

3D Terrain Mapping: A Stochastic Approach

Jeff Leal, Steve Scheduling and Gamini Dissanayake

Australian Centre for Field Robotics
Rose Street Building J04
The University of Sydney, 2006.
{jeff,scheduling,dissa}@acfr.usyd.edu.au

Abstract

In this paper, a novel approach for online 3D multi-resolution surface reconstruction of unknown, unstructured environments is presented. The algorithm is based on a probabilistic framework that enables the capture of the uncertainties involved throughout the knowledge gain process. The main contribution of this paper is a stochastic representation for information. We present some results obtained in unstructured environments.

1 Introduction

For an agent to interact with its environment it needs some sort of perceived representation of its surroundings. An agent is defined in this work as an entity that can interact with its surroundings. The perception that each agent has of the environment depends on the type of sensors used to obtain information and the agents objectives. These objectives can be manifold (e.g. navigate, detect features, path-planning, etc.). Consequently, the way each agent makes use of the information about the reconstructed environment will be different. Hence constructing the environment with the relevant information and presenting it differently to the various agents will improve their decision making capabilities.

Most previous approaches to surface reconstruction in robotics applications only address a single objective, i.e., a map is generated for a single purpose. One application for surface reconstruction is to generate a meshed model in order to perform differential computation on localized nodes of that object [Chew, 1993]; further employment of such techniques are for visualization purposes [Hoppe *et al.*, 1992]. There is a large amount of literature on the subject of mesh generation, and surveys on various surface reconstruction and meshing techniques can be found in [Bern and Eppstein, 1992; Heckbert and Garland, 1995; Owen, 1998]. In the work presented the objective is

to generalize the surface reconstruction process into a multi-objective problem. Hence the issue now becomes one of maintaining a knowledge base of sensed information in order to maximize the *utility* of the representation at any given time.

Other issues present in most previous algorithms that generate surfaces is one of not quantifying the errors and uncertainties in the observed information and one of accepting the generated representation as a “true” model of the surface. This makes performing decisions (e.g., decimation, localization, navigation, etc.) a difficult or even unfeasible task. Sensors are not perfect and usually provide incomplete information about the features of interest. If this is not accounted for, the result of making decisions under the assumption of the representation being correct may result in equipment loss and mission failure. The approach presented here introduces a representation for the summary of the uncertainties involved in the reconstruction and information gain processes of map-building. Therefore, decisions can be performed under the knowledge of the amount of uncertainty in the surface. The algorithm also enables for fusing data from multiple viewpoints and different types of ranging sensors (e.g. Laser, Radar, Stereo vision, etc.).

The following section describes the methodology used in this approach. Section 3 describes the representation of information. Section 4 describes the surface representation. Finally, some experimental results using this approach are presented, followed by conclusions.

2 Methodology

To fit an approximated surface to data obtained from multiple viewpoints and from various sensors, it is necessary to know with some degree of certainty the sensor locations. With the work presented in this paper, sensor location is assumed known. As will be seen in Section 3, uncertainty in sensor location can be dealt with easily.

Another issue that is necessary to account for is the uncertainties in the sensor measurements. These can be modelled with some degree of accuracy with the information obtained from the manufacturers *spec sheets* or through empirical testing of the sensor. Although bounds can be chosen on these sources of uncertainties, there still exists the issue of representing these uncertainties in a manner that can be useful for making decisions about the true surface location. For this, a particle-based representation is used to encapsulate the uncertainties from these sources. Particles are spawned from a likelihood distribution that represents the uncertainty model (e.g. sensor observation uncertainty model, sensor location uncertainty model, etc.). This information is then maintained in a database and decisions are performed on the knowledge contained in that database. A Delaunay Triangulation (DT) is fit to this data using a Weighted Principal Component Analysis method (WPCA)[Massy, 1965; Vijayakumar, 1998]. The two representations, the information summary, and the DT are completely independent. The DT is locally updated with actions that are the outcome from utilities defined over the particle representation. Figure 1 illustrates the general steps of the proposed algorithm for online multi-resolution map-building.

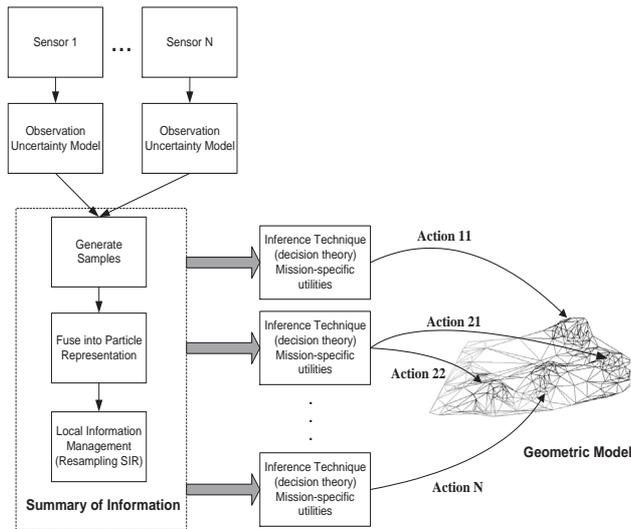


Figure 1: Flowchart illustrating the main steps of the proposed algorithm.

The particle representation chosen for the summary of information enables straightforward fusing of measurements from other sensors and other viewpoints. The following section illustrates this concept in more detail.

3 Information Representation

The summary of information consists of a set of particles which are used as approximations to an observation uncertainty distribution. This introduces

some advantages over other representations (e.g. State Space). Since the state space representation has been dropped, there is no need for linearisation [Julier and Uhlmann, 1997], consequently reducing approximation errors for sufficiently large number of particles. Additionally, this method does not require a predefined array or occupancy grid where states are determined by locations in that array, thus increasing the efficiency by directing computation to specific nodes rather than to every location in the predefined array. Finally, there is no need for the assumption of any model for the surface, hence avoiding the introduction of more sources of errors.

From this particle representation, two sources of knowledge can be extracted. One consists of the knowledge of the uncertainty in the surface from all previous observations, and the other consists of the knowledge of the amount of information in a particular region of the surface. Additionally, from these sources of knowledge a decision mechanism can be formulated. These issues are further developed in the following subsections.

3.1 Information Uncertainty

In order to distinguish between particles that represent states with lower observation uncertainty and states with higher observation uncertainty, the notion of likelihood is introduced. The probabilistic information contained in an event, z_k about the variable m is described by a distribution $p(z_k | m)$, known as the likelihood function. Sampling particles from the likelihood function allows to extract an approximation of the uncertainty distribution of a particular observation.

An example of a particle distribution for one observation, (r_{obs}, θ_{obs}) is presented in Figure 2. This was obtained from a range-bearing observation model. The figure contains 1000 particles spawned from the joint distribution in range, $N(r_{obs}, \sigma_r)$ and in bearing, $N(\theta_{obs}, \sigma_\theta)$. Each particle is then weighted with the likelihood $p(z_k | m)$:

$$p_i(z_k = (r_{obs}, \theta_{obs}) | m) = \frac{e^{-\frac{(r-r_{obs})^2}{2\sigma_r^2}}}{\sqrt{2\pi}\sigma_r} \cdot \frac{e^{-\frac{(\theta-\theta_{obs})^2}{2\sigma_\theta^2}}}{\sqrt{2\pi}\sigma_\theta} \quad (1)$$

where $m \in M$, the surface location space and $z_k \in Z$, the observation space.

The result of every observation is then a group of particles that represent possible surface locations and the likelihood of these possibilities. This characterizes not only the local surface but also the observation uncertainties. This methodology can be easily extended to sensor location uncertainty. As location uncertainty increases(decreases) the uncertainty in measurements increases(decreases). This allows to propagate location uncertainty through the observation model and into observation uncertainty.

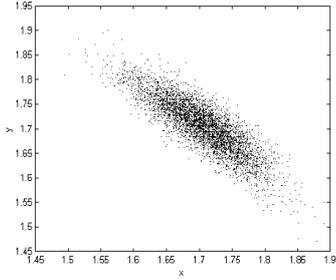


Figure 2: Particle-based representation for the likelihood distribution from a range-bearing observation model(the density of particles depict regions of higher likelihood).

3.2 Information Gain

The increase in information is simply a summary of new information with previous gained information. At a global scale, the information about the entire surface increases as new observations are made. Subsection 3.3 focuses on the issue of limiting the amount of global information, but still enable the fusing of new information. At a local scale, the amount of information is represented by the amount of particles in that region. Additionally, the quality of that knowledge is depicted by the particle weights.

As a result of this *information* knowledge, the ability to define inference techniques in order to make decisions (e.g. active sensing, path planning, feature extraction, etc.) is enabled.

3.3 Information Management

Due to limitations that might be imposed on the system, such as time and storage space, it may be necessary to reduce the amount of information in the particle representation. For this, resampling techniques are employed. There are various resampling techniques that have different outcomes. In the context of this work, the objective is to reduce the amount of knowledge but still maintaining the quality in that knowledge. One technique that performs this in a stochastic manner is Sampling Importance Resampling (SIR)[Fearnheard, 1998; Smith and Gelfand, 1992].

The idea is to sample a new set of N particles from the prior $p(m | Z^{k-1}, M_k)$ to obtain a set of samples that represent the posterior $p(m | Z^k, M_k)$. M_k denotes a subset of the particles, in a neighboring region, to a point m on the surface distribution and the prior is represented by the particle distribution in that region.

Let P denote the number of samples which are distributed from a continuous density function $g(x)$, and that samples are required from the probability density function (pdf) proportional to $l(x).g(x)$, where $l(x)$ is a known function. The sampling importance resampling method states that N samples drawn from the discrete distribution $g(x)$, with probability mass,

$$w(\vec{m}_i) = \frac{l(\vec{m}_i)}{\sum_{M_k} l(\vec{m}_i)} \quad (2)$$

tends to the required density as N tends to infinity.

If $g(x)$ is taken as $p(m | Z^{k-1}, M_k)$ and $l(x)$ as $p(z_k | m, M_k)$, then the theorem provides a means of obtaining a new set of N samples, which represent the new posterior distribution. Proof of this result can be found in [Smith and Gelfand, 1992].

Taking $N < P$, is equivalent to reducing the number of particles that represent the posterior distribution, hence reducing the amount of knowledge.

4 Surface Representation

In order to visualize a surface or to perform some simple decisions (e.g., distance to surface, local/global surface resolution, etc.), an approximation to the true surface must be chosen. As mentioned previously a Delaunay Triangulation was employed for the work presented in this paper, although there is no restriction on the model chosen. The reason for this choice is due to the DT properties of well formed triangles and minimized global amount of triangles over a surface[Guibas and Jorge, 1985].

Due to the fact that a summary of previously obtained information is stored, an online multi-resolution surface approximation can now be formulated and fit to the surface at any time. This property enables the representation of a surface that assists with mission-specific goals. Additionally, due to conflicting representations that may exist for different mission objectives (e.g. Navigation, Path Planning, Localization, etc.) multiple surfaces can be generated from the same underlying source of information, the particle representation.

In order to approximate triangles to a set of weighted particles, a Weighted Principal Component Analysis is employed.

4.1 Weighted PCA

This technique is used for approximating a triangle to the current description of the true surface (the particle representation) in a given region. In this work, the degree of optimality in the fit of an approximated surface to a true surface is defined as to minimize the mean square error (MMSE) between both surfaces. 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].

Due to the fact that the true surface is unknown and that there are uncertainties in the measurements, the MMSE will have uncertainty associated with it. Thus the MMSE is performed by calculating the variability of the information distribution with the WPCA algorithm taking into account the weights in the particle distribution. This enables particles with more likelihood to belong to the true surface to have more emphasis on the outcome of the WPCA.

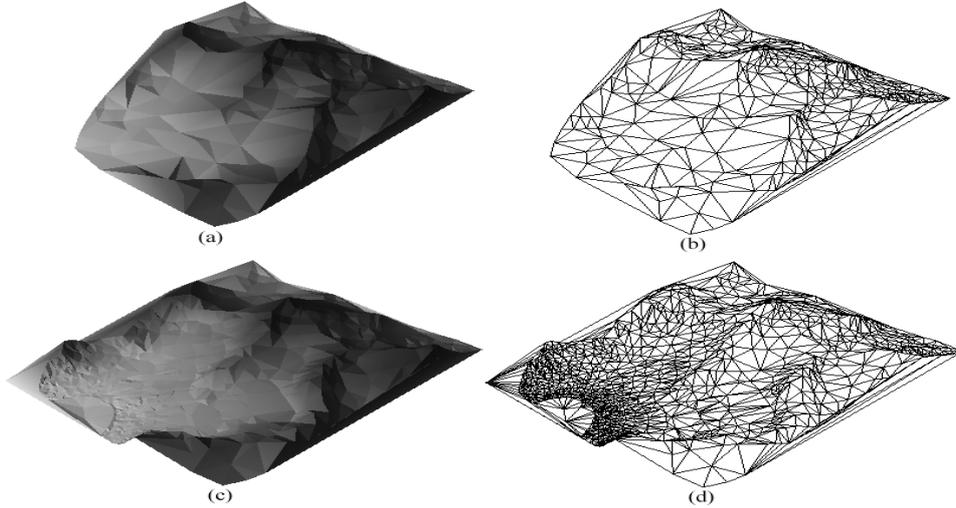


Figure 3: The triangulated reconstructed surface from multiple viewpoints, and with multiple scanning resolutions; (a, b) after a scan from first viewpoint with a two degree scanning resolution in pan and 0.5 degree in the tilt axis; (c, d) after an additional scan from a second viewpoint with a 1.5 degree scanning resolution in both axis. The MMSE threshold of the local particle representation over triangles was set to $\sigma^2 = 0m^2$.

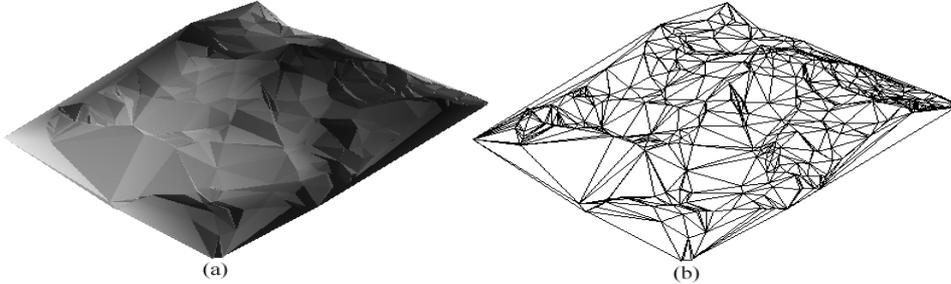


Figure 4: Triangulated reconstruction of the surface from multiple viewpoints, and with multiple scanning resolutions. Obtained from the same particle representation as in Figure 3. In this case, the MMSE threshold of the local particle representation over triangles was set to $\sigma^2 = 0.0005m^2$.

The WPCA algorithm enables the decomposition of any number of particles into N maximum variation directions in N -dimensions (principal components). In 3D, this consists of decomposing local particle distributions into three major directions. In order to perform this, a local covariance matrix is defined as,

$$C_l = \frac{\sum_{i=1}^N w(\vec{m}_i) \cdot (\vec{m}_i - \bar{m}) \cdot (\vec{m}_i - \bar{m})^T}{\sum_{i=1}^N w(\vec{m}_i)} \quad (3)$$

where $w(\vec{m}_i)$ are the weights defined by Equation 1 and \vec{m}_i are the locations of the particles in 3D space. The principal components are consequently given by the eigenvalues of Equation 3 and their directions are given by the corresponding eigenvectors. The smallest principal component is the projection on the direction in which the variance of the weighted particles is minimized. This direction, corresponds to the normal of the particle set. Although, in some under-observed cases this normal has not converged to a reasonable direction. Triangles are then updated to reflect the changes

in the particle distribution by adjusting the normals of the triangles with the normals obtained by the WPCA minimum component.

4.2 Surface Management

In order to approximate a set of triangles to the particle distribution with higher detail, it may be necessary, at some stage, to increase the number of triangles. This is performed by splitting triangles and re-fitting these triangles to the particle distribution. Additionally, it may be necessary to reduce the number of triangles in order to reduce or eliminate unnecessary detail. This is performed by merging triangles into lesser triangles and re-fitting them to the particle distribution.

The moment for choosing a decision of a particular action to be performed (split, merge, update) can now be defined in a decision-theoretic framework. With the definition of utilities, it is possible to achieve surfaces that are more useful to specific applications. This technique was shown to be efficient in 2D map-building [Leal *et al.*, 2000]. At the current stage of this work,

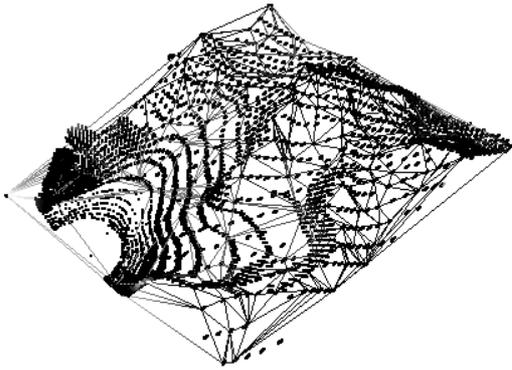


Figure 5: The representation of two levels of knowledge about the surface: The particle representation, and the triangulated surface (DT).

only simple threshold decimation utilities were implemented in 3D.

5 Experimental Results

Simulation experiments were generated from real DEM data. Results from using this algorithm with data obtained from multiple viewpoints are shown in Figures 3 and 4. Figure 3 depicts the reconstructed terrain from various viewpoints. Figures 3(a)(b), were obtained at a low resolution from a distant viewpoint. Figures 3(c)(d) were obtained after the first scan by adding information from another sensor in a closer location providing higher resolution data. Additionally, the surface management layer was designed to maximize the resolution of the reconstructed surface, by limiting the normal error of every triangle to zero variance. This resulted in a highly triangulated terrain. On the other hand, in Figure 4, the utilities were set to obtain a lower resolution triangulation, by limiting the normal error of every triangle to a 0.0005 variance. This small deviation from the first example allowed the result of a less triangulated surface without loss of the main features of the map. This is due to a higher triangulated terrain where surface variations are greater and lower triangulated regions, where there is less variation (flatter regions).

Figure 5 depicts the two levels of representation in the knowledge base. One is the geometric model (DT) that is designed to fit in a Least Mean Square (LMS) sense to the other, the particle representation. The utilities defined in these examples are to minimize local errors, forcing a threshold on the global error in the final reconstructed surface.

Figure 6 demonstrates the decisions performed at the time where new knowledge is obtained. Also represented is the number of triangles at every step. An interesting result is seen, up to the stage where the first scan, from a distant viewpoint is performed (approximately observation 750). In this region, the increase in triangles is greater. This is because initially there is

no information in most regions of the environment, so any sensorial information is knowledge gain, hence triangulating these regions. After this, new information is fused and the increase in knowledge isn't as large. There is also a noticeable plateau in the number of splits at the beginning of the second scan. This is due to the fact that the observed data at this time corresponds to the flat region (easier seen as the lighter region in Figure 3c)). Additionally, around observation 1500^th , the larger variation in the terrain resumes the decimation of the surface.

6 Conclusions and Future Work

This paper presented an approach to surface reconstruction that encapsulates the uncertainties in the sensorial data. In addition, a stochastic framework was described to deal with such uncertainties and to maintain a consistent knowledge representation. Experimental results were shown, in a large, unstructured and expandable environment demonstrating the success of this approach in the surface reconstruction problem.

Work is currently directed in defining utilities that describe multi-objective agents. This will allow for the surface to alter its state in order to maximize the the outcome of an action performed by a particular agent.

References

- [Bern and Eppstein, 1992] M. Bern and D. Eppstein. Mesh generation and optimal triangulation. Csl-92-1, Xerox PARC, March 1992.
- [Chew, 1993] P.L. Chew. Guaranteed-quality mesh generation for curved surfaces. *The 9th Annual Computational Geometry*, 1993.
- [Fearnheard, 1998] Paul Fearnheard. *Sequential Monte Carlo Methods in Filter Theory*. PhD thesis, Merton College, University of Oxford, 1998.
- [Guibas and Jorge, 1985] L. Guibas and S. Jorge. Primitives for the manipulation of general subdivisions and the computation of voronoi diagrams in filter theory. *ACM Transaction on Graphics*, 4(2):74–123, April 1985.
- [Heckbert and Garland, 1995] P.S. Heckbert and M. Garland. Survey of surface approximation algorithms. Technical Report CMU-CS-95-194, CS Department., Carnegie Mellon University, 1995.
- [Hoppe *et al.*, 1992] Hugues Hoppe, Tony DeRose, Tom Duchamp, John McDonald, and Werner Stuetzle. Surface reconstruction from unorganized points. *Computer Graphics (SIGGRAPH)*, 26(2):71–78, 1992.
- [Jolliffe, 1986] I.T. Jolliffe. *Principal Component Analysis*. Springer-Verlag, 1986.
- [Julier and Uhlmann, 1997] S.J. Julier and J.K. Uhlmann. A consistent, unbiased method for converting between polar and cartesian coordinate

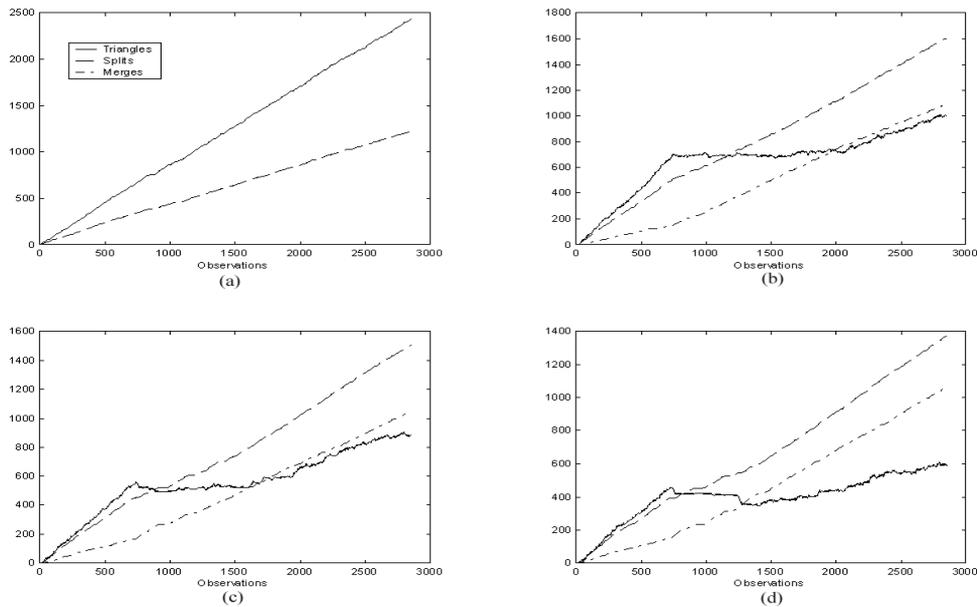


Figure 6: Number of operations and number of triangles at every stage of new information fusion with normal variance set to (a) 0; (b) 0.0005; (c) 0.001; (d) 0.002.

systems. *The Proceedings of AeroSense*, March 1997. Orlando, Florida.

[Leal *et al.*, 2000] J. Leal, S. Scheduling, and G. Dis-anayake. Probabilistic 2d mapping in unstructured environments. *Proceedings of the Australian Conference on Robotics and Automation*, pages 19–24, August 2000. Melbourne, Australia.

[Massy, 1965] W.F. Massy. Principal component regression in exploratory statistical research. *Journal of the American Statistical Association*, 60:234–246, 1965.

[Owen, 1998] S.J. Owen. A survey of unstructured mesh generation technology. *7th Proceedings of the International Meshing Roundtable*, pages 239–267, 1998.

[Smith and Gelfand, 1992] A.F.M. Smith and A.E. Gelfand. Bayesian statistics without tears: A sampling-resampling perspective. *The American Statistician*, 46(2):84–88, May 1992.

[Thrun *et al.*, 2000] S. Thrun, W. Burgard, and D. Fox. A real-time algorithm for mobile robot mapping with applications to multi-robot and 3d mapping. *IEEE International Conference on Robotics and Automation*, pages 321–328, April 2000. San Francisco, CA.

[Vijayakumar, 1998] S. Vijayakumar. *Computational Theory of Incremental and Active Learning for Optimal Generalization*. 95d38108, Tokyo Institute of Technology, 1998.