriate for a given application. The utility system described by:

$$U^2_{split}(s_i) = \begin{cases} f(O_1(s_i), O_2(s_i), \alpha.l_{\min}), & s_i \cap \text{VF} \\ f(O_1(s_i), O_2(s_i), l_{\min}), & \text{otherwise} \end{cases}$$

$$U^2_{merge}(s_i) = \begin{cases} g(O_1(s_i), O_2(s_i), \alpha.l_{\min}), & s_i \cap \text{VF} \\ g(O_1(s_i), O_2(s_i), l_{\min}), & \text{otherwise} \end{cases} \qquad (5)$$

$$U^2_{nothing}(s_i) = \begin{cases} h(O_1(s_i), O_2(s_i), \alpha.l_{\min}), & s_i \cap \text{VF} \\ h(O_1(s_i), O_2(s_i), l_{\min}), & \text{otherwise} \end{cases}$$

is a for a viewpoint dependant application, where the goal is to have higher detail of the surface S^* in the proximity of an observer, whilst reverting to the viewpoint invariant utility in other regions. Equation 5 also demonstrates how utilities can be combined, in order to achieve multiple goals.

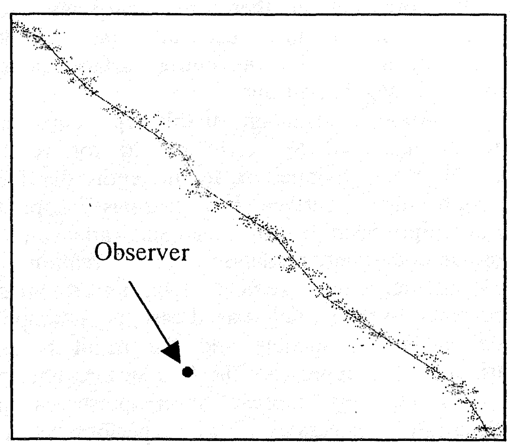

Figure 5. Utility-driven approximation of the surface. Area of map: 7m by 8m.

Figure 5 shows how a viewpoint dependant utility enforces higher detail than the viewpoint invariant scenario of Figure 4. In the case shown in Figure 5 the settings were $l_{\min} = 1.5m$ and $\alpha = 0.53$.

4.3 Data Association

Given the fact that observations only provide local information about the true surface it is unnecessary to update the entire surface when a local observation is made. For this reason, the particle representation is locally distributed across every segment. This reduces computation and data storage. However, this makes it necessary to perform data association to match an observation to its corresponding particle representation, P_i and hence segment s_i.

The data association method used in this paper is a simple point-in-polygon method. The mean of each sampled observation is tested with bounding polygons of every segment.

4.4 Sampling / Resampling

Each observation is sampled from its uncertainty model M times. In order to incorporate into the map the new observation, the particle distribution P_i of segment s_i relies on a resampling technique.

To reduce storage space every particle distribution P_i has a limit to the amount of particles it contains. This limit is proportional to the length of the segment with an upper bound dependant on $l_{\min}$.

The new particle representation for a particular segment is resampled from a uniformly distributed function over the number of prior samples that exist in that segment. By choosing N_i large, this is equivalently (in

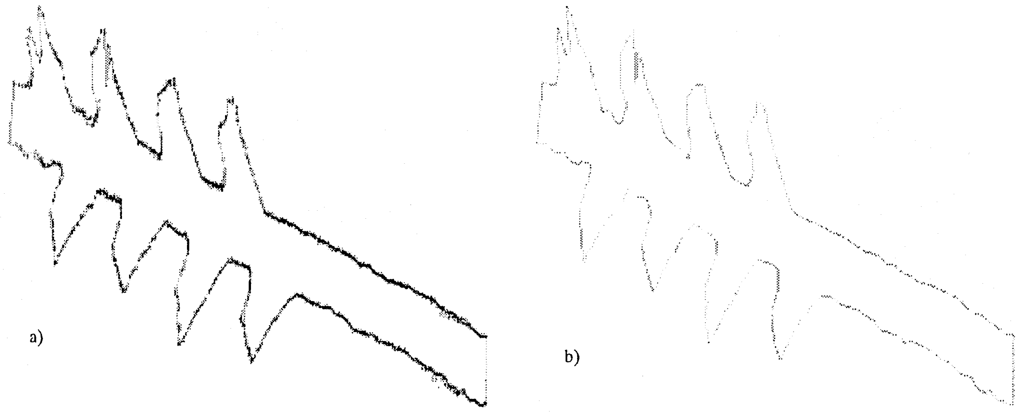

Figure 6. (a) Viewpoint invariant distribution map. (b) Segment set (without the particles)

the limit) to removing less likely distributed particles. The resampling technique consists of randomly removing M particles from the discrete uniform probability density function F_u.

$$F_u^i(n) = \frac{1}{M + N_i} \quad , \quad n \in [0, M + N_i - 1] \quad (6)$$

This resampling technique together with the PCA algorithm bears a resemblance to the bootstrap method [Efron and Tibshirani, 1993]. The difference lies in that the sampling is done without replacement. This was found to produce slightly better results.

5 Results

Various tests were performed in an unstructured environment to evaluate the proposed map building algorithm. The data was obtained from a sick laser range finder mounted on a 4WD. In this section, maps obtained using the viewpoint invariant and the viewpoint dependant utilities are shown.

The mapped environment corresponds to an area of 90m by 30m. The minimum segment length that was found to be optimal was approximately 1m. M was set to 30 sampled particles from the observation model of the sensor. A maximum of 300 particles per meter was used in the representation of the environment. The results obtained are shown in Figure 6.

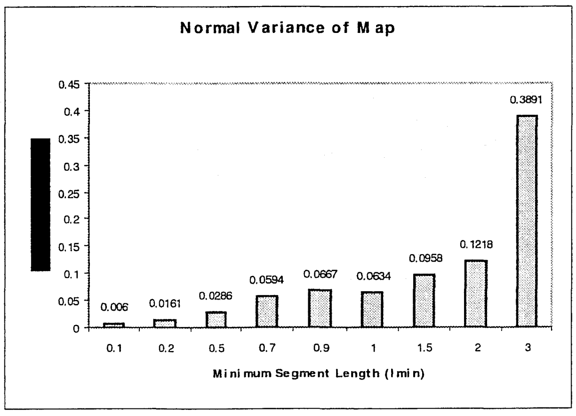

Figure 7. A global measurement of the maps quality with different resolutions.

Figure 7 shows the normal variance of the particle-based distribution given the final map S^* shown in figure 6. In this particular environment and with the sensor used it was found that by using $l_{\min} \leq 0.7m$, the maps that were obtained fitted S^* to the errors in our observation models and not to the walls of the environment. Also for values of $l_{\min} \geq 1.5m$ the error in the map was beginning to be large.

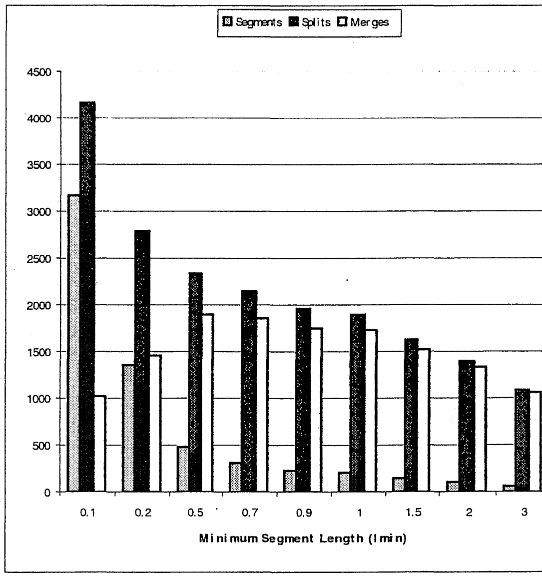

Figure 8. Total number of segments, splitting and merging for the different map resolutions.

Figure 8 demonstrates how the choice of the resolution of the map affects the amount of segments and operations performed during the map building process.

Figure 9 demonstrates a run using a viewpoint dependant utility as discussed in Section 4.2. As can be seen from Figure 9(a), S^* is not a fine approximation to the underlying particle distribution other than in the right region of the map where the walls are mainly straight. In Figure 9(b) a point observer was moved through the map and resulted with a higher detailed map in the proximity of

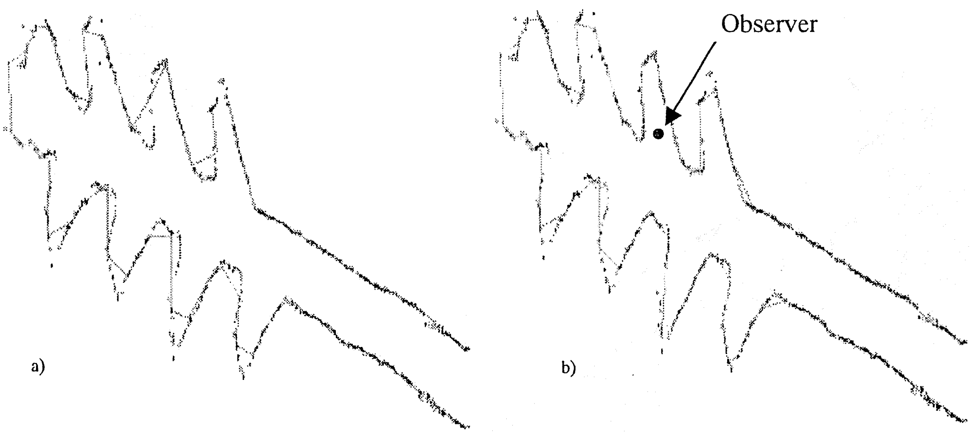

Figure 9. (a) Map with minimum segment length of 3m. (b) Utility run with a=0.33

the observers' location.

5 Conclusion

This paper presented an approach to map building that encapsulates the global properties of the sensed environment, such as sparse observations and uncertainties in the sensors, within the map's representation. It enforced the necessity of a probabilistic representation of the environment in order to perform substantiated operations and to be able to measure the quality of the map at any given time. It then demonstrated the use of this information in order to set up a utility-based framework to be able to apply the map to different applications. Experimental results were shown, in a large, unstructured and expandable environment demonstrating the success of this approach to map building. We furthermore demonstrated results with the use of different utility expressions. The key idea of having the particle-based distribution representation in our approach is to be able to preserve the uncertainties in all the stages of the map building process.

6 Future Work

Although the promising results obtained with this approach, there is room for development. In future work, adaptive sampling techniques are to be implemented in order to spread the particle-based distribution to the most informative regions of the sensed environment. This will drastically reduce data storage and improve map quality and information.

The main objective is to take this map building approach to three dimensions.

References

[Carpenter *et al.*, 1999] Carpenter, J. Clifford, P. Fearnhead, P. *Improved particle filter for non-linear problems*. IEE Proceeding on Radar, Sonar and Navigation. Vol. 146. No.1, February 1999.

[Efron and Tibshirani, 1993] Efron, B. and Tibshirani, R. J. *An Introduction to the Bootstrap*. Chapman and Hall, 1993.

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

[Fox *et al.*, 1998] Fox D. Burgard, W. and Thrun S. *Probabilistic approach to concurrent mapping and localization for mobile robots.* Machine Learning 31, 1998

[Gordon *et al.*, 1993] N.J. Gordon, D.J. Salmond, A.F.M. Smith. *Novel approach to non-linear/ non-gaussian Bayesian state estimation.* IEE Proceedings-F Vol. 140, No. 2, 107-113, 1993

[Jolliffe, 1986] I.T. Jolliffe. *Principal Component Analysis*. Springer-Verlag, 1986.

[Julier and Uhlmann, 1997] S.J. Julier and J.K. Uhlmann. *A Consistent, Debiased Method for Converting Between Polar and Cartesian Coordinate Systems.* The Proceedings of AeroSense: The 11th International Symposium on Aerospace/Defense Sensing, Simulation and Controls, Orlando, Florida, March 1997.

[Kendall, 1975] M. Kendall. *Multivariate Analysis.* Charles Griffin&Co., 1975.

[Klein *et al.*] R. Klein, D. Cohen-Or, T. Hüttner. *Incremental View-dependant Multiresolution Triangulation of Terrain.* November 1997

[Moravec and Elfes, 1985] H. Moravec and A. Elfes. *High-resolution maps from wide-angle sonar*. In *Proc. IEEE Int'l Conf. on Robotics and Automation*, St. Louis, Missouri, March 1985.

[Pitt and Shephard, 1997] Pitt, M.K. Shephard, N. *Filtering via simulation: auxiliary particle filters.* October 22, 1997.