

although interpolation techniques have been applied to improve it for point targets. This makes it difficult to generate elevation maps because objects of different heights are illuminated at the same time, and it prevents the use of geometric or shape algorithms, such as those commonly used with lasers. In general, the “alternate” image of the scene provided by the radar may be difficult to interpret because its modality, resolution, and perspective are very different from those of visual images.

Furthermore, radar propagation to some extent and particularly scattering are different from those of optical-based sensors, such as laser, stereo, or sonar; thus existing

sensor models are inadequate. For example, laser data typically return the range to the first target detected along the beam, although last pulse-based lasers solve this problem to some extent and are becoming more common. In contrast, radar outputs power-downrange arrays, i.e., a single beam contains information from multiple targets mainly due to the wider beamwidth (about 2–3 deg compared with 0.1 deg for the laser). A single sensor sweep, therefore, outputs data containing n samples at discrete range increments dR along each azimuth or scan angle. As an example, Figure 1(a) shows a bidimensional intensity graph of the radar data (radar image) acquired from a large, relatively




Figure 1. A sample radar image acquired from a large flat area: azimuth angle-range image (polar coordinates, radar frame) (a), same radar image in Cartesian coordinates (radar frame) (b), and camera image approximately collocated with the radar (c). Note the rich information content of the radar map due to its ability to sample reflectivity at multiple ranges for a single scan angle. Please refer to the online version of the paper for a color view of the radar image.

flat area [Figure 1(c)]. The abscissas in Figure 1(a) represent the horizontal scanning angle. The ordinates represent the range measured by the sensor. Amplitude values above the noise level suggest the presence of objects with significant reflectivity. Amplitude close to or below the noise level generally corresponds to the absence of objects, but exceptions exist. These include specular reflecting surface aligned to reflect the signal away, a highly absorbing material, or a total occlusion of radiation.

One interesting feature of the radar image is the ground echo, i.e., the intensity return scattered back from the portion of terrain that is illuminated by the sensor beam. In the presence of relatively flat terrain, the ground echo appears as a high-intensity parabolic sector in the radar image [see Figure 1(a)]. This sector is referred to as the radar image background throughout the paper. The ability to automatically identify and extract radar data pertaining to the ground and project them onto the vehicle body frame or navigation frame would result in an enabling technology for all visibility-condition navigation systems. In this research, a theoretical model describing the geometric and intensity properties of the ground echo in radar images is described. It serves as a basis for the development of a novel method for radar ground segmentation (RGS), which allows classification of observed ground returns in three broad categories, namely ground, nonground, and unknown. The RGS system also improves the accuracy in range estimation of the detected ground for enhanced environment mapping. Persistent ground segmentation is critical for a robot to improve perception in natural terrain under all visibility conditions, with many important applications including scene interpretation, environment classification, and autonomous navigation, as ground is often taken as the most likely traversable terrain. It should be noted that this research focuses specifically on the analysis of the radar image background. The full radar data can also be used to detect objects present in the foreground. However, this is out of the scope of the paper and it is not expressly addressed here.

In this investigation, a mechanically scanned MMW radar, designed for perception and navigation in low-visibility conditions, is employed. The sensor, custom built at the Australian Centre for Field Robotics (ACFR), is a 95-GHz frequency-modulated continuous wave (FMCW) MMW radar that reports the amplitude of echoes at ranges between 1 and 120 m. The wavelength is $\lambda = 3$ mm, and the 3-dB beamwidth is about 3.0 deg in elevation and azimuth. The antenna scans across the angular range of 360 deg. The important technical properties of the sensor are collected in Table I.

For the extensive testing of the system during its development, the CAS Outdoor Research Demonstrator (CORD) was employed, which is an eight-wheel, skid-steering all-terrain unmanned ground vehicle (UGV) [see Figure 2(a)]. In Figure 2(b), the radar is visible, mounted to a frame

Table I. Radar technical properties.

Property	Value
Model	ACFR custom-built
Max. range (m)	120
Raw range resolution (m)	0.25
FOV (deg)	
Horizontal	360
Instantaneous	3.0×3.0
Angle scan rate (rps)	≈ 3
Chirp period (ms)	2



(a)



(b)

Figure 2. The CORD UGV employed in this research (a) and its sensor suite (b).

attached to the vehicle's body and tilted forward so that the center of the beam intersects the ground at a look-ahead distance of about 11.4 m in front of the vehicle. The robot is also equipped with other sensors, including four 2D SICK laser range scanners, a mono-charge-coupled device (CCD) color camera, a thermal infrared camera, and a real-time kinematic/differential global positioning system/inertial navigation system (RTK DGPS/INS) unit that provides accurate position and tilt estimation of the vehicle during the experiments.

The paper is organized as follows. A description of related literature is provided in Section 2, and basic principles of radar sensing are recalled in Section 3. The model of ground echo and the RGS method are described in detail in Sections 4 and 5, respectively. In Section 6, the RGS system is proved to be effective and robust to visibility conditions in field tests performed with the CORD UGV. Section 7 concludes this paper.

2. RELATED WORK

Persistent navigation is one of the basic problems for autonomous mobile robots. The extraction of information about the surrounding environment from sensory data to build up a map, while at the same time keeping track of its current position, in what is usually referred to as simultaneous localization and mapping (SLAM) (Dissanayake, Newman, Clark, Durrant-Whyte, & Csorba, 2001), is a key component for a mobile robot to accomplish the assigned task. Hence, fast and reliable algorithms capable of extracting features from a large set of noisy data are critical in such applications, especially for outdoor environments. Specifically, detection and segmentation of the ground in a sensor-generated image is a challenging problem with many applications in perception. This is a key requirement for scene interpretation, segmentation, and classification (Douillard, Underwood, Kuntz, Vlaskine, Quadros, et al., 2011), and it is important for autonomous navigation (Reina, Ishigami, Nagatani, & Yoshida, 2010). For the purposes of obstacle detection and avoidance, current systems usually rely on ranging sensors such as vision, laser, or radar to survey the three-dimensional (3D) shape of the terrain. Some features of the terrain, including slope, roughness, and discontinuities, are analyzed to segment the traversable regions from the obstacles (Pagnot & Grandjea, 1995; Singh, Simmons, Smith, Stentz, Verma, et al., 2000). In addition, some visual cues, such as color, shape, and height above the ground, were employed for segmentation in DeSouza and Kak (2002) and Jochem, Pomerleau, and Thorpe (1995). Other methods used motion cue of stereo images (Ohnishi & Imiya, 2006; Zhou & Baoxin, 2006). The ground is detected and segmented by computing a disparity map, i.e., the point correspondences between two images. The image motion induced by the ground can be characterized by a homography, which differs from that induced by obstacles. Thus, the ground pixels and obstacle pixels can be identified by evaluating whether their displacements are consistent or inconsistent with the motion specified by the estimated homograph.

In addition, the ground region is often assumed to be larger than that of obstacles, which facilitates the estimation of the homography. These methods work well on indoor images in which the scene is basically composed of floor (ground) and planes perpendicular to the floor (nonground), as the difference of the image motions between the ground and nonground is substantial in

such scenarios. However, this approach cannot work successfully for outdoor environments in which the surface normal changes smoothly in many cases, causing only a slight difference in image motion between ground and nonground. An alternative approach was proposed using the appearance information via some learning techniques for ground segmentation (Pomerleau, 1989). Various types of visual cues including color and texture have been used for detection of ground (Ulrich & Nourbakhsh, 2000). Other information such as corners and edges has also been exploited in some applications (Poppinga, Birk, & Pathak, 2008; Vosselman & Dijkman, 2001).

Relatively limited research has been devoted to investigate MMW radar for short-range perception and 3D terrain mapping. For example, previous work presented the implementation of radar-based obstacle avoidance on large mining trucks (League & Lay, 1996). In other work, a MMW radar-based navigation system detected and matched artificial beacons for localization in a 2D scan (Clark & Durrant-Whyte, 1997). Pulsed radar with a narrow beam and high sampling rate produced dense 3D terrain maps (Boehmke, Bares, Mutschler, & Lay, 1998). However, the resulting sensor size is excessive for most robotic applications. MMW radar has been used on a large autonomous guided vehicle for cargo handling (Durrant-Whyte, 2002); the radar is scanned horizontally and measures range and bearing to a set of trihedral aluminum reflectors. The reflectors may be covered by a polarizing grating to enable discrimination from other objects. Radar capability was demonstrated in a polar environment (Foessel-Bunting, Chheda, & Apostolopoulos, 1999) and for mining applications (Brooker, Hennesy, Lobsey, Bishop, & Widzyk-Capehart, 2007). Mullane, Adams, and Wijesoma (2009) used a MMW radar for occupancy mapping within a probabilistic framework.

A body of research also exists in the automotive community related to road and obstacle detection in radar images. These works relied on constant (Kaliyaperumal, Lakshmanan, & Kluge, 2001) or adaptive (Jiang, Wu, Wu, & Sun, 2001) thresholding but achieved only marginal performance on good paved roads and are unsuitable for off-highway driving.

3. RADAR BACKGROUND

The working principle of the radar used in this research is based on a continuous-wave (CW) signal that is modulated in frequency to produce a linear chirp that is radiated toward a target through an antenna. The echo received from a target nanoseconds later by the same antenna is mixed with a portion of the transmitted signal to produce a beat signal at a frequency that will be proportional to the round-trip time. A spectrum analyzer can then be used to produce an amplitude-range profile that represents this target as a spectral peak. Any other targets within the antenna beam

would simultaneously produce their own peaks, giving the radar a true multitarget detection capability from a single observation (Brooker, Hennessey, Bishop, Lobsey, Durrant-Whyte, et al., 2006). For robot perception, an accurate range map of the environment can be constructed through the scanning of a narrow beam, which is usually referred to as a pencil beam. MMW radar provides a pencil beam with relatively small antenna apertures. The beam width is proportional to the wavelength and is inversely proportional to the antenna aperture. A constant uniformly illuminated antenna aperture shapes narrower beams at shorter wavelengths according to Brooker (2005):

$$\theta_e = 1.02 \frac{\lambda}{D}. \tag{1}$$

Equation (1) relates the elevation beamwidth θ_e as a function of the wavelength λ and the antenna aperture D . For example, at 95 GHz and a wavelength of 3 mm, a 1.17-deg beam results from a 150-mm antenna aperture. In reality, the aperture is not uniformly illuminated, as that results in higher sidelobes, and the beamwidth resulting from a weighted illumination is significantly wider, typically greater than 1.5 deg. A narrower beam produces more accurate terrain maps and obstacle detection. However, antenna size is limited by robot size and spatial constraints. Most airborne radar applications sense targets in the antenna far-field region where the radiated power density varies inversely with the square of the distance and where the antenna pattern remains constant for each angle. The far-field region is commonly considered to begin approximately at a distance R_{\min} (Skolnik, 1981)

$$R_{\min} \simeq 2 \frac{D^2}{\lambda}. \tag{2}$$

Relation (2) results in $R_{\min} \simeq 16$ m for the radar used in this study. Therefore, UGV-based ground perception tasks require short-range sensing because most of the targets fall within the near-field region, where the antenna pattern is range dependent and the average energy density remains fairly constant at different distances from the antenna (Slater, 1991). Near-field effects do not prevent sensing close to the radar but suggest complex beam geometries and more difficulty in data interpretation.

In this research, the radar is directed at the front of the vehicle with a constant forward pitch to produce a grazing angle β of about 11 deg in order to survey the environment, as shown in Figure 3. The origin of the beam at the center of the antenna is O. The proximal and distal borders of the footprint area illuminated by the divergence beam are denoted with A and B, respectively. The height of the beam origin with respect to the ground plane is h . The slant range of the radar bore sight is R_0 , and the range to the proximal and distal borders is denoted with R_1 and R_2 , respectively. Near-grazing angles stretch the pencil-beam footprint, resulting in range-echo spread, as illustrated in Figure 4. In principle, the computation of the area on the




Figure 3. Scheme of a pencil beam (of beamwidth θ_e) sensing terrain at grazing angle β .

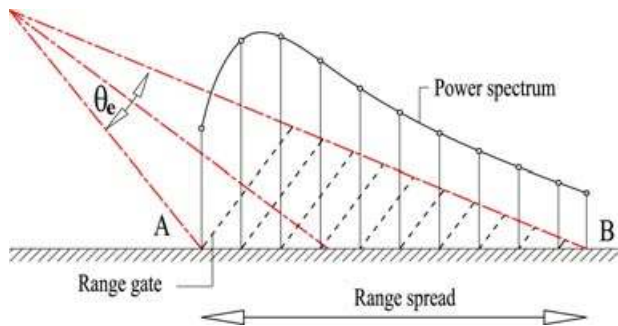


Figure 4. Expected power return from grazing angle perception.

ground surface that is instantaneously illuminated by the radar depends on the geometry of the radar boresight, elevation beamwidth, resolution, and angle of incidence to the local surface. Two different geometries can be defined: the pulse length-limited case and the beamwidth-limited case. The pulse-limited case is defined in which the so-defined radar's range resolution dR , projected onto the surface, is smaller than the range extent of the total illuminated area. This is the case of the ground echo, as explained in Figure 3. For this geometry, the instantaneously illuminated resolution cell is approximately rectangular. Conversely, the beam-limited case is defined for a geometry in which the radar range resolution projected onto the ground is larger than the range extent of the area illuminated by the radar beam. The resolution cell in this case is elliptical.

4. GROUND ECHO MODELING

A theoretical model of the ground echo in the radar image is developed. It provides prediction of the range spread of the ground return along with the expected power spectrum.

4.1. Ground Echo: Geometry

Based on the configuration of the radar onboard the vehicle, a good estimate of the expected range spread of the ground return can be obtained by modeling the radar as a conical pencil beam. For our radar system, the cone aperture is found experimentally as the half-power point. This results approximately in $\theta_e = 3$ deg. By intersecting the pencil beam with the ground plane, it is possible to obtain a prediction of the ground range spread as a function of the azimuth angle α and the tilt of the vehicle. It is assumed that the position of the radar relative to the vehicle frame

$$R_V^W = \begin{pmatrix} \cos \psi \cdot \cos \theta & \cos \psi \cdot \sin \theta \cdot \sin \phi - \sin \psi \cdot \cos \phi & \cos \psi \cdot \sin \theta \cdot \cos \phi + \sin \psi \cdot \sin \phi \\ \sin \psi \cdot \cos \theta & \sin \psi \cdot \sin \theta \cdot \sin \phi + \cos \psi \cdot \cos \phi & \sin \psi \cdot \sin \theta \cdot \cos \phi - \cos \psi \cdot \sin \phi \\ -\sin \theta & \cos \theta \cdot \sin \phi & \cos \theta \cdot \cos \phi \end{pmatrix}, \quad (4)$$

is known by initial calibration and fixed during travel (Underwood, Hill, Peynot, & Scheduling, 2010). With reference to Figure 5, five important reference frames can be defined: a world reference frame (WRF) $\{O_w, X_w Y_w Z_w\}$, a vehicle reference frame (VRF) $\{O_v, X_v Y_v Z_v\}$, a base radar reference frame (BRF) $\{O, X_{R0} Y_{R0} Z_{R0}\}$, and a radar reference frame (RRF) $\{O, X_R Y_R Z_R\}$. Each coordinate frame represents one step in the kinematic chain from the world to the RRF. The RRF is characterized by the X_R axis instantaneously coincident with the sensor boresight. During the radar sweep, the RRF will be rotated of the scan angle α around the Z_{R0} axis of the BRF. In general, a given point P of the environment can be defined in the RRF by the coordinates $P^R = (P_x^R, P_y^R, P_z^R)$. The same point will have coordinates in the WRF $P^W = (P_x^W, P_y^W, P_z^W)$ that can be determined as

$$P^W = R_R^W \cdot P^R + t_R^W, \quad (3)$$

where R_R^W is the rotation matrix of the RRF with respect to the WRF and t_R^W represents the coordinates of the origin O of the RRF in the WRF. R_R^W can be expressed as a composition of successive rotations along the kinematics chain:




Figure 5. Nomenclature for the reference frames.

$R_R^W = R_V^W \cdot R_{R0}^V \cdot R_R^{R0}$, R_i^{i-1} being the rotation from a coordinate frame i to a previous frame $i - 1$. Without loss of generality, we can simplify Eq. (3) under the assumption of BRF perfectly aligned with the VRF, i.e., $R_{R0}^V = I_3$. The rotation matrix R_V^W can be expressed in terms of a set of three independent angles known as Euler angles. In our implementation, we adopted the so-called ZYX Euler angles ϕ , θ , and ψ , usually referred to as roll, pitch, and yaw angles, respectively. The matrix R_R^{R0} describes the simple rotation of α around the Z_{R0} axis of the RRF with respect to the BRF:

$$R_R^{R0} = \begin{pmatrix} \cos \alpha & -\sin \alpha & 0 \\ \sin \alpha & \cos \alpha & 0 \\ 0 & 0 & 1 \end{pmatrix}. \quad (5)$$

We can now derive explicitly the third row of Eq. (3) as

$$P_z^W = r_{3,1} \cdot P_x^R + r_{3,2} \cdot P_y^R + r_{3,2} \cdot P_z^R (t_R^W)_z, \quad (6)$$

where $r_{i,j}$ are the components of the compound rotation matrix R_R^W . The intersection of the radar boresight with the ground that we assume as perfectly planar will have coordinates in the RRF $P^R = (R_0, 0, 0)$, R_0 being the slant range. The projection of P along the Z_w axis is null, i.e., $P_z^W = 0$, and $(t_R^W)_z = h$ is the z coordinate of the radar frame origin in the WRF. Replacing into Eq. (6), one gets the expected slant distance R_0 as a function of the azimuth angle and the tilt of the robot:

$$R_0 = \frac{h}{r_{3,1}} = \frac{h}{\cos \theta \cdot \sin \alpha \cdot \sin \phi - \sin \theta \cdot \cos \alpha}. \quad (7)$$

Similarly, the range of the proximal and distal borders R_1 and R_2 (points A and B in Figure 3) can be estimated by adding to the kinematic chain a final rotation of $\pm\theta_{el}$, where $\theta_{el} = \theta_e/2$, around the Y_R axis of the RRF,

$$R_1 = \frac{h}{\cos \theta_{el} (\cos \theta \sin \alpha \sin \phi - \sin \theta \cos \alpha) - \cos \theta \cos \phi \sin \theta_{el}}, \quad (8)$$

$$R_2 = \frac{h}{\cos \theta_{el} (\cos \theta \sin \alpha \sin \phi - \sin \theta \cos \alpha) + \cos \theta \cos \phi \sin \theta_{el}}. \quad (9)$$

In Figure 6, the ground echo spread as obtained by the geometric model is overlaid on the radar image using black dots. Figure 6 refers to the previous scenario considered in Figure 1 when the vehicle surveys a large area without

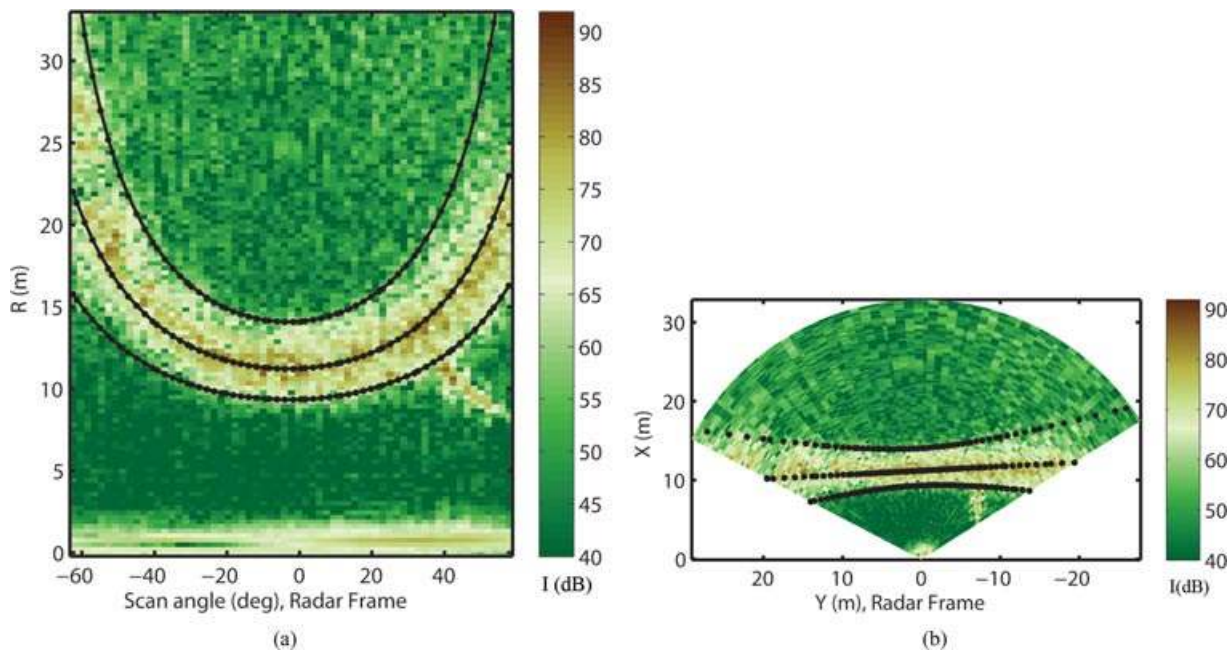


Figure 6. Ground echo spread obtained by the theoretical model (black dots) overlaid on the radar image expressed in polar (a) and Cartesian (b) coordinates. Note that the azimuth angle is defined positive when clockwise in the convention used in this research.

experiencing any significant tilt. The model matches the radar data very well. In Figure 7, a more complex radar image is shown in which the vehicle travels at a speed of 0.5 m/s with roll and pitch angles of 6.1 and 4.5 deg, respectively. The prediction of the model is again very good. For completeness, both radar images are also expressed in the Cartesian coordinates of the radar frame.

Thus, the geometric model provides a useful means to accurately predict the spread of the ground return under the assumption of globally planar ground. This, in turn, allows the definition of a region of interest in the radar image to improve ground segmentation.

4.2. Ground Echo: Power Spectrum

When a radar pulse is radiated into space, it will propagate at the speed of light, with most of the power constrained within a cone defined by the antenna beam pattern. If the beam intersects with an object, then some of the power will be scattered back toward the radar, where it can be detected. For a monostatic radar, in which the transmitter and receiver antenna are colocated, the received power P_r from a target at distance R can be defined by the radar equation (Brooker, 2005)

$$P_r = \left(\frac{P_t G}{4\pi R^2} \right) \cdot \left(\frac{\sigma^0 \Delta A}{4\pi R^2} \right) \cdot \left(\frac{G\lambda^2}{4\pi} \right), \tag{10}$$

where P_t is the transmitted power, G the antenna gain, λ the radar wavelength, σ^0 the backscatter coefficient (also known as the reflectivity), and ΔA the illuminated area. The first grouped term in the radar equation represents the power density (watts per square meter) that the radar transmitter produces at the target. This power density is intercepted by the target with radar cross section $\sigma^0 \Delta A$, which has units of area (square meters). Thus, the product of the first two terms represents the reflected power density at the radar receiver (again, watts per square meter). The receiver antenna then collects this power density with an effective area described by the third term. This interpretation of the radar equation will be useful later. The computation of the area ΔA , which is instantaneously illuminated, depends on the location of the radar onboard the vehicle and on the physical and geometric properties of the radar beam, as described in the preceding section (see Figure 3). Under the assumption of a RRF aligned with the ground, the length of the illuminated patch is given by

$$\overline{\Delta A} = \frac{dR}{\cos \beta}, \tag{11}$$

where dR is the so-defined range resolution (Brooker, et al., 2006) and β the grazing angle. It should be noted that the power return from the ground is extended over a range of grazing angles and that the backscatter coefficient and the gain depend on grazing angle. However, the variation

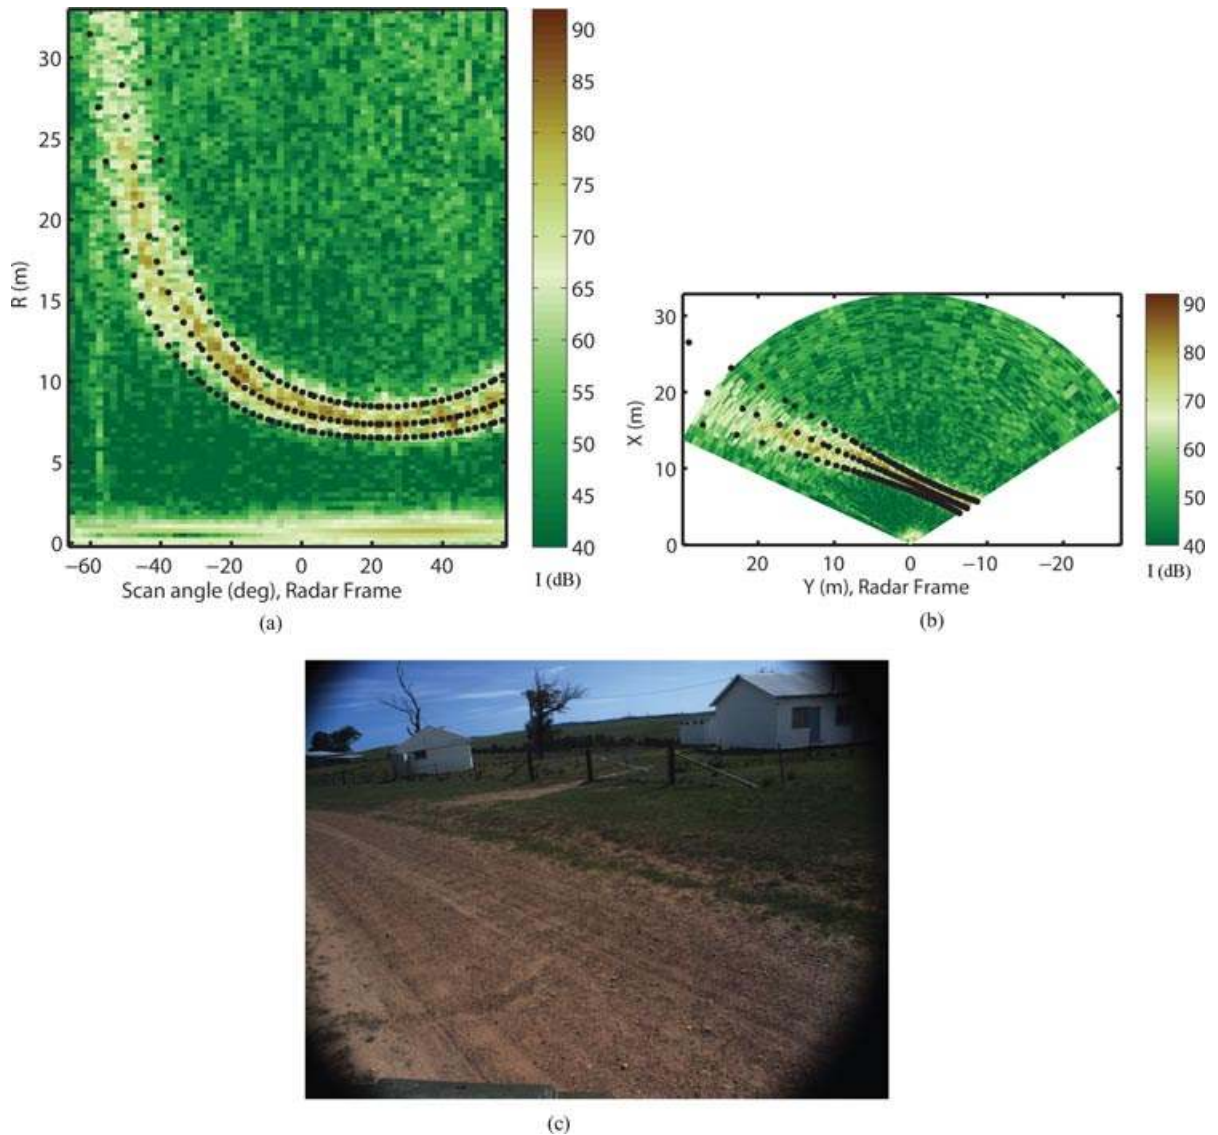


Figure 7. Ground echo spread obtained by the theoretical model (black dots) overlaid on the radar image expressed in polar (a) and Cartesian (b) coordinates, in the presence of significant vehicle tilt (roll and pitch angles of 6.1 and 4.5 deg, respectively). Camera image collocated with the radar (c).

range of the grazing angle is small (about 3 deg in our case), making the dependence on the backscatter coefficient weak. This assumption is based on the Georgia Tech land clutter model (Currie, Hayes, & Trebits, 1992, Chapter 3, p. 142) that for the configuration used in this research gives a backscatter coefficient falling in a plateau region with a little variation in σ^0 with grazing angle [less than 2 dB (m^2/m^2)]. Therefore, σ^0 is treated as a constant. The change in the gain across the beam must be taken into account. The gain is maximum when the target is located along the antenna's boresight. However, the gain reduces with angle off boresight, as defined by the antenna's radiation pattern. In

our case, a Gaussian antenna pattern can be adopted, as shown in Figure 8. If we introduce the elevation angle θ_{el} , measured from the radar bore sight and defined as

$$\theta_{\text{el}} = \arcsin\left(\frac{h}{R}\right) - \arcsin\left(\frac{h}{R_0}\right) \quad (12)$$

and denoting with $\theta_{3\text{dB}}$ the 3-dB beamwidth ($\theta_{3\text{dB}} = 3$ deg for our radar), then the antenna gain can be approximated by (Brooker, 2005)

$$G = e^{-2.776\left(\frac{\theta_{\text{el}}}{\theta_{3\text{dB}}}\right)^2}. \quad (13)$$




Figure 8. Gaussian antenna pattern.

It should be noted that Eq. (10) holds in the far field. However, the illuminated ground falls within the near-field region where the antenna pattern is range dependent and the average energy density remains fairly constant at different distances from the antenna (Slater, 1991). This allows one to simplify Eq. (10) as

$$P_r(R, R_0, k) = k \frac{G(R, R_0)^2}{\cos \beta}. \quad (14)$$

Equation (14) describes the power return of the ground echo in the radar image as a function of the range R . Note that for large tilt of the robot, the near-field region assumption may be violated and Eq. (14) would lose its validity. Figure 9 shows a simulated wide pulse of the ground return using Eq. (14). The model is defined by the two parameters k and R_0 that can be determined in practice by fitting the model to experimental data, as explained later. Finally, it is worth mentioning that for ground clutter, the sidelobe influence is in general negligible. However, if a target with a large radar cross section appears in sidelobes, it may produce significant effects.

In summary, Eqs. (8) and (14) represent two pieces of information defining the theoretical ground echo in the radar image. Any deviation in the spread or intensity shape suggests low likelihood of ground return in a given radar observation.

5. THE RADAR GROUND SEGMENTATION SYSTEM

A radar image can be thought of as composed of a foreground and a background. The background is referred to as the part of the image that contains reflections from the terrain. Radar observations belonging to the background show a wide pulse produced by the high-incident-angle surface, but exceptions may exist due to local unevenness or occlusion produced by obstacles of large cross sections




Figure 9. Simulated power return of the ground echo. The following parameters were adopted in the simulation: $k = 70$ dB, $R_0 = 11.3$ m, $h = 2.2$ m.

in the foreground. In this section, the RGS system is presented. It aims to assess ground by looking at the image background obtained by a MMW radar mounted on a mobile robot.

The RGS module performs two main tasks:

- Background extraction from the radar image
- Analysis of the power spectrum across the background to perform ground segmentation

In the remainder of this section, each stage is described in detail.

5.1. Background Extraction

A prediction of the range spread of the ground echo as a function of the azimuth angle and the tilt of the vehicle can be obtained using the geometric model presented in Section 4.1. It should be recalled that the model is based on the assumption of globally flat ground. Therefore, discrepancies in the radar observations may be produced by the presence of local irregularities or obstacles in the radar-illuminated area. To relax the assumption of global planarity and compensate for these effects, a change detection algorithm is applied in the vicinity of the model prediction. Specifically, the cumulative SUM (CUSUM) test is used, which is based on the cumulative sums charts to detect systematic changes over time in a measured stationary variable (Page, 1954). The CUSUM test is computationally very simple and intuitively easy to understand and can be motivated to be fairly robust to different types of changes (abrupt or incipient). In words, the CUSUM test looks at the prediction errors ϵ_t of the power intensity value. Under the assumption of normally distributed data, $\epsilon_t = (x_t - \bar{x}_t)/\sigma$, where x_t is the

last point monitored, \bar{x}_t the mean of the process, and σ the standard deviation. ϵ_t is a measure of the deviation of the observation from the target: the farther the observation is away from the target, the larger ϵ_t . The CUSUM test gives an alarm when the recent prediction errors have been sufficiently positive for a while. Mathematically, the test is formulated as the following time recursion:

$$\begin{aligned} g_0 &= 0, \\ g_t &= g_{t-1} + \epsilon_t - \nu, \\ g_t &= \max(0, g_t), \\ \text{if } g_t &> th, \text{ then alarm and } g_t &= 0, \end{aligned} \quad (15)$$

where ν and th are design parameters. The CUSUM test expressed by Eqs. (15) gives an alarm only if the power intensity increases. When negative changes need to be found as well, the min operation should be used instead, and this time a change has been found when the value of g_t is below the (negative) value of the threshold value th . The combination of two such CUSUM detectors is referred to as a two-sided CUSUM test. A typical result from the change detection algorithm is shown in Figure 10(a). The radar signal obtained from a single azimuth observation ($\alpha = 32$ deg) is denoted by a solid gray line. The theoretical prediction of the range spread of the ground return is shown by black points at the bottom of Figure 10(a) representing the range of the central beam, the proximal and distal borders, points R_0 , R_1 , and R_2 , respectively. When a positive change in the radar signal is found in the vicinity of the proximal border (in practice within a 1-m window centered on R_1), a flag is raised (dotted black line). The alarm is lowered when a negative change is detected in the vicinity of the distal border. The ground echo can then be extracted (portion of the signal denoted in black) from the given observation. The process can be repeated for the whole scanning range, and the background of the radar image can effectively be extracted, as shown in Figure 10(c) referring to the running example of Figure 1.

5.2. Ground Segmentation

The image background contains ground candidates. To define a degree of confidence in actual ground echo, the power return model, presented in Section 4.2, can be fitted to a given radar observation. The hypothesis is that a good match between the parametric model and the data attests to a high likelihood of ground. Conversely, a poor goodness of fit suggests low likelihood due, for example, to the presence of an obstacle or to irregular terrain. We recall that $P_r(R)$ is a function defined by the parameters R_0 and k . k can be interpreted as the power return at the slant range R_0 , and both parameters can be estimated by data fitting for the given azimuth angle. By continuously updating the parameters across the image background, the model

can be adjusted to local ground roughness and produce a more accurate estimation of R_0 , as shown later in the paper. A nonlinear least-squares approach using the Gauss–Newton–Marquardt method is adopted for data fitting (Seber & Wild, 1989). The initial parameter estimates are chosen as the maximum measured power value and the predicted range of the central beam as expressed by Eq. (7), respectively, limiting the problems of ill conditioning and divergence.

Output from the fitting process are the updated parameters \bar{R}_0 and \bar{k} as well as an estimate of the goodness of fit. The coefficient of efficiency was found to be well suited for this application (Nash & Sutcliffe, 2006):

$$E = 1 - \frac{\sum(t - y)^2}{\sum(t - \bar{t})^2}, \quad (16)$$

t being the data point, \bar{t} the mean of the observations, and y the output from the regression model. E ranges from $-\infty$ to 1, as the best possible value. E reaches 0 when the square of the differences between measured and estimated values is as large as the variability in the measured data. In case of negative E values, the measured mean is a better predictor than the model. By evaluating the coefficient of efficiency and the model parameters, ground segmentation can be effectively performed and radar observations can be labeled as ground, unknown, and nonground object (i.e., obstacle).

Two typical results are shown in Figure 11. Specifically, in Figure 11(a), the model matches very well the experimental data with a high coefficient of efficiency $E = 0.96$, thus attesting to the presence of ground. Conversely, Figure 11(b) shows an example in which the goodness of fit is poor ($E < 0$); in this case a low confidence in ground echo is associated with the given observation. In practice, a threshold Th_E is experimentally determined and the observation i is labeled as ground if E_i exceeds Th_E .

However, relying on the coefficient of efficiency only may be misleading in some cases. Figure 12(a) shows an example in which a ground patch would be seemingly detected according to the high coefficient of efficiency ($E = 0.91$), when there is actually no ground return. To solve this issue, a physical consistency check can be performed by looking at the updated value of the proximal and central range as estimated by the fitting process. For this case, they are almost coincident ($\bar{R}_0 = 10.82$ m and $\bar{R}_1 = 10.42$ m, respectively) and certainly not physically consistent with the model described in Section 4. Therefore, the radar observation is labeled as uncertain ground if the difference between the central and proximal range is lower than an experimentally defined threshold Th_R . An analogue comparison is done between the distal and central border as well.

In case of uncertain terrain, an additional check is performed to detect possible obstacles present in the region of interest, which would appear as narrow pulses of high intensity. In this respect, it should be noted that, during operation, the RGS system records the value of \bar{k} , defining

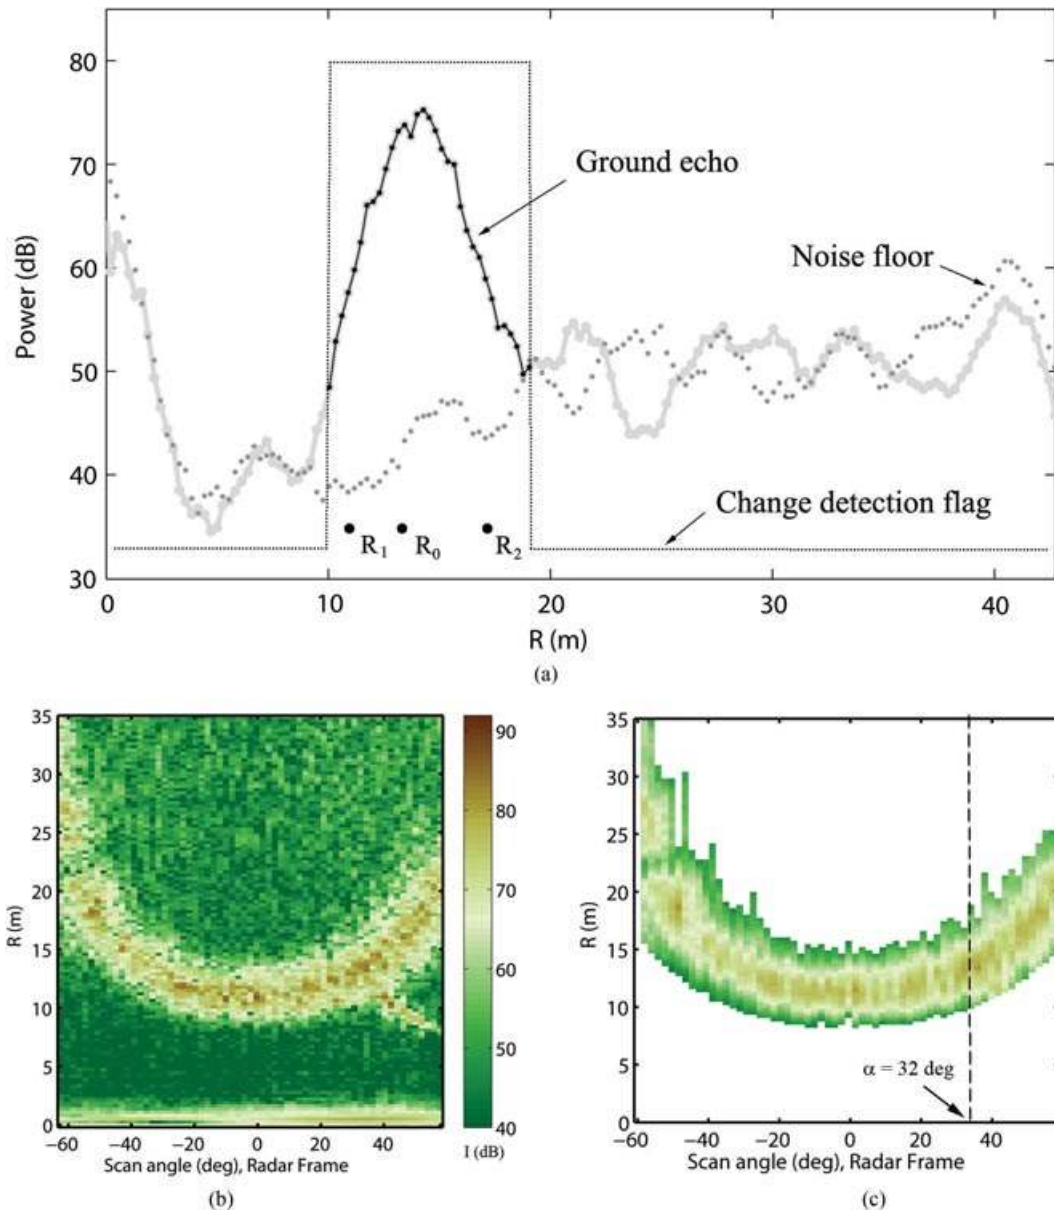


Figure 10. Ground echo extraction in the radar signal through a change detection approach (a): radar signal at scan angle $\alpha = 32$ deg (gray solid line), extracted ground echo (solid black line), change detection flag (dotted black line). Note that the opposite (i.e., 180-deg scan angle difference) radar signal is also plotted (gray dotted line); it points skyward, and no obstacle is detected in it, thus showing the typical noise floor in the radar measurement. Original radar image (b); radar image background (c).


a variation range for the ground-labeled observation. Typically, \bar{k} was found to range from 73 to 76 dB. If a percent relative change in the maximum intensity value between the uncertain-labeled observation t_{\max} and the model y_{\max} is defined, $\Delta P = (t_{\max} - y_{\max})/t_{\max}$, then an obstacle is detected when ΔP exceeds an experimentally defined threshold Th_p , and, at the same time, t_{\max} is greater than the maximum value of \bar{k} . An example of obstacle (labeled as nonground) detection is shown in Figure 12(b).

In summary, the classification approach described in Table II can be defined. Those rules express our physical understanding of the problem. The rule set is not unique; new rules may be thought of and implemented to improve the output of the system.

It should be noted that when the radar observation from a given azimuth angle is successfully labeled as ground or obstacle, an estimate of its range distance can also be obtained. In the case of ground, the whole range

Table II. Set of classification rules for the RGS method.

Class	Goodness of fit ($Th_E = 0.8$)	Parameters of the regression model ($Th_R = 1.5$ m, $Th_P = 10\%$)			
	E	$(\bar{R}_0 - \bar{R}_1)$	$(\bar{R}_2 - \bar{R}_0)$	ΔP	\bar{k} (dB)
Ground	$\geq Th_E$	$> Th_R$	$> Th_R$	$< Th_P$	73–76
Unknown	$< Th_E$	—	—	$< Th_P$	—
Non-ground	$< Th_E$	—	—	$\geq Th_P$	$> \bar{k}_{max}$


Figure 11. Ground segmentation by model fitting: good fit labeled as ground (a); poor fit labeled as uncertain ground (b).

spread of the ground echo can be theoretically evaluated along this azimuth angle (see Figure 4), thus providing a 3D point cloud with a single radar observation. However, for simplicity's sake, we refer only to the range estimation of the radar bore sight R_0 in this paper. If an obstacle is flagged instead, the RGS system outputs the range corresponding to the maximum intensity value of the detected pulse. Conversely, no range estimation is possible in the case of uncertain classification.

6. EXPERIMENTAL RESULTS

In this section, experimental results are presented to validate our approach for ground segmentation using a MMW radar. The RGS system was tested in the field using the CORD UGV (see Figure 2). The test field was located in a rural environment at the University of Sydney's test facility near Marulan, NSW, Australia. It was mainly composed of a relatively flat ground with sparse low grass delimited by fences, a static car, a trailer, and a metallic shed, as shown in Figure 13(a). During the experiment, the CORD vehicle was remotely controlled to follow an approximately closed-

loop path with an average travel speed of about 0.5 m/s and a maximum speed of 1.5 m/s. Variable yaw rates were achieved with a maximum of 1.12 rad/s (i.e., 64 deg/s) and roll and pitch angles of up to 5 deg. Along the path, the robot encountered a 40-cm trench and various slopes of medium–low inclination. In this experiment, the RTK DGPS/INS unit and a high-precision 2D SICK laser range scanner provided the ground truth with an average standard deviation per point of approximately 0.053 m [more details can be found in Underwood et al. (2010)]. The full data set is public and available online (Peynot, Scheduling, & Terho, 2010). The path followed by the robot is shown in Figure 13(b) as estimated by the onboard RTK DGPS-INS unit. It resulted in a total distance of 210 m traveled in about 6.5 min.

6.1. Ground Segmentation

Figures 14–16 show some typical results obtained during the experiment. Specifically, Figure 14 refers to the instant $T_1 = 34.1$ s [see Figure 13(b)] when the vehicle traversed a large, relatively flat area delimited by fences to the right




Figure 12. Ground segmentation by model fitting: seemingly high fit labeled as uncertain ground due to physics inconsistency with the model (a); narrow pulse labeled as obstacle (b).




Figure 13. The Marulan test field (a) and the path followed by the UGV during the experiment (b).

and a car to the far left. Figure 14(a) shows the radar intensity image overlaid with the results obtained from the RGS system. Ground labels are denoted by black dots (or by red dots in the online version of the paper), a black cross marks uncertain terrain, and nonground (i.e., obstacle) is denoted by a black (red) triangle. In Figure 14(b), the results are projected over the image plane of the camera for visualization purposes only. Finally, a comparison with the laser-based ground truth is provided in Figure 14(c), which demonstrates the effectiveness of the proposed approach for ground segmentation. As can be seen from these figures, the RGS system correctly detected the flat ground area in front of the robot and the obstacle to the left.

Figure 15 shows a different scenario at time $T_2 = 320$ s [see Figure 13(b)] with a stationary car to the right of the robot. The RGS method was successful in labeling the ground. Uncertain terrain was flagged along the portion of the background occluded by the car and to the far left due

to the presence of highly irregular terrain. Finally, at time $T_3 = 350$ s the robot faced a metallic shed with irregular terrain to the far left, as shown in Figure 16. In this scene for completeness, nonground labels detected in the foreground as high-intensity narrow pulses are also shown using black outlined triangles (or red outlined markers in the online version of the paper). The RGS module was correct in segmenting the ground to the left, indicating the presence of a large obstacle in front of the robot and of low confidence in ground to the right and far left of the scene.

Overall, the RGS system was tested over 1,100 radar images, each containing 63 azimuth observations for a total of 69,300 classifications. As a measure of the segmentation performance, the false-positive and false-negative rates incurred by the system during classification of ground and nonground were evaluated by comparison with the ground-truth laser data. To this aim, a previously proposed method for segmentation of laser data [GP-INSAC;

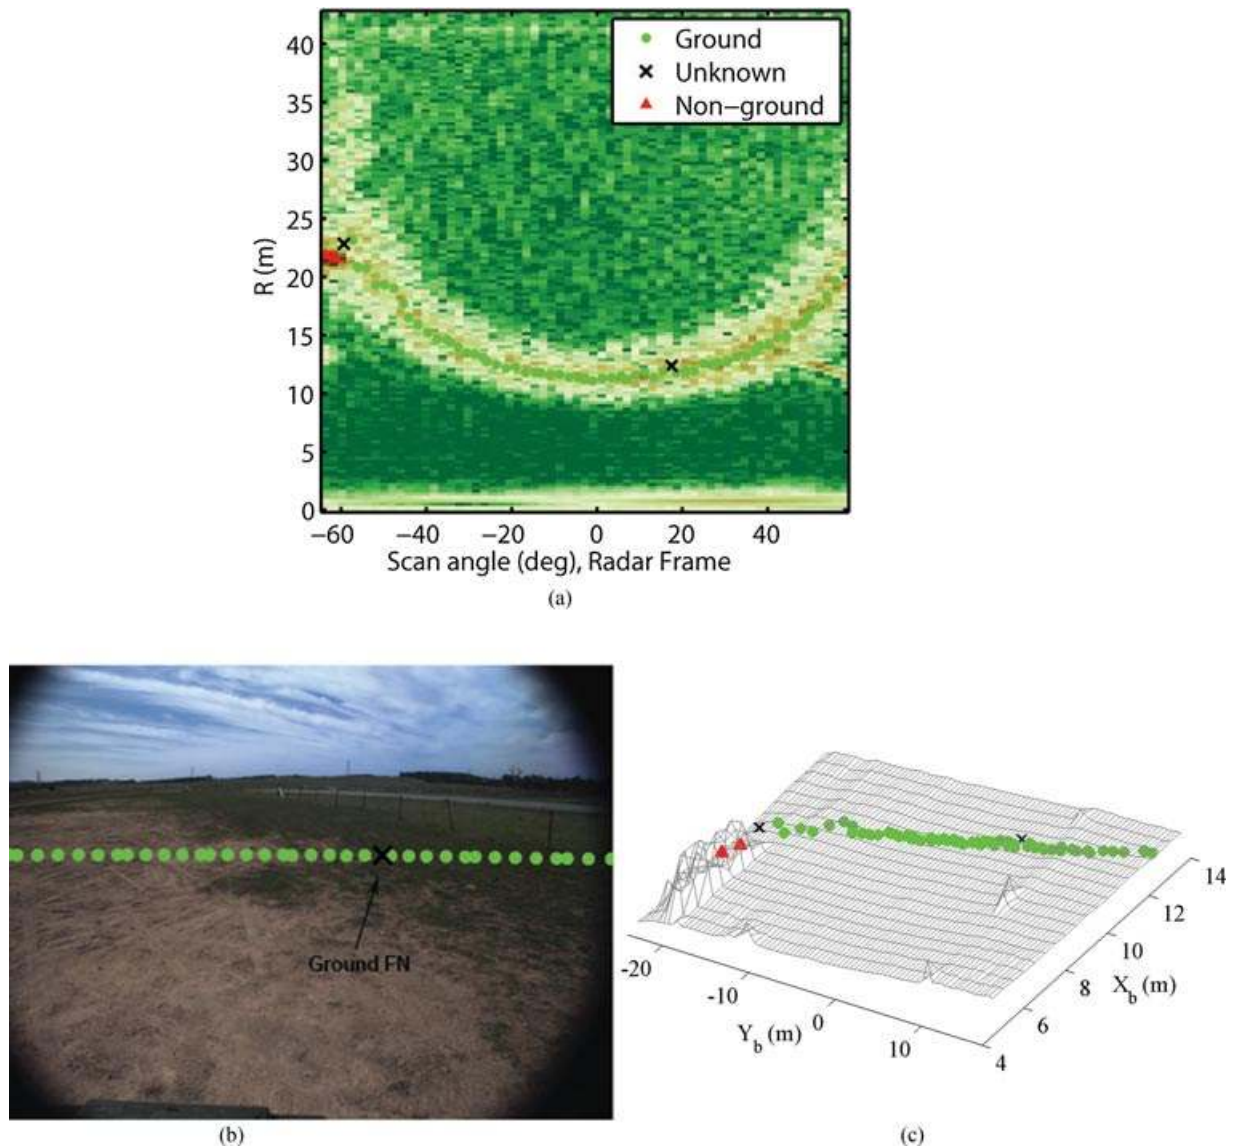


Figure 14. Results obtained from the RGS system for a relatively flat scenario. A large obstacle is detected to the far left: output of the RGS system (a); results overlaid on the camera image (b) and on the laser-generated ground-truth map (c). Note that the map is referred to the VRF. Also note that the uncertain observations are ranged indicatively according to the prediction of the model described in Section 4.1.

Douillard, Underwood, Kuntz, Vlaskine, Quadros, et al., (2011)] was applied to the ground-truth map to extract the true ground and true obstacles. As described in Section 5.2, whenever the RGS system labels data along a particular scan azimuth as ground or nonground, a range reading is returned. When combined with the localization estimation of the vehicle, this provides a 3D georeferenced position for the labeled point (see also Section 6.2).

A ground-labeled observation is counted as a false positive if a closest neighbor cannot be found in the true ground data within a minimum distance (less than 0.5 m in

our case). Similarly, a non-ground-labeled point is counted as a false positive if it is not sufficiently close to the nearest true obstacle datum.

False negatives arise when the ground/nonground is present in the image but the RGS system is not able to label it, returning the label of unknown instead. Because the system cannot provide a range measurement for the unknown-labeled observations, the rate of false negatives can be evaluated only by manual inspection in each radar image.

The results are collected in Table III. The rate of false positives in ground-labeled observations was 2.1%, likely


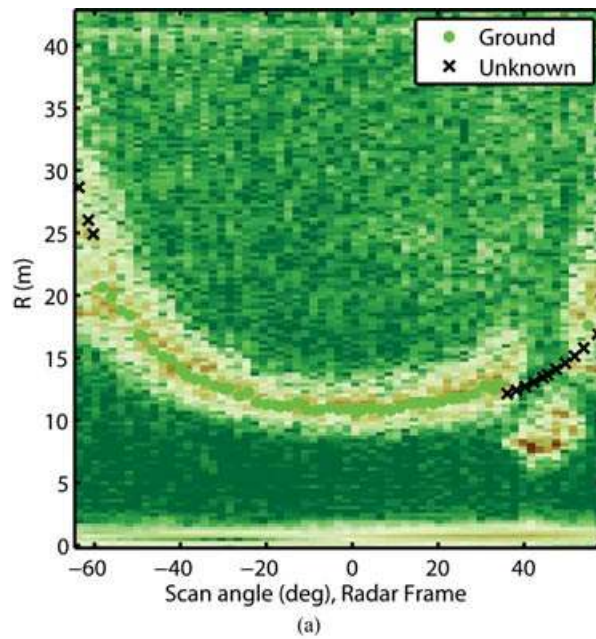


Figure 15. Results obtained from the RGS system in the presence of a large obstacle to the right: output of the RGS system (a); results overlaid on the camera image (b) and on the laser-generated ground-truth map (c).

Table III. Segmentation results obtained from the RGS system.

Class	Observations	False positives (%)	False negatives (%)	Accuracy
Ground	40,150	2.1	3.5	$E_z = 0.051$ m
Nonground	657	0.0	5.6	$E_{xy} = 0.065$ m

The false-positive rate was evaluated by comparison with the ground-truth laser data. The percentage of false negatives was obtained by manual inspection. See Section 6.2 for more details on the definition of E_z and E_{xy} .

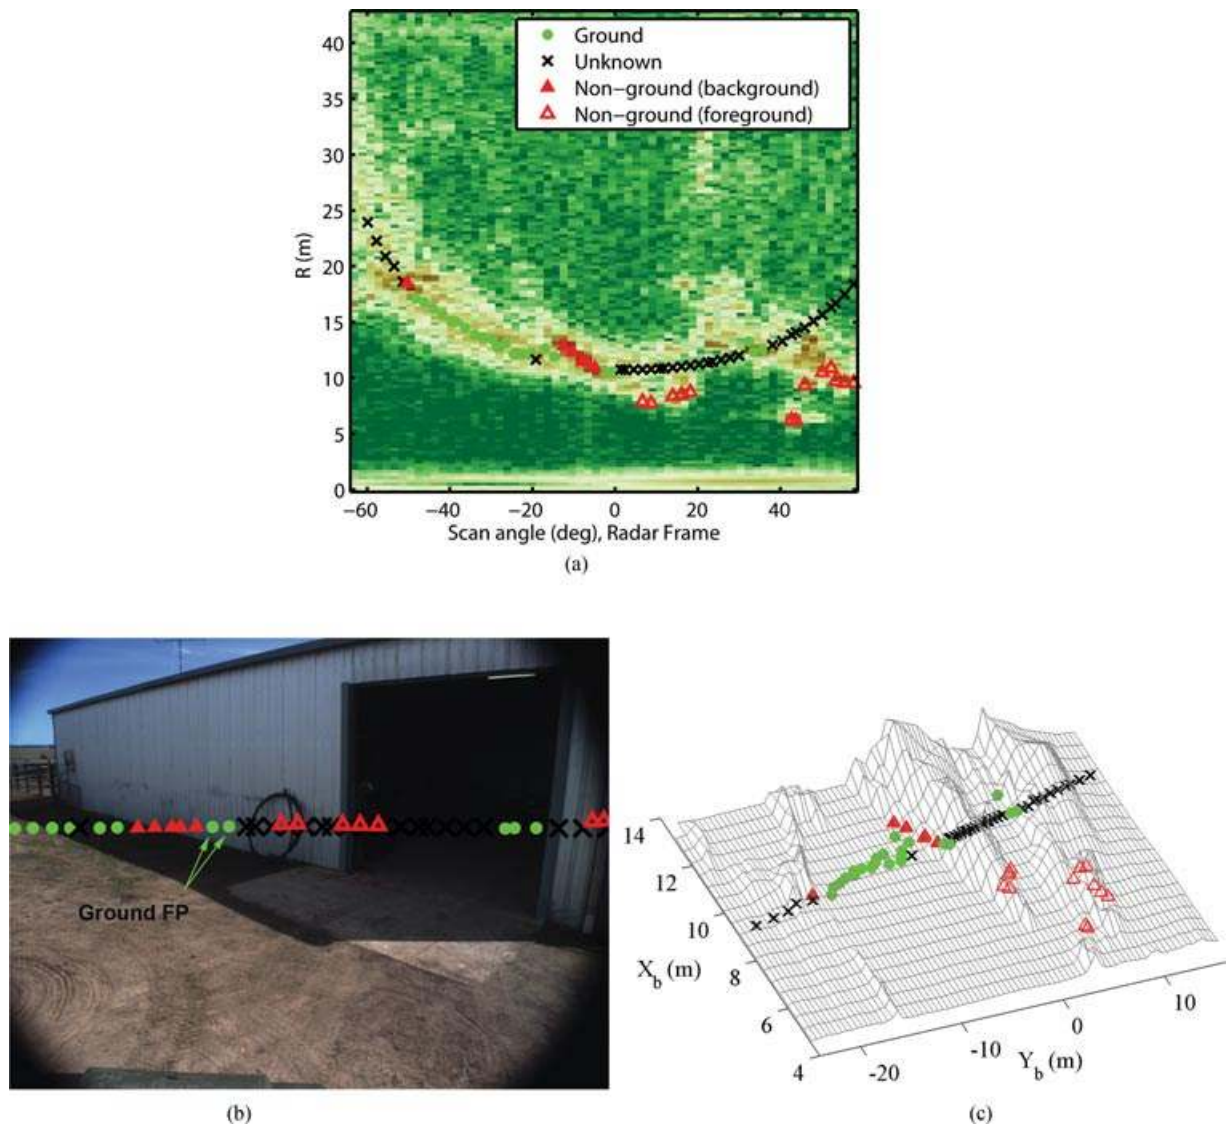


Figure 16. Results obtained from the RGS system with the robot facing a metallic shed: output of the RGS system (a); results overlaid on the camera image (b) and on the laser-generated ground-truth map (c).

due to seeming matches produced by obstacles present in the illuminated area. Typical examples of ground false positives are marked in Figure 16(b). No false positives were detected in the non-ground-labeled observations.

The false-negative rate for the ground-labeled observations was 3.5%. A typical example shown in Figure 14(b) is due to low power return of the radar observation. For non-ground-labeled observations, the false-negative rate was 5.6%.

The average rate of unknown labels in a single radar image was 19.8%, including occluded areas and false negatives due to radar misreading or low resolution (i.e., footprint overlapping part of an object and ground). It should be noted that false negatives mostly appear in the radar

image as spurious observations that do not affect the general understanding of the scene, as shown in the example of Figure 14(b).

6.2. Accuracy Analysis

The accuracy of the RGS system in ranging the ground was assessed through comparison with the true ground map. For the ground-labeled observation i , the RGS system outputs the relative slant range $\bar{R}_{0,i}$. Through geometric transformation, as described in Section 4, it is possible to estimate the corresponding 3D point in the WRF P_i and to compare it to the closest neighbor of the ground-truth map P_i^{gt} . Because the laser-generated map is available as a

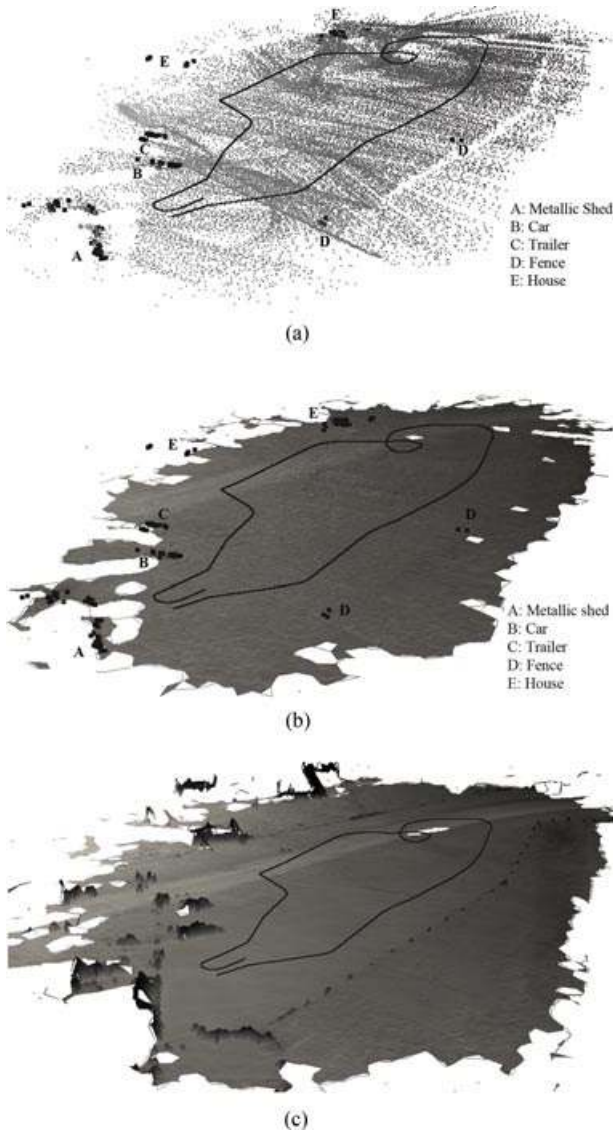


Figure 17. Segmentation results for the entire test: radar-generated map, shown as raw data obtained from the RGS system (a); same radar data after Delaunay triangulation (b); and laser-generated map after Delaunay triangulation (c). Note that the RGS system outputs only the location of the detected obstacles and not their elevation.

regularly sampled grid with square cells of 0.3 m, where the center of the grid represents the average height of the cell points, a mean square error in the elevation can be defined as

$$E_z = \frac{1}{n} \sqrt{\sum_{i=1}^n (P_{z,i} - P_{z,i}^{gt})^2}. \quad (17)$$




Figure 18. Field test for the low-visibility experiments.

In this experiment, the RGS system detected ground returns in $n = 40,150$ observations with an error of $E_z = 0.051$ m and an associated variance of $\sigma_z = 0.002$ m². If the value of R_0 is measured conventionally taking the intensity peak of the ground return, the error grows to $E_z = 0.251$ m and $\sigma_z = 0.181$ m². Similarly, the accuracy of the system in measuring the position of detected obstacles can be evaluated by comparison with the nearest datum in the true obstacle map. A mean square error can be defined this time as

$$E_{xy} = \frac{1}{n_1} \sqrt{\sum_{i=1}^{n_1} (P_{x,i} - P_{x,i}^{gt})^2 + (P_{y,i} - P_{y,i}^{gt})^2}. \quad (18)$$

The RGS system measured nonground returns in $n_1 = 657$ observations with an error of $E_{xy} = 0.065$ m and a variance of $\sigma_{xy} = 0.0015$ m², thus proving the effectiveness of the proposed approach.

For a complete overview of the system performance, the results obtained from the RGS module along the entire experiment are used to build a map of the environment, as shown in Figure 17(a). The ground-labeled observations are denoted by gray scale dots colored according to the elevation, whereas the obstacle-labeled points are shown by black points for higher contrast. The path followed by the robot is also shown by a solid black line. Figure 17(b) depicts the same data after a postprocessing step applying a Delaunay triangulation. Finally, in Figure 17(c) the laser-generated map is shown for comparison using the same color scale and Delaunay triangulation. This figure demonstrates that the RGS system is capable of providing a clear understanding of the environment, suitable for robotic applications including scene interpretation and autonomous navigation.

6.3. Persistence

To prove the robustness to low-visibility conditions, the RGS system was applied to a second data set. In this

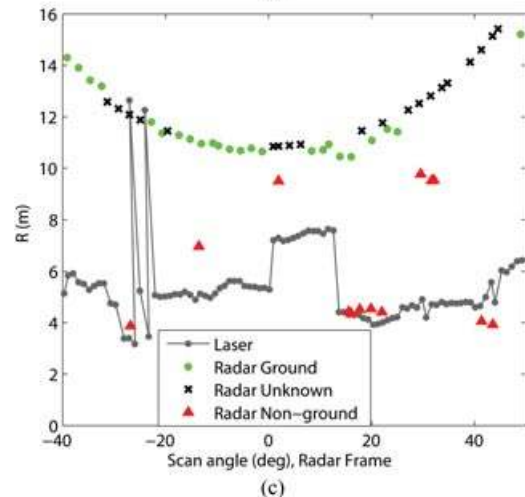
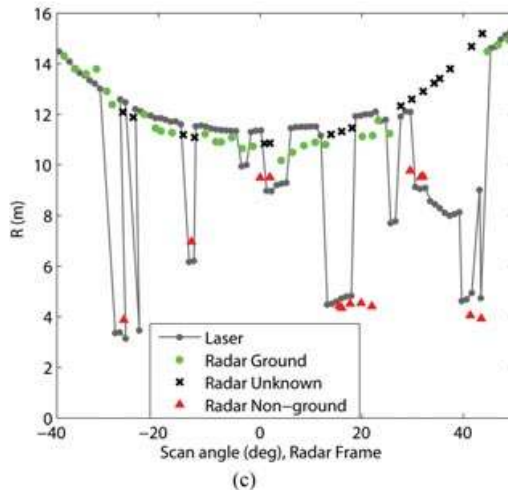



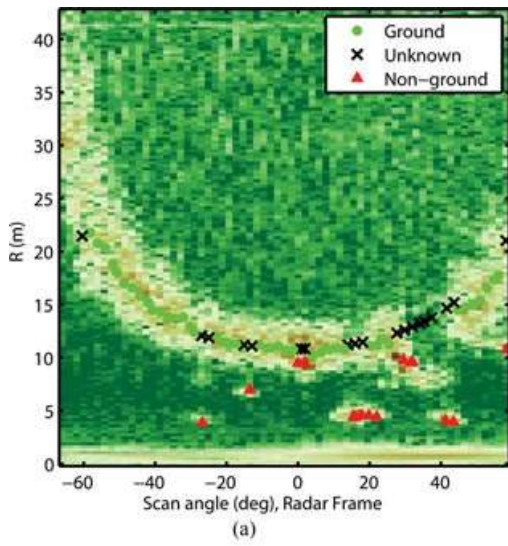


Figure 19. Results obtained from the RGS system with the robot facing a fixed area in clear visibility condition: output of the RGS system (a), results overlaid on the camera image (b), and laser readings projected onto the radar sensor reference frame (c).

Figure 20. Results obtained from the RGS system in the presence of heavy dust: output of the RGS system (a), results overlaid on the camera image (b), and laser readings projected onto the radar sensor reference frame (c).

experiment, the vehicle was stationary while facing a fixed area in changing environmental conditions. The purpose was to assess the persistence of the CORD UGV sensor suite in the presence of a heavy dust cloud. Figure 18 shows the test field including a metal frame, artificial objects of various shapes and geometry, and some equipment, as well as natural features such as a tree branch (see Peynot, et al., 2010, for a complete description of the object sizes and positions). The test started in clear visibility conditions, and afterward a dust cloud was artificially generated by blowing air to a dusty soil pile using a high-power air compressor. The dust was carried by the wind from the left to the right of the scene, and it moved between the test area and the sensors, obstructing significantly their field of view (FOV).

A typical result obtained from the RGS system in nominal visibility condition (absence of obscurants) at time $T_1 = 5$ s is shown in Figure 19(a) overlaid on the radar intensity image. Ground labels are denoted by green dots; a black cross marks uncertain terrain. Nonground labels, present in the foreground as high-intensity narrow pulses, are also shown, and they are denoted by red triangles. In Figure 19(b), the same results are projected over the image plane of the camera for visualization purposes only. Finally, a comparison with one of the 2D SICK lasers [*Laser-Horizontal* in Figure 2(b)] is provided in Figure 19(c), where the data were transformed into the radar sensor reference frame. The discrepancies between the two sensors can be explained when considering that their scan planes do not intersect perfectly on the ground and they are generally not aligned; as a result they “look” at different sections of terrain and at objects from different heights. As can be seen from these figures, the RGS system correctly labeled the ground area and most of the obstacles present in the scene. The accuracy was also consistent with the values obtained in Section 6.2: the mean square error in the elevation of the ground-labeled observation was $E_z = 0.050$ m ($\sigma_z = 0.002$ m²), whereas the error in the position of the obstacles was $E_{xy} = 0.06$ m ($\sigma_{xy} = 0.156$ m²).

The effects of the dust cloud on the robot perception are shown in Figure 20 for a typical sample scan at time $T_2 = 44$ s. The laser was largely affected by the dust, with most of the readings unusable due to reflections from airborne particles [Figure 20(c)]. Similarly, vision data were significantly corrupted due to dust occlusion, as shown in Figure 20(b). In contrast, radar was less susceptible to the presence of obscurants and the output of the RGS was almost identical to the one obtained in clear condition. The change in the range accuracy of the ground- and nonground-labeled measurement was also negligible, thus proving the robustness and persistency of the proposed approach to challenging conditions.

7. CONCLUSIONS

In this paper, a novel method for performing ground segmentation was presented using a MMW radar mounted on

an off-road vehicle. It is based on the development of a physical model of the ground echo that is compared against a given radar observation to assess the membership confidence to one of the three broad categories of ground, nonground, and unknown. In addition, the RGS system provided improved range estimation of the ground-labeled data for more accurate environment mapping when compared to the standard highest-intensity-based approach. A comprehensive experiment in the field demonstrated the overall effectiveness of the proposed approach. The RGS method was able to correctly label ground with rates of false positive and false negative of 2.1% and 3.5%, respectively, and with 0.051-m accuracy, which is a considerable improvement over the reference of 0.251 m. The system was also proved to be effective in heavy-dust conditions, showing its merits over laser and vision sensing.

This technique can be successfully applied to enhance perception for autonomous off-road vehicles in natural scenarios or more generally for ground-based MMW radar terrain sensing applications.

ACKNOWLEDGMENTS

The authors are thankful to the Australian Department of Education, Employment and Workplace Relations for supporting the project through the 2010 Endeavour Research Fellowship 1745.2010. This research was undertaken through the Centre for Intelligent Mobile Systems (CIMS) and was funded by BAE Systems as part of an ongoing partnership with the University of Sydney. The financial support of the ERA-NET ICT-AGRI through the grant Ambient Awareness for Autonomous Agricultural Vehicles (QUAD-AV) is also gratefully acknowledged.

REFERENCES

- Boehmke, S., Bares, J., Mutschler, E., & Lay, K. (1998, May). A high speed 3D radar scanner for automation. In IEEE International Conference on Robotics and Automation, Leuven, Belgium.
- Brooker, G. (2005). Introduction to sensors. New Delhi, India: SciTech Publishing.
- Brooker, G., Hennessey, R., Bishop, M., Lobsey, C., Durrant-Whyte, H., & Birch, D. (2006). High-resolution millimeter-wave radar systems for visualization of unstructured outdoor environments. *Journal of Field Robotics*, 23(10), 891–912.
- Brooker, G., Hennessey, R., Lobsey, C., Bishop, M., & Widzyk-Capehart, E. (2007). Seeing through dust and water vapor: Millimeter wave radar sensors for mining applications. *Journal of Field Robotics*, 24(7), 527–557.
- Clark, S., & Durrant-Whyte, H. F. (1997, December). The design of a high performance MMW radar system for autonomous land vehicle navigation. In International Conference Field and Service Robotics, Sydney, Australia.
- Currie, N. C., Hayes, R. D., & Trebits, R. N. (1992). Millimeter-wave radar clutter. Norwood, MA: Artech House.

- DeSouza, G., & Kak, A. (2002). Vision for mobile robot navigation: A survey. *IEEE Transactions on Pattern Analysis and Machine Intelligence*, 24(2), 237–267.
- Dissanayake, G., Newman, P., Clark, S., Durrant-Whyte, H. F., & Csorba, M. (2001). A solution to the simultaneous localization and map building (SLAM) problem. *IEEE Transactions on Robotics and Automation*, 17(3), 229–241.
- Douillard, B., Underwood, J., Kuntz, K., Vlaskine, V., Quadros, A., Morton, P., & Frenkel, A. (2011, May). On the segmentation of 3-D lidar point clouds. In *IEEE International Conference on Robotics and Automation*, Shanghai, China.
- Durrant-Whyte, H. F. (2002). An autonomous guided vehicle for cargo handling applications. *International Journal on Robotics Research*, 15(5), 407–441.
- Foessel-Bunting, A. (2000, November). Radar sensor model for three dimensional map building. In *Proceedings SPIE, Mobile Robots XV and Telemanipulator and Telepresence Technologies VII*, Boston.
- Foessel-Bunting, A., Chheda, S., & Apostolopoulos, D. (1999, August). Short-range millimeter-wave radar perception in a polar environment. In *International Conference on Field and Service Robotics*, Leuven, Belgium.
- Jiang, T. Z., Wu, H., Wu, K., & Sun, X. W. (2001). Thresholding method of CFAR for millimeter-wave collision warning radar. *Journal of Infrared and Millimeter Waves*, 24(3), 217–220.
- Jochem, T., Pomerleau, D., & Thorpe, C. (1995, August). Vision-based neural network road and intersection detection and traversal. In *IEEE/RSJ International Conference on Intelligent Robots and Systems*, Osaka, Japan.
- Kaliyaperumal, K., Lakshmanan, S., & Kluge, K. (2001). An algorithm for detecting roads and obstacles in radar images. *IEEE Transactions on Vehicular Technology*, 50(1), 170–182.
- League, R., & Lay, N. (1996). System and method for tracking objects using a detection system. U.S. Patent no. 5.587.929.
- Mullane, J., Adams, D. M., & Wijesoma, W. S. (2009). Robotic mapping using measurement likelihood filtering. *International Journal of Robotics Research*, 28(2), 172–190.
- Nash, J. E., & Sutcliffe, J. V. (2006). River flow forecasting through conceptual models part 1: A discussion of principles. *Journal of Hydrology*, 10(3), 282–290.
- Ohnishi, N., & Imiya, A. (2006). Dominant plane detection from optical flow for robot navigation. *Pattern Recognition Letters*, 27(9), 1009–1021.
- Page, E. S. (1954). Continuous inspection schemes. *Biometrika*, 41(1–2), 100–115.
- Pagnot, R., & Grandjea, P. (1995, May). Fast cross-country navigation on fair terrains. In *IEEE International Conference on Robotics and Automation*, Nagoya, Japan.
- Peynot, T., Scheduling, S., & Terho, S. (2010). The Marulan data sets: Multi-sensor perception in a natural environment with challenging conditions. *International Journal of Robotics Research*, 29(13), 1602–1607.
- Peynot, T., Underwood, J., & Scheduling, S. (2009, October). Towards reliable perception for unmanned ground vehicles in challenging conditions. In *IEEE/RSJ International Conference on Intelligent Robots and Systems*, St. Louis, MO.
- Pomerleau, D. (1989). ALVINN: An autonomous land vehicle in a neural network. In *Advances in neural information processing systems 1*. San Francisco, CA: Morgan Kaufmann.
- Poppinga, J., Birk, A., & Pathak, K. (2008). Hough based terrain classification for realtime detection of drivable ground. *Journal of Field Robotics*, 25(1–2), 67–88.
- Reina, G., Ishigami, G., Nagatani, K., & Yoshida, K. (2010). Odometry correction using visual slip-angle estimation for planetary exploration rovers. *Advanced Robotics*, 24(3), 359–385.
- Seber, G. A., & Wild, C. J. (1989). *Nonlinear regression*. New York: Wiley.
- Singh, S., Simmons, R., Smith, T., Stentz, A., Verma, V., Yahja, A., & Schwehr, K. (2000, October–November). Recent progress in local and global traversability for planetary rovers. In *IEEE International Conference on Robotics and Automation*, San Francisco, CA.
- Skolnik, M. I. (1981). *Introduction to radar systems*. New York: McGraw Hill.
- Slater, D. (1991). *Near field antenna measurements*. Boston: Artech House.
- Ulrich, I., & Nourbakhsh, I. (2000, July–August). Appearance-based obstacle detection with monocular color vision. In *AAAI National Conference on Artificial Intelligence*, Austin, TX.
- Underwood, J. P., Hill, A., Peynot, T., & Scheduling, S. J. (2010). Error modeling and calibration of exteroceptive sensors for accurate mapping applications. *Journal of Field Robotics*, 27(1), 2–20.
- Vandapel, N., Moorehead, S., Whittaker, W., Chatila, R., & Murrieta-Cid, R. (1999, March). Preliminary results on the use of stereo, color cameras and laser sensors in Antarctica. In *International Symposium on Experimental Robotics*, Sydney, Australia.
- Vosselman, G., & Dijkman, S. (2001). 3D building model reconstruction from point clouds and ground planes. *International Archives of Photogrammetry and Remote Sensing*, 34(3/W4), 37–43.
- Zhou, J., & Baoxin, L. (2006, October). Robust ground plane detection with normalized homography in monocular sequences from a robot platform. In *IEEE International Conference on Image Processing*, Atlanta, GA.