

# Real-Time Map Building from a Stereo Camera under Unconstrained 3D Motion

Heiko Hirschmüller  
School of Computing,  
De Montfort University, Leicester, LE1 9BH, UK.  
hhm@dmu.ac.uk.

## Abstract

*In this paper, a method is proposed for building occupancy grid based maps from sequences of stereo images in real time. The stereo camera can move and rotate unconstrained in three dimensions. This distinguishes the method from all other mapping approaches, which constrain the sensor to move in two dimensions. This freedom of motion allows to perform mapping in many new application areas. Additionally, three dimensional information can be maintained in layers or fused into a two dimensional overview of the environment. This transition is based on changeable rules rather than constraints. The method uses fuzzy logic and eliminates many problems that arise from traditional, probabilistic approaches for occupancy grids. It is shown that the proposed method performs well, is memory efficient and suitable for incremental use in real-time. The map will be used to assist navigation of a tele-operated mobile robot and real-time virtual walkthroughs, which are generated from the same stereo and motion data.*

**Keywords:** Map building, stereo vision, unconstrained camera motion, occupancy grid, fuzzy logic, real-time.

## 1. Introduction

### 1.1. Motivation

Maps are very useful tools for navigation in mobile robot applications. The movement of a robot is usually two dimensional on a ground plane. The most common robot sensors like laser scanner and ultrasonic sensors are designed to determine the distance of obstacles from the robot. Of interest are usually areas through which the robot can travel safely and those which are inhabited by obstacles. Thus, the problem can be considered as being two dimensional.

The real world however, is three dimensional and applications can benefit from knowledge about the height of obstacles or the existence of free space under an obstacle (e.g. table). Furthermore, it can be a strong restriction if the sensor must move in two dimensions. The contribution of this research is to allow the sensor to move and rotate unconstrained in three dimensions and to model three dimensional information properly within a two dimensional map. This allows to use mapping in new applications, e.g. walking or flying robots, sensors on manipulators or even in non-robotics applications like as aid for disabled people.

The focus of the current research is a tele-operated mobile robot that will be fitted with a stereo camera, possibly on the manipulator. A fast stereo vision system delivers dense depth images [1], while a camera motion estimation method determines the position and orientation of the camera, exclusively from stereo images [2]. This paper describes the incremental creation of a map from stereo and motion data. The map will be represented as a two dimensional overview to the operator for navigation of the robot. Furthermore, the map will also help to navigate inside a virtual environment, which will be created from the stereo and motion information.

It is important to note that the reduction to two dimensions is based on changeable rules and not on sensor or methodological limitations. The presented method can maintain the height information internally and represent it in different ways, if needed.

### 1.2. Review of Existing Methods

Maps can be represented using geometric primitives like lines. Data from accurate sensors like laser scanners may be used to identify line segments and fuse them into a global map [3]. Feature based stereo vision has also been used to create three dimensional sparse representations of a scene [4]. Even rather inaccurate ultrasonic sensor data can be used to create maps based on geometric primitives [5]. The uncertainty in the measurements is in all cases modelled in the parameters of the geometric primitives.

Another choice for map building are occupancy grids, which have been developed by Moravec and Elfes [6, 7]. The environment is divided into cells, which are either at least partly occupied or completely empty. Occupancy grids are especially suitable for fusing inaccurate ultrasonic sensor data and producing a rather accurate map.

The knowledge about the occupancy of an individual cell is commonly modelled as a single probability, which is updated with Bayes rule. However, the use of probability theory has been criticised for several reasons. Firstly, it is difficult to create accurate sensor models for new sensors. The characteristics of ultrasonic sensors are well known, but unrealistic simplifications are needed to model the complex behaviour of stereo vision. Hence, some authors even decided to skip probability theory and to invent an own update rule [8]. Secondly, a single probability does not allow

to distinguish between *unknown* and *uncertain* occupancy. Thus, it can not be determined whether an area has not been scanned at all (e.g. due to occlusion) or the sensor data was unreliable. The Dempster-Shafer theory has been used to tackle this particular problem [9]. Finally, the violation of strong assumptions can lead to large errors. The assumption of independence of different scans is usually violated for a slowly moving robot. Enforcing this independence artificially improves map building [10]. Another implicit assumption is that of evenly distributed occupied cells. Occupied cells appear usually clustered in natural scenes [11].

Gambino et al. compare probability theory and fuzzy logic for map building from ultrasonic sensors and conclude that the former method is very sensitive and prone to serious errors like those mentioned above [11]. Fuzzy logic was identified as performing better, which is attributed to less constraining hypotheses. This paper is based on a refinement of their method [12]. However, the idea had to be modified and extended to work with a completely different sensor type (i.e. stereo vision instead of ultrasonic sensors) and not in two, but three dimensions. It is shown that the choices of occupancy grids and fuzzy logic adapt well to the characteristics of stereo vision.

### 1.3. Overview over the whole System

Figure 1 gives an overview of the system. Fast stereo vision delivers dense disparity images [1]. A camera motion estimation method utilises the knowledge about the three dimensional position of features (i.e. corners, which are detected in the left image) to robustly and accurately determine motion of the camera [2]. There is no estimate from other sensors used or constraint on the motion needed, other than motion must not be too fast in respect to the frame rate (i.e. at least half of consecutive images should overlap). In fact, the system is usually tested by just taking the camera in the hand and walking around in the room.

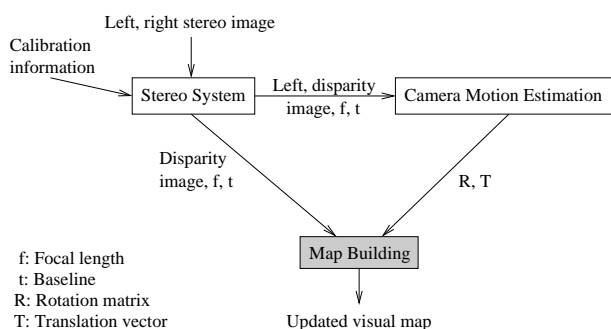


Figure 1: Overview over the whole system.

The complete system without mapping runs with  $\approx 10$ fps on a 1.2GHz Athlon, using a resolution of  $320 \times 240$  pixel with a search range over 32 disparities. This paper concentrates on creating a visual map from stereo and motion data using the FLOG (Fuzzy Logic based Layered Occupancy Grid) method. The input are disparity images and the corresponding absolute position and orientation of the

camera. It seems important for stereo vision to reduce general matching errors, which manifest themselves as spikes and disturb map building. A simple segmentation filter that removes areas below a certain size in the disparity image eliminates almost all outliers [8].

The new map building method is described in section 2. Section 3 discusses results from two kinds of experiments, while section 4 concludes the paper.

## 2. The FLOG Mapping Method

### 2.1. Geometrical Considerations

The easiest way to use stereo vision for map building is by reducing its data immediately from three to two dimensions, e.g. by either using only one scan-line or by constructing a scan-line by selecting the closest point in every column of the disparity image [8]. However, this restricts possible movements of the camera to two dimensions as the scan-line is expected to represent measurements parallel to the ground. Furthermore, the latter variant only works due to the fact that stereo vision can not detect distances on objects without texture and the ground is usually un-textured. Otherwise, the closest point on the ground would be constantly picked up as obstacle. This is clearly a problem if the ground is textured.

Modelling stereo vision three dimensionally [9], leads to very memory consuming three dimensional occupancy grids with cubes as cells. However, the vertical resolution (i.e. height above ground) can often be chosen much lower than in the other two dimensions, which saves very much memory. Figure 2 depicts this case. Grid cells are modelled as boxes. This will be called a *layered occupancy grid*, to distinguish it from three dimensional grids with cubic cells. The height of each layer can be chosen application dependent. Applications that do not require explicit knowledge about the height of obstacles may choose only one big layer.




Figure 2: The layered occupancy grid.

The disadvantage of introducing an asymmetry in the size of cells is that the orientation of the grid needs to be aligned with the environment. Camera positions are determined relative to each other. The transformation from the first camera position into the world coordinate system can be defined by manually aligning the first camera position with the world coordinate system (e.g. optical axis of the

first camera is parallel to the ground and the horizon is horizontal in the image).

The stereo data is modelled by reconstructing the ray of light that passes through each pixel in the pinhole camera model, i.e. from the optical centre of the camera through every pixel up to the measured distance of that pixel (figure 2). All cells through which the ray passes are modified. It is important to consider the difference to other sensors. The usually modelled ultrasonic sensors emit very wide beams. Each cell that is inside this beam can be considered as empty. In contrast, the width of a ray is very small in comparison to the size of a cell. Thus, if a ray passes through a cell it can only be concluded that a small part of the cell is empty.

The position  $\vec{p} = [XYZ]^T$  in camera coordinates can be reconstructed from corresponding points  $x_l, y_l$  and  $x_r, y_r$  in the left and right image plane using equation (1).

$$X = \frac{x_l t}{x_l - x_r} \quad Y = \frac{t(y_l + y_r)}{2(x_l - x_r)} \quad Z = \frac{f t}{x_l - x_r} \quad (1)$$

The disparity is  $d = x_l - x_r$  (i.e. assuming rectified images with epipolar lines matching scan-lines). The focal length is  $f$  and  $t$  is the baseline. Camera motion data is used to transform  $\vec{p}$  into the world coordinate system  $\vec{q} = R\vec{p} + \vec{T}$ . The ray goes from  $\vec{T}$  to  $\vec{q}$ , since  $\vec{T}$  is the position of the optical centre in world coordinates.

Important are the error characteristics of stereo that can be derived from these equations using error propagation [2]. Small errors  $\Delta e$  in the position of the corresponding points in the image plane propagate into errors in the position  $\vec{p}$  as,

$$\Delta X = \Delta e \frac{Z}{f t} \sqrt{(t - X)^2 + X^2}, \quad (2)$$

$$\Delta Y = \Delta e \frac{Z}{f t} \sqrt{2Y^2 + \frac{1}{2}t^2}, \quad (3)$$

$$\Delta Z = \Delta e \frac{Z^2}{f t} \sqrt{2}. \quad (4)$$

The errors  $\Delta X$  and  $\Delta Y$  increase linear with distance, while  $\Delta Z$  increases squared. Typically,  $\Delta Z$  is 2 magnitudes higher than  $\Delta X$  and  $\Delta Y$  for moderate distances (e.g. a few meters) and influences several cells in contrast to  $\Delta X$  and  $\Delta Y$ . Thus, rays can be modelled one dimensionally. For simplicity of calculation, the error  $\Delta l$  in the distance (i.e. in the direction of sight)  $l = \sqrt{X^2 + Y^2 + Z^2}$  is approximated with equation (5). This is justified with the dominance of  $\Delta Z$  and  $Z$  on cameras without wide angle lenses.

$$\Delta l = \Delta e \frac{l^2}{f t} \sqrt{2}. \quad (5)$$

## 2.2. Updating of Cell Memberships

Rays are modelled one dimensionally as explained in section 2.1, with an uncertainty in the direction of the ray. Matthies and Shafer [13] argued that the uncertainty around the measured distance  $l$  can be modelled with a Gaussian function. This is approximated by a triangular function, due to computational speed (i.e. there have to be thousands of rays calculated for each stereo image). The maximum error  $3\Delta l$  is calculated with equation (5) with a standard deviation in the image plane of  $\Delta e$  of 0.1 pixels. Thus, the membership functions for emptiness and occupancy for ray  $i$  can be formalised with parameter  $r$  as distance from the optical centre (i.e. positive in the direction of sight).

$$\mu_{E_i}(r) = \begin{cases} 1 & \text{if } r \geq 0 \text{ and } r < l - 3\Delta l, \\ \frac{l-r}{3\Delta l} & \text{if } r \geq l - 3\Delta l \text{ and } r < l, \\ 0 & \text{otherwise.} \end{cases} \quad (6)$$

$$\mu_{O_i}(r) = \begin{cases} 1 - \frac{l-r}{3\Delta l} & \text{if } r \geq l - 3\Delta l \text{ and } r < l, \\ 1 - \frac{r-l}{3\Delta l} & \text{if } r \geq l \text{ and } r < l + 3\Delta l, \\ 0 & \text{otherwise.} \end{cases} \quad (7)$$

The membership functions are visualised in figure 3. These functions are mapped into the occupancy grid by considering only cells  $C$  through which the ray passes and determining for each cell the parameter  $r$  that is closest to the middle of the cell.




Figure 3: Used Memberships functions.

The sensor data is in contrast to the approach of Oriolo et al. [12] directly fused into the global occupancy grid. Oriolo et al. demonstrated the advantage of using an averaging calculation instead of a union operator so that memberships of cells do not only increase but also decrease. This is important to compensate errors with new data and account for partly dynamic environments. Therefore, this choice has been assimilated. The memberships to the *empty* ( $\mu_E(C)$ ) and *occupied* fuzzy set ( $\mu_O(C)$ ) of cell  $C$  are updated from ray  $i$  that passes through cell  $C$  by,

$$\mu_E^*(C) = \frac{(N(C) - 1)\mu_E(C) + \mu_{E_i}(C)}{N(C)}, \quad (8)$$

$$\mu_O^*(C) = \frac{(N(C) - 1)\mu_O(C) + \mu_{O_i}(C)}{N(C)}. \quad (9)$$

The number of rays through cell  $C$  is counted in  $N(C)$ . An advantage of this formulation as Oriolo et al. [12] point out is that  $N(C)$  can be limited using a maximum threshold  $N_f$ . This effectively limits the influence of past measurements and allows to adapt the map quickly to changes in the environment from new data.

An extension for stereo vision is the calculation of the membership of cell  $C$  to the *confidence* fuzzy set  $K$  using the number of measurements in a cell,

$$\mu_K(C) = \frac{N(C)}{N_K}. \quad (10)$$

Parameter  $N_K$  represents the number of rays that are required to give full confidence in the correctness of  $\mu_E$  and  $\mu_O$ . This formulation of confidence coincides nicely with the behaviour of stereo vision. The number of rays that can pass through a cell increase squared with decreasing distance of the cell from the camera. The same is true for the accuracy of stereo as discussed in section 2.1. For this study,  $N_K$  has been chosen as the number of rays that pass through a fully visible cell observed at closest possible distance. The threshold for counter  $N_i$  has been set to  $3N_K$ .

There are two issues which require consideration. Firstly, the calculations of emptiness, occupancy and confidence are only correct for fully visible cells. Partly visible cells introduce errors, because the measurements of a part of a cell do not necessarily reflect measurements of the whole cell. Smaller cells are less likely to be partly visible, but they lead to an increased memory consumption. The main importance of this issue is to choose the height of layers so that they are not too big, i.e. a camera that is close to a cell should still be able to observe the cell completely. Secondly, multiple stereo images from the same camera position do not contain additional information, but often the same systematic errors, e.g. increased size of objects due to object border problems. Therefore, as a heuristic, new views are only added if the translation between the last added view is at least 4 grid cells or if the rotation is more than half of the field of view (i.e.  $21^\circ$  for the current cameras).

It is important to note the slight change in meaning of empty and occupied memberships of a cell. For ultrasonic sensors with a beam that is much larger than cells, the memberships adjust to the *believe* that a cell is completely empty or partly occupied. In stereo vision the memberships reflect the *ratio* to which a cell is empty or occupied. This has an impact on processing of these values in section 2.3.

### 2.3. Creation of a Visual Map

One strength of the fuzzy logic approach is that maps can be defined for different purposes. Oriolo et al. define a map that is *safe for planning* and one that is *safe for motion* [12]. The difference is the treatment of undetermined areas. The former map treats them as being empty, while the latter one as being occupied. For the visualisation of the map to the operator, undetermined areas will be treated as if they are empty. Calculation of the visual map is derived from the approach of Oriolo et al. Differences are that it does not make sense to calculate a measure of *ambiguity* for cells, since cells that are partly empty and partly occupied are explicitly modelled (section 2.2). Furthermore, the measure

of *certainty* is new. Finally, multiple layers of two dimensional grids are used, that have to be fused.

The map value is calculated for each layer separately. First, a measure of *indeterminateness*  $I$  is calculated as not empty and not occupied. Next, the map value  $M_k$  for layer  $k$  is defined.

$$M_k = \overline{(E \cap \bar{O}) \cup I} \cap K \quad I = \bar{E} \cap \bar{O} \quad (11)$$

Layers could be represented using different colours to visualise different heights. They can also be fused if height information is not important, i.e.

$$M = \cup_k M_k. \quad (12)$$

The resulting value  $M$ , which is between 0 and 1 is mapped to intensities between white and black for visualisation. It is important to note that the fusion of several layers is superior to using only one layer, because increasing the height of cells introduces errors if cells are often only partly observed, as discussed at the end of section 2.2.

### 2.4. Considerations about the Implementation

A strength of occupancy grids is that their memory consumption can be easily predicted and can be kept constant, because it depends only on the size of the environment and the chosen size of grid cells. However, this is controversial as the size of the visited environment must be known in advance. Otherwise, time expensive reallocation and memory copy operations are required for extending the size. Furthermore, non-rectangular areas are either difficult to model or memory is wasted by defining the grid rectangular.

A good compromise for all of these problems is to model the grid as consisting out of tiles. Each tile has a fixed size and holds all cells of a squared part of the grid, for all layers. Figure 4 visualises the situation. A table of pointers is used as an index structure. Tiles are allocated as soon as cells inside the area that they cover are accessed. This data structure allows to overestimate the size of the visited environment without wasting memory or to adapt the size efficiently, since only the table of pointers needs to be allocated or adapted.




Figure 4: Used flexible data structure of the occupancy grid.

Furthermore, the speed to access individual grid cells can be greatly increased by choosing the size of tiles to be a power of 2 (e.g.  $64 \times 64$ ). Thus, only shift or mask operations are needed to address tiles and cells inside tiles.

Otherwise, multiplications and divisions would be needed, which are much more time consuming on most processors.

### 3. Experimental Results

#### 3.1. Mapping of a Simple Scene

A simple experiment has been designed to test the main characteristics of mapping. The setup is shown in figure 5. The camera has been circled in 32 steps around the centre of the big box. Only one map layer has been used for simplicity, which covers the height of the bigger box. The size of each grid cell has been set to  $10 \times 10 \times 420$  mm. Thus, the space over and under the box (e.g. the table) are excluded from the map.




Figure 5: Setup of the first experiment with two boxes (left) and one of the 32 recorded stereo images (right).

The previously developed stereo vision and camera motion estimation methods have been used to create disparity images and determine camera positions. It seems important to use the segmentation filter to remove small disparity patches, which represent usually outliers (section 1.3). Camera motion has been determined, without any pre-knowledge about the movement, exclusively from consecutive stereo images.

The left image in figure 6 shows a map that was created from close distance, i.e. the distance from the camera to centre of the big box was constantly  $1.1m$ . Camera motion has been determined from the last to the first stereo images as well to calculate the overall error. The camera was only  $12mm$  and  $1.3^\circ$  degrees off its correct position and orientation. The right image in figure 6 shows mapping of the same objects but from greater distance, i.e. between  $1.6m$  and  $2.0m$  distance from centre of the big box. The camera was in this case after closing the cycle  $63mm$  and  $2.6^\circ$  degrees away from its correct position and orientation.

It can be seen from both maps that two sides of the smaller box are very weak, which is due to obstruction from the bigger box. Interesting is how the sides of the boxes are represented thicker when the distance of the camera increases. This is a direct result of equation (5), which is used to create the membership functions (i.e. the uncertainty in distance  $\Delta l$  increases squared with distance  $l$  to an obstacle). Thus, grid cells around occupied cells are also declared as *possibly* occupied in absence of more accurate close range data.




Figure 6: Circling at close distance (left) and bigger distance (right) around two boxes.

#### 3.2. Mapping of Real Scenes

Mapping has been performed in 3 different indoor and one outdoor environment. Some experiments used an absolutely arbitrary three dimensional camera movement (i.e. including rotation around the optical axis). Figure 7 shows a typical example. The simplified, manually created map is seen on the left. The raw stereo sequence is 50s long and was taken by hand with general camera movements. The cell size of the grid has been set to  $20 \times 20 \times 700$  mm with 3 layers. The automatically created map can be seen on the right in figure 7. The boundaries of objects appear more fuzzy than in the simple example, because the camera is usually further away from objects. However, mapping seems suitable to give a human operator an overview over an unknown environment, i.e. obstacles and free space can be identified very well.

Incremental map building has been done on a 1.2GHz Athlon in real-time, including rectification, stereo correlation, camera motion estimation and map updating with  $\approx 7.8$ fps, which corresponds to 390 processed stereo views. However, updating of the map has been done from only 112 selected views by skipping all views that were less than  $100mm$  in translation (i.e. 4 cell sizes) and  $21^\circ$  degrees in orientation (i.e. halve field of view) away from the previous mapped view. Mapping of one view took 59ms, even though only every second disparity row was considered for mapping. However, this corresponds to just 17ms per processed stereo view, which seems acceptable. Speed can be further increased by choosing bigger grid cells or using only every n-th row of the disparity image, which has little effect on the created map.

### 4. Conclusion

A new method has been described to incrementally build occupancy grid based maps from sequences of stereo images in real time. Advantageous are that height information can be maintained or certain height regions selected. Thus, excluding unimportant or even disturbing information. Furthermore, movements of the stereo camera are

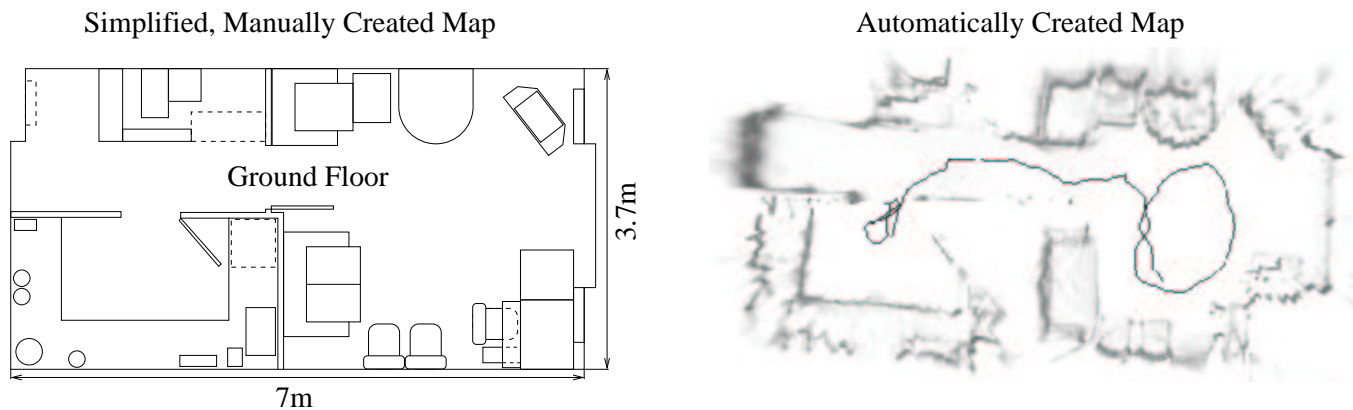


Figure 7: The true layout of a room (left) and the generated real-time map with camera path (right).

unconstrained in three dimensions, which opens the way to use maps for more, possibly even non-robotics applications.

Furthermore, problems that arise from the traditional probabilistic approach (section 1.2) have been addressed. Firstly, it is difficult to create accurate probabilistic sensor models for complex sensors like stereo vision. However, section 2.1 explained a straight forward approach for fuzzy logic. Secondly, the probabilistic approach is unable to distinguish between unknown and uncertain areas, while the fuzzy logic solution offers a much richer description. The visual map is created according to rules (e.g. treats unknown areas as empty or occupied, etc) and not by constraints. Thirdly, the independence of scans is accounted for by enforcing different camera positions or orientations. Finally, other implicit assumptions like the even distribution of occupied cells do not exist and thus can not be violated.

The general impression is that the incremental FLOG mapping method seems suitable to provide a human operator with an overview over an unknown environment and permits to identify obstacles and free space. Thus, it can be used to assist navigation of a tele-operated mobile robot and virtual walkthroughs, which are generated from the same stereo and motion data.

## Acknowledgements

I would like to thank QinetiQ for their financial support and Dr. Peter R. Innocent and Dr. Robert I. John from De Montfort University for their valuable advice.

## References

- [1] H. Hirschmüller, P. R. Innocent, and J. M. Garibaldi, "Real-time correlation-based stereo vision with reduced border errors," *International Journal of Computer Vision*, vol. 47, pp. 229–246, April–June 2002.
- [2] H. Hirschmüller, P. R. Innocent, and J. M. Garibaldi, "Fast, unconstrained camera motion estimation from stereo without tracking and robust statistics," in *Seventh International Conference on Control, Automation, Robotics and Vision*, (Singapore), pp. 1099–1104, 2–5 December 2002.
- [3] J. Vandorpe, H. Van Brussel, and H. Xu, "Exact dynamic map building for a mobile robot using geometrical primitives produced by a 2d range finder," in *Proceedings of IEEE International Conference on Robotics and Automation*, (Minneapolis, Minnesota), pp. 901–908, April 1996.
- [4] Z. Zhang and O. Faugeras, "A 3d world model builder with a mobile robot," *International Journal of Robotics Research*, vol. 11, pp. 269–285, August 1992.
- [5] J. Gasós and A. Martín, "Mobile robot localization using fuzzy maps," in *Fuzzy Logic in AI — Selected papers from the IJCAI 95 Workshop* (T. Martin and A. Ralescu, eds.), pp. 207–224, Springer-Verlag, 1997.
- [6] A. Elfes and H. P. Moravec, "High resolution maps from wide angle sonar," in *Proceedings of the IEEE International Conference on Robotics and Automation*, (St. Louis, MO), pp. 116–121, 1985.
- [7] A. Elfes, "Using occupancy grids for mobile robot perception and navigation," *Computer*, vol. 22, pp. 46–57, June 1989.
- [8] D. Murray and J. Little, "Using real-time stereo vision for mobile robot navigation," *Autonomous Robots*, vol. 8, no. 2, pp. 161–171, 2000.
- [9] A. P. Tirumalai, B. G. Schunk, and R. C. Jain, "Evidential reasoning for building environment maps," *IEEE Transactions on System, Man and Cybernetics*, vol. 25, pp. 10–20, January 1995.
- [10] K. Konolige, "Improved occupancy grids for map building," *Autonomous Robots*, vol. 4, no. 4, pp. 351–367, 1997.
- [11] F. Gambino, G. Oriolo, and G. Ulivi, "A comparison of three uncertainty calculus techniques for ultrasonic map building," in *Proceedings of the 1996 SPIE International Symposium on Aerospace/Defense Sensing and Control*, (Orlando, Florida, USA), pp. 249–260, 1996.
- [12] G. Oriolo, G. Ulivi, and M. Vendittelli, "Real-time map building and navigation for autonomous robots in unknown environments," *IEEE Transactions on Systems, Man, and Cybernetics*, vol. 28, no. 3, pp. 316–333, 1998.
- [13] L. Matthies and S. A. Shafer, "Error modeling in stereo navigation," *IEEE Journal on Robotics and Automation*, vol. 3, pp. 239–248, June 1987.