THREE-DIMENSIONAL MAP BUILDING WITH MMW RADAR

Alex Foessel-Bunting, John Bares, William Whittaker
The Robotics Institute, Carnegie Mellon University, Pittsburgh, PA 15213, USA
Phone +1 (412) 681-9527, Fax +1 (412) 681-6961, Email {afoessel, bares, red}@ri.cmu.edu

Abstract: Radar offers advantages as a robotic perception modality because it is not as vulnerable to the vacuum, dust, fog, rain, snow and light conditions found in construction, mining, agricultural and planetary-exploration environments. However radar has shortcomings such as a large footprint, sidelobes, specularity effects and limited range resolution—all of which result in poor environment maps. The fusion of successive radar observations can alleviate radar shortcomings and improve map fidelity. Sensor models exist for the fusion of sonar, laser and stereo into evidence grids. However radar outputs richer data and cannot use those existing models. This paper presents a sensor model that uses constant false alarm for occupancy detection and incorporates heuristic rules to approach occlusions. The resulting radar-based map of outdoors can suit robot obstacle avoidance, navigation and tool deployment. The limited map detail achieved suggests the need for a more rigorous probabilistic approach to encode the dependencies and estimate the model parameters.

Keywords: Millimeter-wave radar, evidence grids, robot sensing, mobile robots, FMCW radar.

1 INTRODUCTION

Current research on capable robotic vehicles focuses on the mining, earth-moving and agricultural industries as well as on planetary-exploration applications. Imaging sensors provide obstacle avoidance, task-specific target detection and generation of terrain maps for navigation. In addition the sensors provide a means of detecting humans, animals and vehicles entering the workspace. Visibility conditions are often poor in field-robotic scenarios. Day/night cycles change illumination conditions. Weather phenomena such as fog, rain, snow and hail impede visual perception. Dust clouds rise in excavation sites, agricultural fields and planetary exploration. Also smoke compromises visibility in fire emergencies and battlefield operations.

Laser and stereo are common visually based sensors affected by these conditions [Vandapel et al., 1999]. The sizes of dust particles, fog droplets and snowflakes are comparable to the wavelength of visual light so clouds of particles block and disperse the laser beams impeding perception. Stereo depends on the texture of objects and on an illumination source. Sonar is a common sensor not affected by visibility restrictions. However sonar suffers from a limited maximum range, poor angular resolution, and reflections by specular surfaces; sonar is of limited utility for field robots.

Millimeter-wave radar provides consistent range measurements for the environmental imaging needed to perform autonomous operations in dusty, foggy, blizzard-blinding and poorly lit environments [Foessel-Bunting et al., 1999]. Radar overcomes the shortcomings of laser, stereo and sonar. In addition radar can provide information of distributed targets and of multiple targets that appear in a single observation. However current developments of short-range radar imaging are not satisfactory. High angular resolution can only be obtained with inconveniently large antenna apertures, and downrange resolution has hardware limitations. Signal amplitude varies significantly with object material and surface orientation and wide beams result in large footprints.

Evidence grids have been useful for the integration of uncertain and noisy sensor information. The method has fused laser, stereo or sonar, as well as combinations of those sensors. The evidence grid framework uses sensor models to capture the interaction of the sensor with the environment. However radar propagation and scattering are different from those of optical-based sensors thus making existing sensor models inadequate. In addition, frequency-modulated continuous-wave (FMCW) radar outputs downrange-amplitude arrays, which difficult the adoption of existing single input sonar-sensor models designed. The development of a radar-sensor model for building three-dimensional perception models from radar motivates this research.

This paper presents development of an FMCW-radar sensor model that can be used to build three-dimensional evidence-grid maps. The next section revises related work in the areas of radar robot perception and utilization of evidence grids for map building. Section 3 explores radar phenomena and detection methods, proposes heuristic rules for signal-amplitude interpretation, and presents a sensor model derived from the rules. Section 4 describes the experiments and presents the results of outdoor scene representation from simulated and real data. Section 5 discusses the results and investigates advantages and shortcomings of the approach. Section 6 concludes this work with summary and suggestions of future research.

2 RELATED WORK

Limited precedence exists for applying millimeter-wave radar for short-range three-dimensional terrain mapping. For example, previous work by Lange and Detlefsen [1991] shows indoors imaging at frequencies of 35 GHz and 94 GHz. Processing speed limited real-time imaging to only two-dimensions. The approach divides the area of interest in cells containing a binary value that represents the radar

evidence of occupation in the area. Another research effort presents the implementation of radar-based obstacle avoidance on large mining trucks [League and Lay, 1996]. A technique of aging cells collects the information of sequential frames, thus reinforcing real targets and rejecting noise. The radar, which mounts on the front of the truck, detects persons, obstacles and the sides of dirt roads. In another development, a millimeter-wave radar-based navigation system detects and matches artificial beacons for localization in a two-dimensional scan [Clark and Durran-Whyte, 1997]. The publication indicates that an improved method would eliminate the need for an infrastructure through recognition of features and map building during navigation. Pulsed radar with a narrow beam and high-sampling rate produces dense three-dimensional terrain maps [Boehmke et al., 1998]. However, the resulting sensor size is excessive for most robotic applications and the pulsed radar does not deal with distributed or multiple targets.

Previous and ongoing research show successful application of evidence grids for building terrain and scene maps. This approach divides the space of interest in regular cells and accumulates observations from sonar [Moravec and Elfes, 1985], stereo and laser sensing [Yamauchi et al. 1998]. Each cell stores the accumulated evidence of occupancy for the corresponding area or volume as provided by the sensor observations. Although most researchers utilize two-dimensional grids, there are three-dimensional implementations to fuse stereo data for terrain [Moravec, 1996] and underwater mapping [Stewart, 1996]. Evidence-grid extensions use more than one value per cell to include confidence values or evidence of other hypotheses. Other implementations lift restrictive assumptions and consider the existence of multiple objects [Yamauchi et al. 1998] or include surface orientation hypotheses [Konolige, 1997]. However most common evidence grid implementations use single value per cell: the logarithm of the odds of occupancy [Martin and Moravec, 1996]. Most of the aforementioned evidence-grid implementations take as input a singles range measurement. Some cases of underwater mapping use a sensor model that accounts for multiple and distributed targets contained in a single measurement [Rosenblum and Kamgar-Parsi, 1992].

Robotic tasks such as planning, navigation and collision avoidance can use evidence-grid representation directly.

3 APPROACH AND METHODOLOGY

3.1 Sensor Model for Range Arrays

Robotic tasks such as navigation, obstacle avoidance and tool deployment require knowledge about the presence of all solid objects in the workspace. Task execution most times requires an existing map of the workspace thus imposing time constraints to the sensing and perception. Rapid map building benefits from using all the information provided by the sensors. However typical sensors suffer from occlusion and only provide range to the first object in the propagation axis. The corresponding sensor models, necessary for accumulating information in a world representation, accept only such single value. This is the case of sensor models for sonar, laser and stereo range readings.

But radar does not suffer of occlusion as much and each radar observation can provide information of multiple or distributed objects. Trees, poles and other small radar-cross-section objects in the foreground that do not occlude the background are examples of multiple-object observation. Low-grazing sensing of surfaces with radar mounted on vehicles is a common example of distributed-object observation. The adoption of simple techniques such as first pulse detection would discard the rich information that radar can provide. The use of the totality of the radar output array motivates the investigation a sensor model that accepts the radar output array.

3.2 Radar Sensing

Radar transmits electromagnetic radiation in a narrow beam by means of an antenna. The radiated waves propagate along the beam axis and objects reflect, refract and absorb portions of the incidence energy. The remaining portion of the radiated energy propagates past the object until encounters another object. This process continues downrange until all the radiation is absorbed, back scattered or reflected.

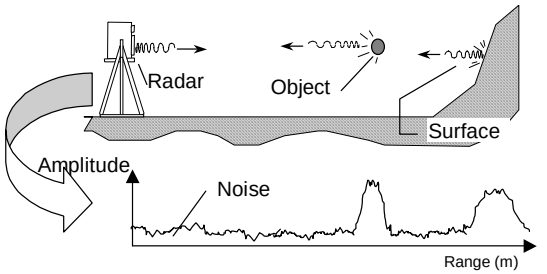

Figure 1. Radar radiates and receives the back scatter of objects along the beam. The output is an array with the energy received from the corresponding range.

The radar receiver converts the returning energy into an array of amplitude values. The amplitude in each cell is proportional to the energy back scattered by one or more objects located at the corresponding range and illuminated by the beam. The energy returned is a function of the objects radar cross section. Amplitude values above the noise level suggest the presence of objects with significant radar cross section. Similarly, amplitude values close or below the noise level generally correspond to the absence of objects, but exceptions exist. The lack of a signal can correspond to a specular reflecting surface, to a highly absorbing material, or to total occlusion of radiation. Occlusions vary the amount of energy that illuminates an object and the back scatter that reaches the antenna. Radar can suffer total occlusion

caused by large objects as well from unnoticeable perturbations caused by point-like objects.

Other factors that affect the radar output are range propagation attenuation, noise, sidelobes and other sources of radiation. Although the propagation attenuates the energy received by the radar, a filter corrects the output so that similar objects produce similar signal strengths at different distances. Electronic noise, ambience radar radiation and frequency-sweep nonlinearity contribute with noise. The study assumes an exponential distribution of the noise power. Sidelobes, specular reflections and other radar sources can introduce false targets. This research assumes that the radar antenna does not have strong sidelobes and that no specular surfaces or other nearby radar sources exist. Under these assumptions, pulses above the noise level imply the presence of objects at the corresponding range with certainty.

The number of elements of the radar output array relates to the radar maximum range and to the range resolution. There are techniques to improve the range resolution; summary of such techniques appears in [Foessel-Bunting, 2000].

The following subsection considers this discussion of the radar sensing phenomena and derives qualitative rules to convert the radar output array into an array of evidence of occupancy.

3.3 Heuristic Rules

Empirical rules can help the interpretation of the radar-output array. Radar observations in typical mobile-robots environments contain three main features: narrow pulses, wide pulses, and pulse absence or noise segments. Narrow pulses correspond to point-like objects or to surfaces normal to the beam axis. Point-like objects generally do not occlude more distant objects, surfaces generally do. A wide pulse typically represents a surface with high incident angle. Long objects aligned with the beam can also appear as wide pulses by they likely appear as a narrow pulse corresponding to the end closer to the sensor.

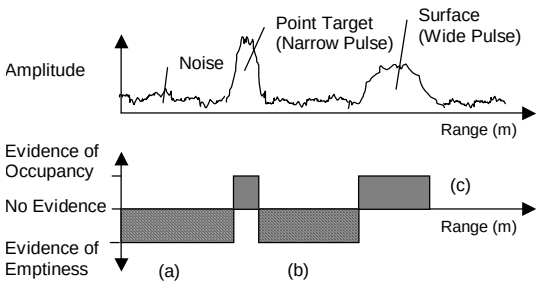

Figure 2. Evidence of emptiness is strong between pulses and very weak after a wide pulse.

Under the assumptions of this study—no sidelobes and no other sources of radiation—all pulses indicate object presence. But the absence of noise can result from the absence of objects or to the absence of radiation due to occlusion. This dependency changes the interpretation of the noise segments based on the features found closer to the sensor.

A surface captures most of the radiated energy and occludes more distant objects. Groups of nearby objects result in wide pulses and also have a similar occlusion effect. Thus wide pulses *generally* occlude distant objects; only noise follows wide pulses downrange because of the lack of radiation reaching distant zones. Lack of radiation does not provide information and the evidence of empty space after a wide pulse is very weak, as shown in Figure 2 (c).

Noise between pulses strongly suggests empty space, as shown in Figure 2 (b). Low values between echoes do imply emptiness. If the radar signal shows more than one echo, all the space between the sensor and the first echo, and all the spaces between echoes provide strong evidence of emptiness. The existence of an echo discards the possibility of occlusion.

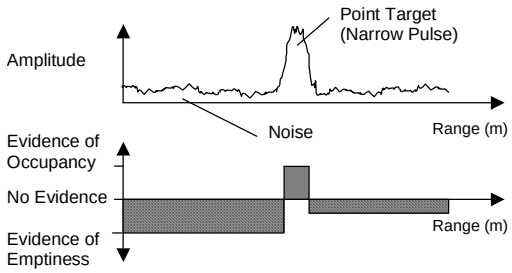

Figure 3. Evidence of emptiness after a narrow pulse is weaker because of the possible occlusion produced by a surface normal to the beam axis.

Noise after a point target implies emptiness because a point target does not completely occlude downrange objects. This rule is of dubious application because a wall perpendicular to the sensor will also appear as a point target. Therefore the evidence of emptiness after a narrow pulse is weak. Figure 3 illustrates this case.

3.4 Discrimination of Pulses and Noise

Radar-output-array interpretation and the implementation of empirical rules require the identification of pulses and the estimation of pulse width. Radar research has developed techniques for pulse detection with variable noise and signal conditions. Constant false-alarm rate (CFAR) is a technique that estimates the noise level near the test range bin to control the detection threshold. This results in an approximately constant rate of false pulse detections. One implementation is cell-averaging CFAR in which the average of the range-bin amplitudes controls the threshold. This noise estimation assumes that the noise distribution is exponential and that the selected threshold level results in a constant false alarm rate. The averaging window calculates the mean value of several range bins for characterization of the noise exponential distribution and for the result a noise estimation n_i. A table indi-

cates the threshold T_{CFAR} for a given false-alarm rate. The test window compares the threshold with the amplitude of the signal s_i to classify the range bin as pulse or noise and assigns a signal-to-noise ratio value. The technique infrequently fails to classify the signal with probability $1–P_d$ because the pulse with noise added can fall below the threshold.

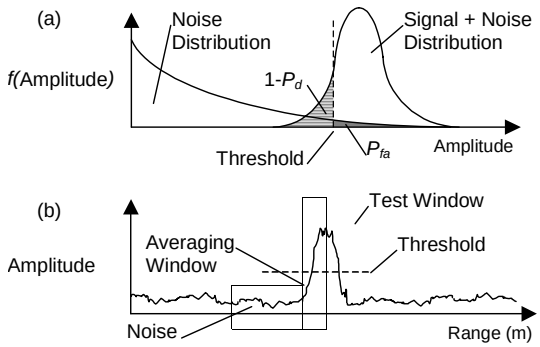

Figure 4. Constant false alarm rate provides a threshold for pulse detection. The graph (a) shows threshold selection in the distribution over the amplitude. Graph (b) shows the technique applied to classify a range bin as a pulse or noise.

The CFAR threshold classifies the signal as *noise* or *pulse* and the signal-to-noise ratio provides *weight* for the evidence of occupancy. Intuitively, a higher signal-to-noise provides stronger evidence of occupancy for a given cell. The sensor model adopts a logarithmic relation between the signal-to-noise ration and occupancy evidence. Noise level however does not weigh the evidence of emptiness.

$$\log\lambda(r_i|M)=\begin{cases} k_{occ}\log\left(\frac{r_i}{n_i}\right), \text{ if } \left(\frac{r_i}{n_i}>CFAR\right) \\ E_h, \text{ if } \left(\frac{r_i}{n_i}>CFAR\right)\wedge \text{before wide pulse} \\ E_l, \text{ if } \left(\frac{r_i}{n_i}\le CFAR\right)\wedge \text{after wide pulse} \end{cases}$$

Figure 5. The sensor-model down-range interpretation rules. E_h and E_l are negative values.

Under the assumptions of this study, signal noise relates only to the transceiver internal noise and does not provide information about the environment. The aforementioned empirical rules condition the evidence of emptiness to the position of the range-bin in relation to wide pulses. A pulse is considered *wide* if the extension at the CFAR threshold is equal to or more than the extension of the distributed echo of flat terrain at the lowest elevation of the radar (–14º for the outdoors scene data gathering).

The sensor model assigns a constant emptiness-evidence level to range-bins classified as noise: E_h for range bins between the sensor and the first wide pulse, and E_l for those behind. This paper refers to this function as the "down-range interpreter." The representation of the interpreter output and the beam geometry is referred as "sensor model." This model is represented by the relation in Figure 5.

3.5 Sensor-model angular uncertainty

Radar wide beams and the resulting large footprint cause angular uncertainty. A radar pulse at a given range can result from an object anywhere in the footprint at that range. The graph on the right of Figure 6 shows the two-way signal strength of a constant radar-cross-section object as a function of the angle measured from the beam center. An intuitive interpretation considers the normalized curve as the probability distribution of the object cross-range position given a pulse at a given range. The constant radar cross-section assumption appears unrealistic in an environment with a diversity of objects. However this study considers no prior knowledge of the objects and no other assumption can be made.

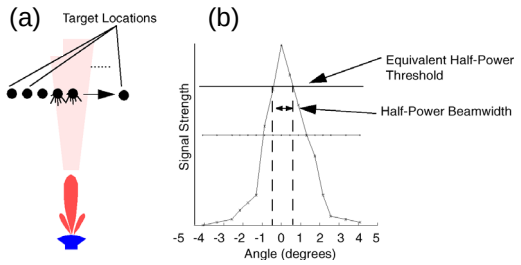

Figure 6. Experimental estimation of beamwidth. Scheme (a) presents the experimental configuration for beamwidth estimation. Graph (b) shows the amplitude of the pulse as the object is displaced cross range

4 EXPERIMENTS AND RESULTS

The experimental work focused first on the evaluation of the downrange resolution and on the characterization of radar for the development of a sensor model. Later the investigation gathered observations of an outdoor scene to demonstrate the fusion of several radar observations to build a map. Figure 8 shows such outdoors scene.

4.1 The Sensor

This investigation used a mechanically scanned millimeter-wave radar designed for obstacle avoidance and road navigation in heavy-dust conditions. The sensor is a 77-GHz FMCW millimeter-wave radar. The wavelength is $\lambda = 3.8\text{mm}$. The beamwidth is 2° in elevation and 1° in azimuth. The bistatic antenna scans horizontally four stacked beams across the angular range of 64°. The four stacked beams cover an elevation range of 8°. One scanning cycle—or radar frame—covers a 64° by 8° area of the scene.

A stand with a tilt axis allows the sensor to aim at different elevation angles. The sensor reports the amplitude of echoes at ranges between 1 and 64 m. The signal sampling provides amplitude measurements every 0.5 m.

4.2 Downrange Interpretation

Figure 7(a) shows the radar-output array of a beam that aims with an elevation of –1º, almost parallel to the ground. The pulses correspond to a fence and a tree and there is no ground return. The interpreter outputs strong emptiness evidence (–1) in all those cells that do not belong to a pulse, as shown in inset (b). The function detects the pulses and calculates the logarithm of the signal-to-noise ratio. This value is close to a value of +2, indicating very strong evidence of occupancy. The graph corresponds to the expected interpreter results.

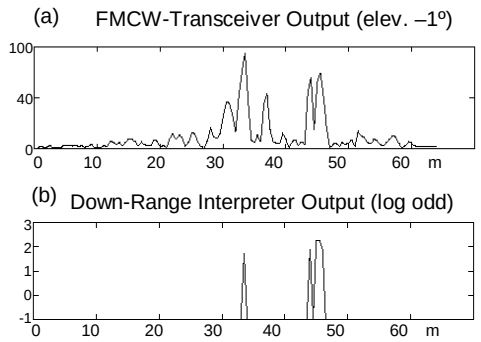

Figure 7. The downrange interpretation rules transform a radar observation into evidence.

4.3 Outdoor scene representation

The test scene contains a variety of objects, some on the fairly flat surface covered with grass; trees, a metal fence, a group of metal bars, and a 55-gallon drum.

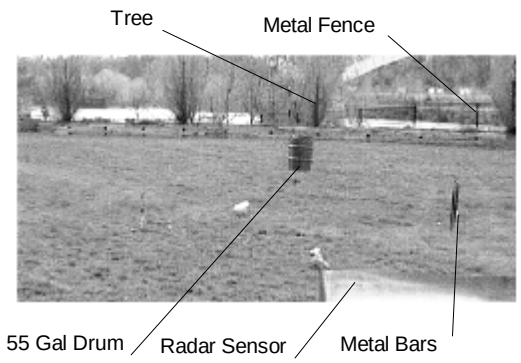

Figure 8. An outdoors scene to demonstrate radar observation fusion with the use of the radar sensor model developed in this work.

The data set consists of observations taken over 15 elevations between –0.2º to –14.2º at intervals of one degree. For each elevation, the four-stacked beams scan 64 positions in azimuth with a total of 3,840 radar observations. The map is a three-dimensional grid 80-m wide, 70-m long and 20-m high. The cells are cubes with 0.1-m sides. The sensor position is such that the ground plane is in the middle of the height of the grid.

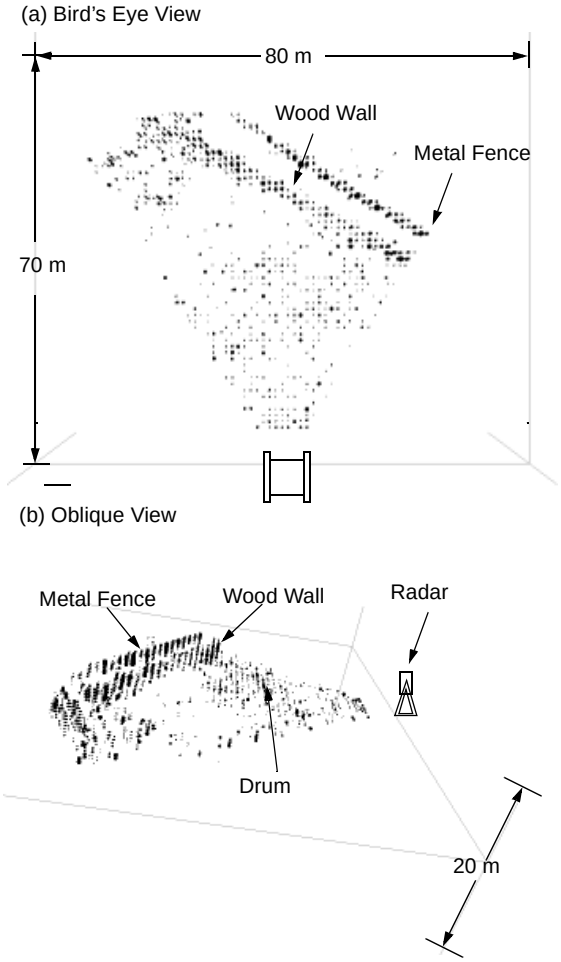

Figure 9. Two perspectives of the grid-based representation of an outdoor scene. The display technique thresholds the grid cells to positive values and represents the accumulated value with a sphere with radius proportional to the grid value.

Figure 9 shows two views of the integration of 960 observations. The inset (a) corresponds to a bird's eye view. The inset (b) shows the same grid from a left and up point of view. The grid visualization consists of displaying a small sphere for each cell that contains odds of occupancy. The sphere radius is proportional to the grid value.

5 DISCUSSION

The visualization shows that the map has a correct geometry and that objects with large radar cross section accumulate the most evidence. Holes in the ground plane suggest that the thresholds are eliminating the weak signals that result from grazing

angles. This indicates reduced interpreter performance in front of weak signals. The performance in the classification depends heavily on the CFAR parameters. Obtaining reasonable representation of the ground surface guided the selection of the CFAR parameters in this study. A more rigorous method for parameter selection is necessary to obtain consistent results. One undesirable consequence of the current parameter selection follows.

Although the resulting map shows the main features in the scene it fails to show smaller objects. The sensor model has several parameters to set the threshold, measure pulse width and for assigning evidence. The parameter selection is not satisfactory because other distinct features, such as a group of metal bars, did not appear in the grid-based representation. Non-optimal choices can explain the poor performance of the map to show smaller objects.

Similarly, the heuristic rules are somewhat arbitrary. The qualitative rules for a downrange signal interpretation account for occlusions—phenomenon that causes strong downrange dependencies. However the rules are dubious and difficult to implement. The evidence weighs given to noise and to pulses should be function of noise and pulse distribution parameters.

6 CONCLUSION

This work develops a radar-sensor model for evidence grids and demonstrates the integration of imaging radar observations to build three-dimensional maps of outdoors. The appreciation of the fidelity of the map is qualitative at this stage. A sensor model provides a means to transform range information and paste it into a grid-based representations.

The sensor model assumes the absence of sidelobes, of specular reflections, and of other radar sources. It does, however, consider the effect of occlusions and attempts to represent variations of evidence when known occlusions exist. A constant false-alarm rate technique classifies pulses above the noise level to classify evidence of emptiness and occupancy.

The experimental results indicate that the combination of qualitative rules and CFAR threshold does classify the signal and provides a means to transform the range-amplitude vector in evidence of occupancy or emptiness. The accumulation of radar observations in a three-dimensional evidence grid demonstrates the effectiveness of the approach.

The model has several shortcomings that indicate several issues worth of investigation. A rigorous probabilistic formulation is necessary; the author currently investigates Bayes networks as a means to represent the dependencies due to occlusion. Experimental estimations of the noise distributions can yield optimal parameter selection.

Finally, a reference scene is essential to evaluate the goodness of new methods and parameters. Learning methods would also benefit from a reference scene.

ACKNOWLEDGMENTS

The Motion-Free Scanning-Radar Alliance at the National Robotics Engineering Consortium, Carnegie Mellon University, sponsored this study.

REFERENCES

[Boehmke et al., 1998] S. Boehmke, J. Bares, E. Mutschler and K. Lay, "A High Speed 3d Radar Scanner for Automation," Proc. IEEE Int. Conf. on Robotics and Automation, Leuven, Belgium, 1998.

[Clark and Durran-Whyte, 1997] S. Clark and H. Durran-Whyte, “The Design of a High Performance MMW Radar System for Autonomous Land Vehicle Navigation,” Proc. Int. Conf. Field and Service Robotics, Ed. Zelinsky, Sydney, Australia, 1997, pp 292–299.

[Foessel-Bunting, 2000] A. Foessel-Bunting, "Radar Sensor Model for Three Dimensional Map Building," Proc. SPIE, Mobile Robots XV and Telemanipulator and Telepresence Technologies VII, SPIE, Vol. 4195, November 2000.

[Foessel-Bunting et al., 1999] A. Foessel-Bunting, S. Chheda and D. Apostolopoulos, "Short-Range Millimeter-Wave Radar Perception in a Polar Environment," Proc. Field and Service Robotics Conference, Pittsburgh PA, 1999.

[Konolige, 1997] K. Konolige, "Improved Occupancy Grids for Map Building," Autonomous Robots, vol. 4, 1996.

[Lange and Detlefsen, 1991] M. Lange and J. Detlefsen, “94 GHz Three-Dimensional Imaging Radar Sensor for Autonomous Vehicles,” IEEE Trans. on Microwave Theory and Techniques, vol. 39, no. 5, 1991.

[League and Lay, 1996] R. League and N. Lay, "System and Method for Tracking Objects Using a Detection System," US Patent 5,587,929, 1996.

[Moravec, 1996] H. Moravec, “Robot Spatial Perception by Stereoscopic Vision and 3D, Evidence Grids,” Technical Report CMU-RI-TR-96-34, Robotics Institute, Carnegie Mellon University, Pittsburgh, 1996.

[Rosenblum and Kamgar-Parsi, 1992] L. Rosenblum and B. Kamgar-Parsi, "3D reconstruction of small underwater objects using high-resolution sonar data," Proc. Sym. Autonomous Underwater Vehicle Technology, 1992.

[Stewart, 1996] W. K. Stewart, "Three-Dimensional Stochastic Modeling Using Sonar Sensing for Undersea Robotics," Autonomous Robots vol. 3, no. 2, 1996.

[Vandapel et al., 1999] N. Vandapel, S. Moorehead, W. Whittaker, R. Chatila and R. Murrieta-Cid, “Preliminary Results on the Use of Stereo, Color Cameras and Laser Sensors in Antarctica,” Proc. Int. Sym. on Experimental Robotics, 1999.

[Yamauchi et al. 1998] B. Yamauchi, A. Schultz and W. Adams, "Mobile robot exploration and map-building with continuous localization," Proc. IEEE Int. Conf. on Robotics and Automation, Leuven, Belgium, 1998.