

Radar Sensor Model for Three-Dimensional Map Building

Alex Foessel-Bunting*

The Robotics Institute, Carnegie Mellon University, Pittsburgh, PA 15213 USA

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, specular effects and limited range resolution—all of which result in poor environment maps. Evidence grids are a flexible and powerful probabilistic method for fusing multiple sensor observations. Sensor models exist for interpreting the range readings of sonar, laser and stereo. However, these existing sensor models do not work with radar because it provides amplitude values for many points downrange. In addition, radar has significant echo signal-to-noise variations between observations as well as limited downrange resolution.

This paper presents the development of a radar sensor model, which can fuse amplitude-vector sensor data into an evidence grid. A study of radar phenomena and of frequency-modulated continuous-wave signal processing suggests rules for signal interpretation. The sensor model uses these interpretation rules and captures the volumetric beam geometry. The results include a three-dimensional map of an outdoor scene. This work is a step towards building high fidelity maps to be used in mobile robot navigation, obstacle avoidance and tool deployment under all visibility conditions.

Keywords: Millimeter-wave radar, constant false alarm, 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.¹ 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 reflections by specular surfaces, a limited maximum range, a long settling time and cross coupling. The beamwidth of sonar can be tens of degrees, which results in poor angular resolution. Software techniques for indoor perception improve the resolution and fidelity of sonar readings with some success. However results are 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.² Radar overcomes the shortcomings of laser, stereo and sonar. However current developments of short-range radar imaging are not satisfactory³. Generating short-range perception models from radar is beyond the state of art. High angular resolution can only be obtained with inconveniently large antenna apertures, and downrange resolution has hardware limitations. Similarly robot position estimation and navigation based on radar succeed only when these techniques rely on artificial landmarks because natural landmarks are hard to distinguish from noise and unwanted signals.^{4,5} The signal-to-noise ratio of scene objects varies significantly with pose change, material and nearby objects. Early attempts to build perception models using short-range radar compare poorly with those built from sonar, laser and stereo sensors.

* Correspondence: Email afoessel@cmu.edu, Phone (412) 681-9527, Fax (412) 681-6961.

Researchers often use evidence grids as a representation for perception models from sonar,¹¹ stereo¹² and laser¹³ sensing. The evidence-grid approach divides the space of interest, which can be two-dimensional or three-dimensional, in regular cells. Each cell stores the accumulated evidence of occupancy for the corresponding area or volume as provided by the sensor observations. The technique takes into account the uncertainty of sensor data through a probabilistic *sensor model*. Evidence-grid extensions use more than one value per cell to include confidence values or evidence of other hypotheses.^{13,15} However, common evidence grid implementations use a single value: the logarithm of the odds of occupancy.¹⁸ The evidence-grid representation can be directly used for robotic tasks such as planning, navigation and collision avoidance. Position-probability grids represent evidence of the robot position in any cell for localization.¹⁷

The sensor model is a key element in the evidence grid framework. It describes the interaction of the sensor with the environment. The sensor model maps a sensor range measurement to a probability distribution of the presence of objects. The sensor model typically accounts for the sensor uncertainty in range and angle. This model considers beam geometry, wide in the case of sonar and narrow for laser and stereo—and sensor orientation to modify the grid cells. The probability distribution also accounts for environment-sensor interaction hypotheses such as the independence of subsequent readings, the existence of only one object per observation and specular or diffuse reflections. Abundant research exists on sensor models for sonar, laser and stereo with various levels of complexity. Some stereo sensor models consider the image disparity as a confidence value to scale the evidence. Other sensor models lift restrictive assumptions and consider the existence of multiple objects¹³ or include surface orientation hypotheses.^{13,15}

Radar presents significant advantages as a perception modality in compromised visibility conditions, and evidence grids have proved useful for the integration of uncertain and noisy sensor information. However no implementations of radar sensor models for evidence grids exist that would allow construction of perception models from radar. Sonar, laser and stereo differ significantly from radar because they output a single range measurement. Frequency-modulated continuous-wave (FMCW) radar outputs a vector with values for each range bin within the sensor range. In addition most existing sensor models do not account for the amplitude of the signal. Exceptions are the aforementioned stereo sensor model that considers the disparity map. Sensor models designed to accept a single range measurement cannot easily adapt to radar signals. The development of a radar-sensor model for building 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 world representations. Experimentation with data from an outdoors scene shows the radar sensor model applicability. The next section revises the signal output for FMCW radar and discusses the inherent sensor limitations in downrange and angular resolution. The section also includes conditions that cause significant signal-to-noise variations, an explanation of the effect of large footprints and a description of the sensor used in the experiments. Section 3 studies signal processing techniques to extract the best resolution from the radar and presents ways to model the beamwidth. It also offers two approaches for signal-amplitude interpretation and the corresponding sensor models. Section 4 describes the experiments and results in downrange improvement and in outdoor scene representation. Section 5 concludes with a discussion of the results and suggestions of future work.

2. FMCW-RADAR BACKGROUND

Radar transmits electromagnetic radiation through an antenna. As a signal propagates, objects reflect, refract and absorb it. The radiation that returns to the receiving antenna converts to a vector of values representing the received-signal amplitude. Large amplitude provides evidence of a large reflecting object. Low amplitude suggests the absence of any reflecting object. Low amplitude could however correspond to a specular reflecting surface or to a highly absorbing material. A filter generally corrects for attenuation due to distance. An ideal filter results in similar signal strengths for similar objects at different distances.

Radar has the ability to measure the range of more than one object downrange and also provides information of distributed objects such as surfaces. Radar cross section represents the ability of an object to return energy to an antenna. This concept is function of three parameters: geometric cross section related to the size of the object; reflectivity, that portion of the energy that is a reflected and not absorbed and that is dependent on the material properties; and directivity, a measure of the energy returned to the antenna—mainly influenced by target geometry. False objects due to specularity exist in radar as well as in sonar and laser perception. For each type of radiation, surface types exist that are specular and that reflect transmitted energy away from surfaces. This inhibits perception of the reflecting surface and introduces inexistent objects that backscatter energy through the inverse reflected path. These inexistent objects appear because they are farther away from the front of the antenna.

Radar imaging is based on a narrow beam. The beam aims at several spots in the scene and obtains range information. This information can be used to build a three-dimensional scene map. The beam shape affects the angular resolution. In addition the beamwidth defines the footprint size, which can be very large for wide beams during observations of surfaces with low incidence angles.

This section revisits those radar properties that are relevant to the development of a sensor model for fusion of radar data into a three-dimensional evidence grid. Reviews of the radar-beam geometry and the FMCW-modulation method as well as a description of the radar unit utilized in the experiments follow.

2.1. Radar-beam geometry

For robot perception, an accurate range map of the environment can be constructed through the scanning of a *pencil beam*. This is an antenna pattern with a single, narrow lobe. Millimeter-wave radar provides a pencil beam with a relatively small antenna aperture.




Figure 1 Graphic and schematic representations of an antenna-radiation pattern. The graph on the left shows the gain of the antenna at each angle from the front. On the right, the length of the shaded areas is proportional to the gain—not to the maximum range.

The antenna-radiation pattern is a geometric property of a radar antenna. It describes the sensitivity of an antenna to a signal at a particular incidence angle. This pattern is usually the same for transmission and reception. A directional antenna concentrates most of its sensitivity in one small frontal area. Similarly, it radiates most of its power in that area. This power concentration is the radar beam. Beamwidth is the width of the antenna mainlobe at some specified level. Typically it is the level at which radiated power density is one half the maximum value on the beam axis. In this case, the beamwidth (θ) relates to the antenna aperture (D) and the wavelength (λ) according to the following equation:

$$\theta = 1.02 \frac{\lambda}{D}. \quad (1)$$

Angular resolution is the result of low beam divergence (i.e., a narrow beam). The radar beamwidth is inversely proportional to antenna aperture for a given frequency. For commercial applications, 77 GHz is a common millimeter-wave radar frequency. A 1° beam is achieved with an antenna aperture of 224 mm at that frequency. This applies to both axes in the case of a symmetrical beam defining minimum width and height for the antenna. Achievement of a narrower beam requires an impractical antenna size. Therefore the angular resolution for point objects is limited.

Most radar applications sense objects 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 beginning of the far-field region occurs at a distance R_{nf} for which

$$R_{nf} = \frac{D^2}{\lambda}. \quad (2)$$

Equation (3) results in $R_{nf} \approx 13.6\text{m}$ for the radar used in this study. Slow operations and actuator-related robot tasks require short-range sensing. Therefore scene objects often fall within the near-field region where the antenna pattern is range-dependent. Earlier work with millimeter-wave radar for robotic perception shows deformation of the beam in near-field conditions.³ Additionally, bistatic radars such as the one used in this study introduce parallax distortions close to the pair of antennas. Near field and parallax effects do not prevent sensing close to the radar but suggest complex beam geometries and increased difficulty in data interpretation.

2.2. FMCW-radar output

The FMCW radar measures the range to objects by comparing the frequency of transmitted and received signals. The transceiver generates a signal of linearly increasing frequency for the frequency-sweep period. The signal propagates from the antenna to a static target and back. The value of the received-signal frequency compared to the transmitted-signal frequency is proportional to the propagation range.




Figure 2 Frequency-modulated continuous-wave radar emits a varying-frequency signal. The signal propagates and backscatters on objects. Some of those echoes reach the antenna. The comparison of transmitted and received signals provides an estimate of the time of flight. This delay is proportional to the range between the sensor and the object. The transceiver output contains frequency components that relate to each object.

The transmitter generates a sine signal at the operation frequency of the radar and increases the signal frequency in time. This increase of frequency is a *frequency sweep* (a.k.a. frequency ramp). The operation frequency, the frequency-sweep period and the frequency sweep range (bandwidth) characterize the frequency sweep. The transmitted signal is expressed by

$$T_x(t) = A \cdot \sin\left(2\pi F_0 \cdot t + \pi \frac{B_s}{T_s} t^2\right) \quad 0 \leq t \leq T_s. \quad (3)$$

The frequency of the transmitted signal changes from the operational frequency at the beginning of the sweep to the operational frequency plus the sweep bandwidth at the end of the sweep. This signal propagates until it illuminates one or several objects. The objects scatter the signal, and a fraction of the energy propagates back to the sensor. The time of travel is proportional to the distance and the propagation speed ($c = 3 \cdot 10^8$ m/s). The round-trip propagation delay for an object at a distance r_t is $t' = 2r_t / c$, which is significantly shorter than the frequency-sweep period (e.g. the delay for an object at 64 m is $t' = 0.43 \mu\text{s}$ and the frequency-sweep period is $T_s = 256 \mu\text{s}$). The received signal is the summation of all the echoes, all with different delays proportional to the range of the reflecting objects. The propagation and reflection result in an attenuation of the amplitude for each echo. The reflection adds a random delay ϕ . The expression for the returned signal is

$$R_x(t) = \sum_{\{i \in \text{objects}\}} A'_i \cdot \sin\left(2\pi F_0 \cdot (t - t'_i) + \pi \frac{B_s}{T_s} (t - t'_i)^2 + \phi\right) \quad 0 \leq t \leq T_s \quad (4)$$

where A'_i is the attenuated amplitude, $t - t'_i$ represents the time shift of the signal and ϕ_i is an arbitrary reflection phase angle. The transceiver compares the two signals by multiplying them in a mixer. The mixer output (after some the application of trigonometric identities) has two series of summands located in opposite sides of the spectrum, as in Equation (5). The instantaneous frequency is the derivative of the cosine argument. The argument of the cosines in the first series is the difference between the two signals' frequencies. There are four summands: three summands do not depend on the time, they thus define the phase of the signal. The summand that is a function of time defines the frequency and has three parameters: T_s , B_s and t' . The values of T_s and B_s are specific to the transceiver configuration.

$$M_o(t) = \sum_{\{i \in \text{objects}\}} \frac{AA'_i}{2} \cdot \cos\left(\frac{2\pi B_s t'_i}{T_s} t + \left(2\pi F_0 t'_i - \pi \frac{B_s}{T_s} t'^2_i - \phi\right)\right) - \sum_{\{i \in \text{objects}\}} \frac{AA'_i}{2} \cdot \cos\left(2\pi F_0 t + \pi \frac{B_s}{T_s} t^2 + 2\pi F_0 (t - t'_i) + \pi \frac{B_s}{T_s} (t - t'_i)^2 + \phi\right) \quad (5)$$

Equation (6)(a) shows the frequency difference related to the time delay for one object. The derivative of the cosine argument in the second series of Equation (5) indicates an instantaneous frequency that approximates to $2F_0$ and appears in Equation (6)(b). Components with such high frequencies are above the response of the circuitry and do not appear in the transceiver output.

$$(a) \frac{B_s t'}{T_s} \qquad (b) \frac{2B_s}{T_s} t + 2F_0 - \frac{B_s t'}{T_s} \approx 2F_0 \qquad (6)$$

In summary the output of the transceiver is a signal with as many sinusoidal components as there are objects in the beam. The frequency of the components is proportional to the range of the target. Once the signal is digitized, a Fast Fourier Transform (FFT) is commonly used to detect the components and to estimate their frequency. The FFT outputs a vector of amplitudes for each range bin along the range of interest. High amplitudes typically represent objects with large RCS (Radar Cross Section). Low amplitudes indicate objects with small RCS or empty space.

Radar has the ability to provide information on objects that would occlude each other under observation with optical sensors. Downrange resolution is a radar property that measures the ability to resolve objects aligned with the main-beam axis. In the case of FMCW radar, object discrimination is possible only if there is a distinct pulse for each object. It is possible to resolve two pulses if they do not overlap. Pulse overlapping occurs when the objects are too close or when the pulses are too wide. The pulse width depends, in the case of FMCW radar, on the bandwidth of the frequency sweep. Equation (7) shows that the pulse width depends only on the frequency sweep bandwidth

$$\Delta R = \frac{c}{2B} . \qquad (7)$$

There are no other parameters we can change to increase the resolution; and this is the best result a radar can achieve. Signal processing should aim at achieving this resolution to provide the best information to the sensor model.

2.3. FMCW-radar unit description

This investigation uses mechanically scanned millimeter-wave radar designed for obstacle avoidance and road navigation in heavy-dust conditions. The sensor is a bistatic 77-GHz FMCW millimeter-wave radar. The wavelength is ~ 3.9 mm. The beamwidth is 2° in elevation and 1° in azimuth. The bistatic-antenna pair scans horizontally across the angular range of 64° . Simultaneously in intervals of 1° in azimuth, switches sequentially select four stacked-beam positions covering a total elevation range of 8° . One scanning cycle—or radar frame—covers a 64° by 8° area of the scene. A stand with a vertical-tilt axis allows the sensor to be aimed at different elevation or depression angles. The sensor range of operation is between 1 and 64 m. The transceiver sweeps 300 MHz in 256 μ s; look-up compensation tables provide linearity for wide temperature-range operations. The acquisition at 1 MHz provides 256 samples per ramp.

3. APPROACH AND METHODOLOGY

Robotic applications—such as navigation, obstacle avoidance and tool deployment—require knowledge about the presence of all solid objects in the workspace. Evidence grids are a flexible and powerful probabilistic method for fusing multiple sensor observations. Sensor models exist for interpreting the range readings of sonar, laser and stereo. However, these existing sensor models do not work with radar. Radar has the unique property of detecting more than one object with a single observation. The use of all the information contained in each radar observation speeds up the construction of a quality model. The distinction between two nearby objects results in higher scene-model quality. With this focus in consideration, the study continues with the presentation of techniques for the extraction of the best downrange resolution. In addition, the interpretation of the range-bin amplitude and the classification of pulses and noise result in expressions for evidence of occupancy or emptiness. Finally each radar observation affects several cells because of the uncertainty in angle. A three-dimensional extension of the downrange interpretation accounts for the beamwidth.

3.1. Evidence-grid framework

The evidence-grid framework divides the volume of interest in M regular cells $M = \{m_1 \dots m_N\}$. Usually, the value in each cell is a binary random variable that represents the probability of occupancy for the corresponding volume in the scene. A sensor provides L observations $R = \{r_1 \dots r_L\}$ that are uncertain, noisy and incomplete. The Bayes' rule can express the random nature of the sensor observations and provide a method to update the grid with each observation. The conditional probability $P(M|R)$ encodes the observation information and is of interest for mobile robot applications. Bayes' rule provides the following expression $P(M|R) = P(R|M) \cdot P(M) / P(R)$, where $P(M)$ and $P(R)$ are the prior probabilities for the map and the observations. The quantity $P(R|M)$ is the sensor model. It represents the probability of the observation R given that the scene is M . Odds are an intuitive representation of the probabilities with values that range from 0 (absolutely impossible) to $+\infty$

(complete certainty). This representation is also convenient because it factors out the priors. The definition of odds and odds-likelihood follows:

$$O(M) \doteq \frac{P(M)}{P(\bar{M})} \longrightarrow O(M|R) = \frac{P(R|M)}{P(R|\bar{M})} O(M) \doteq \lambda(R|M) \cdot O(M). \quad (8)$$

The logarithmic representation of Equation (9) provides a simple update rule. Log odds range from $-\infty$ to $+\infty$. The lower extreme represents an empty cell; the upper extreme an occupied cell. This study adopts the following update rule to fuse radar observations and to update the grid representation:

$$\log O(M|R) = \log \lambda(R|M) + \log O(M). \quad (9)$$

3.2. Signal processing and downrange resolution

The range of interest and the radar FMCW parameters determine the maximum frequency Δf_{max} in the transceiver output as shown in Equation (10). The Nyquist criterion indicates that the acquisition of such a signal requires a sampling frequency of $f_{samp} > 2\Delta f_{max}$. The acquisition period extends for the length of the frequency sweep T_s .

$$\Delta f_{max} = \frac{2B_s}{cT_s} r_{max}. \quad (10)$$

The Fourier transform of a time-limited signal assumes that the signal repeats itself endlessly on both sides. Any difference in the amplitude of the signal ends is interpreted as a high frequency content, and any peaks in the data exhibit spread or leakage. An alternative explanation of the phenomenon considers the transceiver signal as the multiplication of an infinite sine, or combination of sinusoids, and a step function. The transform of the product of two functions is the convolution of the respective transforms. The Fourier transform of the step is the function $\sin(x)/x$, which has significant sidelobes. The Fourier Transform of the sine is an impulse at the sine frequency. The convolution displaces the $\sin(x)/x$ function to the sine frequency, and the resulting spectrum shows the high sidelobes or leakage. This phenomenon deserves attention because the time sidelobes of a strong target can hide nearby weaker targets.




Figure 3 Effect of tapered windows in the frequency spectrum of a signal. The attenuation at the extremes of the signal forces continuity and reduces the spectrum leakage or time sidelobes. However the technique increases the mainlobe width.

A *tapered window* prevents leakage by smoothly attenuating the ends to zero. An undesirable consequence of the application of tapered windows is the increase of the pulse width. Figure 3 shows the effect of the application of a tapered window on the spectra of a signal. The upper-right graph shows the original signal and the discontinuity due to the replication. A window attenuates the extremes of signal vector as shown in upper-left graphs. The two lower graphs show a comparison of the resulting spectra. The Hanning window is a good compromise between leakage reduction and acceptable pulse width.¹² It provides a time-sidelobe attenuation of -31dB and a normalized pulse width of 3.1 over the number of samples.

The FFT is a discrete sampling of the continuous Fourier transform. The frequency interval of the discrete spectrum sampling is the inverse of the frequency-sweep ramp duration, in the case of FMCW radar. Unfortunately, that frequency sampling is similar to the range resolution dictated by the bandwidth in Equation (7) and can prevent the detection of two nearby pulses. Zero padding is the concatenation of the signal vector and a vector with zeros, as it is shown in the following expression:

$$\vec{x} = [x_1 \ \dots \ x_n] \quad n < 2^N \longrightarrow \vec{\tilde{x}} = [x_1 \ \dots \ x_n \ 0_{n+1} \ \dots \ 0_{2^N}]. \quad (11)$$

This technique is usually used to increase the number of signal samples to a power of two to take advantage of the higher speed of the FFT. An additional benefit results from increasing the number of samples: The frequency interval between discrete samples of the spectrum decreases accordingly.




Figure 4 Limited range resolution due to insufficient samples. Inset (a) shows samples every 0.5 m and the impossibility of detecting two pulses nearby. Zero padding the signal vector allows for the detection of the two pulses in most cases (b) by the comparison of consecutive slopes. There are still cases (c) in which detection is compromised because of the lack of slope change between points. Detection is certain with more samples as in (d).

The cost of using zero padding is an increase of the vector length and of the Fourier-Transform processing time. The range resolution is limited by the frequency-sweep bandwidth. Therefore the pursuit of a better range resolution is useless. Two close objects at the best possible resolution distance present two peaks and a valley. The identification of this geometry requires enough samples to detect the double change in the slope of the curve as shown in Figure 4 (b). There is a special case, however, when the samples show a flat pulse top that can correspond to a wider pulse (Figure 4 (c)). The position of the spectrum samples with respect to the peaks and valley is not known. The best possible resolution is guaranteed with more than three samples in a segment defined by Equation (7) and shown in Figure 4 (d).

3.3. Range-amplitude vector interpretation

Radar can be useful as a robotic sensor because it detects and locates objects. Detection and localization with radar are uncertain and are also prone to measurement errors. Redundant observations help to counter radar uncertainty and errors. Radar also shows incompleteness due to occlusions and limited field of view. Complementary observations provide information about different sectors of the volume of interest. Evidence grids are a convenient representation to fuse redundant observations and accumulate complementary observations. The evidence-grid approach divides the space of interest in regular cells. Each cell stores the accumulated evidence of occupancy for the corresponding volume as provided by the sensor observations.

Radar is a unique sensor because each observation returns a vector of information. The vector contains amplitude values for each range bin. There are 512 range bins after the application of zero padding, as described in the previous subsection. The amplitude of one range bin represents the energy returned to the sensor from one or more objects located within the radar beam at such a range and illuminated by the beam. The energy returned is a function of the objects RCS, the objects range and occlusions. Occlusions vary the amount of energy that illuminates an object and the backscatter that reaches the antenna. Radar can suffer total occlusion caused by large objects as well from unnoticeable perturbations caused by point-like objects. The pulse width provides evidence of the object type and of the possible occlusions. Additional factors of the range-bin amplitude are noise and time sidelobes. Time sidelobes are a consequence of the finite bandwidth. Time sidelobes of a large pulse cause unwanted variations of the amplitude in adjacent range bins. Electronic noise, ambience radar radiation and frequency-sweep nonlinearity contribute with noise to the range-bin amplitude. Sidelobes, specularly and other radar sources can introduce false targets. However this research assumes that the radar antenna does not have strong sidelobes and that there are no specular surfaces or other nearby radar sources. Under these assumptions, pulses above the noise level imply the presence of objects at the corresponding range with certainty.

Empirical rules can help the range-bin interpretation. The following empirical rules describe radar in typical mobile-robots environments. The rules presented apply to the range bins that are more distant than a pulse or between pulses. Figure 5 illustrates the phenomenon. 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 more likely represents a surface with high incident angle and not a long object aligned with the beam. Long objects aligned with the beam likely appear as a narrow pulse corresponding to the end closer to the sensor. A surface captures most of the radiated energy and occludes more distant targets. Groups of nearby objects closer than the radar resolution also appear as a long pulse and cause occlusion of more distant targets. Low amplitude and noise after a wide pulse do not imply emptiness. Low values after a point target may imply emptiness. This rule is difficult to apply because a wall perpendicular to the sensor will appear as a point target. 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.

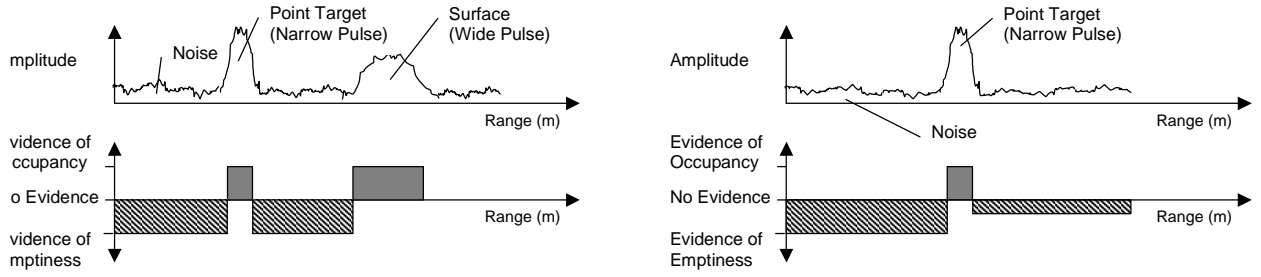


Figure 5 Empirical rules for radar evidence. High amplitudes indicate the existence of objects with high probability. Wide pulses generally correspond to surfaces at high incidence angles. Narrow pulses correspond to point-like objects or to low incidence surfaces. Low amplitudes (noise) between the sensor and the first pulse or between two pulses indicate a high evidence of emptiness, as is shown in the right most inset. Low amplitudes after a wide pulse do not provide evidence because surfaces occlude radar.

The interpretation of range-bin amplitudes 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 indicates 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 misclassifies the signal with probability $1-P_d$ because the pulse with noise added can fall below the threshold.




Figure 6 Constant false-alarm rate method 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 ratio and occupancy evidence. Noise level however does not weigh the evidence of emptiness. 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 (-14° 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. From Equation (9), E_h and E_l are negative values. The sensor-model down-range interpretation rules follow

$$\log \lambda(r_i | M) = \begin{cases} k_{occ} \log \left(\frac{r_i}{n_i} \right) & \text{if } \left(\frac{r_i}{n_i} > CFAR \right) \\ E_h & \text{if } \left(\frac{r_i}{n_i} \leq CFAR \right) \text{ and } (\text{before wide pulse}) \\ E_l & \text{if } \left(\frac{r_i}{n_i} \leq CFAR \right) \text{ and } (\text{after wide pulse}) \end{cases} \quad (12)$$

This paper refers to this function as the "down-range interpreter." The representation of the interpreter output and the beam geometry is referred to as "sensor model."

3.4. Sensor-model geometry and angular uncertainty

Radar wide beams and the resulting large footprint cause angular uncertainty. A radar pulse at a given range can result from a point-like object anywhere in the footprint at that range, from a large object covering the footprint partially or completely or from a combination of the above. The sensor model geometric representation is closely related to the beamwidth because of this uncertainty. Interpretation of the signal classification and the signal-to-noise ratio has the following implications cross range. The graph on the right of Figure 7 shows the two-way signal strength as a function of the angle measured from the beam center. The object remains at a constant range and has a constant RCS. An intuitive interpretation considers the normalized curve as the probability distribution of the object cross-range position given a pulse at a given range, or in a discrete angular space $p(\alpha = k | M_i)$. A normal distribution fits the shape of the beam. The constant RCS 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 7 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. Scheme (c) shows the combined model that represents the near field condition.

Beamwidth is typically measured at half the power level in a one-way transmission (-3 dB). However, the experimental results in Figure 7 indicate that a target placed outside of this region still has a strong echo. The sensor model could consider the update of grid cells with a wider beam to account for the tails of the distribution but that would result in a larger number of computations. This study adopts the commonly accepted half-power definition of beamwidth and disregards any evidence changes outside the volume defined by the beam.

The near-field distortion of the conical beam is another consideration for a model with more fidelity. The cylindrical-conical model proposed by the author for the radar used in this study provides a realistic beam shape for ranges close to the antenna.² Figure 7 (c) shows the combined model.

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.

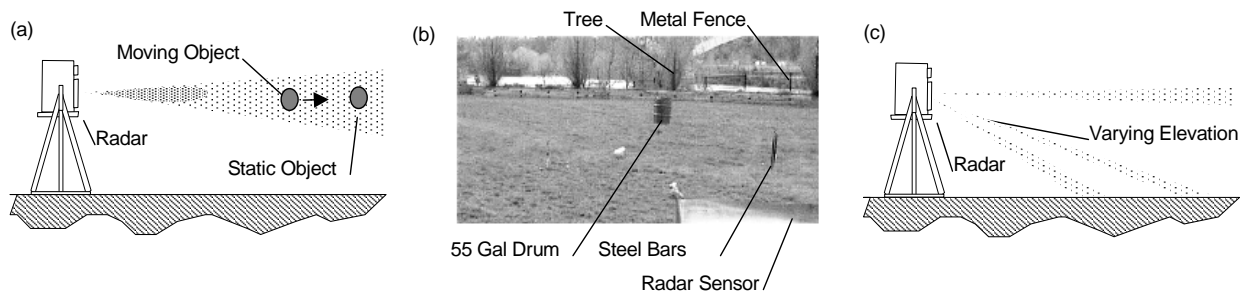


Figure 8 Experimental setup to evaluate range resolution (a). The two point-like objects are in line with the radar antenna. The experiment gathers data for various ranges between the objects. An outdoors scene to demonstrate radar observation fusion with the use of the radar sensor model appears in (b). The radar acquired observations across the sensor azimuth range, at increasing elevations from -14.2° to -0.2° every one-degree.

4.1. Downrange resolution improvement

The experimental setup appears in Figure 8 (a). Figure 9 (a) shows the spectrum of the signal corresponding to two point objects in line 7.0 and 8.0 m away from the sensor. However the spectrum indicates the existence of only one target. The signal has 256 values and no padding is required to operate the FFT. For a ramp duration of $T_s=256 \mu\text{s}$, the minimum detectable frequency is $1/T_s=3.9 \text{ kHz}$. This is also the frequency interval in the discrete spectrum. The frequency interval in the equation $r_t = (cT_s / 2B_s) \cdot \Delta f$ yields a range value every 0.5 m. Equation (7) indicates that the downrange resolution depends solely on the bandwidth; 300 MHz limit the range resolution to objects distant more than 0.5 m. In the specific case of a vector with 256 samples, the concatenation of 256 zeros simulates a total sampling duration of $T_s=512 \mu\text{s}$. The minimum detectable frequency is $1/T_s=1.95 \text{ kHz}$. Now it is possible to detect two peaks 0.5 m away from each other. Figure 9 (b) shows that the spectrum and the two objects appear distinctly. The concatenation of 768 zeros to a total length of 1,024 results in the more continuous spectrum shown in Figure 9 (c). No additional features are evident in this case. Figure 9 also shows a decreasing signal-to-noise ratio with zero padding.




Figure 9 Range signal for objects 1.0 m apart. The graphs show the spectrum with no zero padding (a), with 256 zeros (b) and with 768 zeros (c).

The effect of zero padding is denser discrete sampling of the continuous spectrum. Zero padding does not, however, increase the physical limits. Further experiments demonstrate the use of zero padding on objects 0.5 m apart. The concatenation of 256 or more zeros fails to resolve the objects. The spectrum is smoother, but no distinct features appear.

4.2. Signal interpretation

Figure 10 shows the down-range sensor model evaluate the range amplitude vector for two observations. Figure 10 (a) shows a beam that aims with an elevation of -9° and illuminates the ground with a wide footprint. The interpreter output appears in inset (c). The interpreter detects the noise level and outputs a value of -1 providing strong evidence that the cells are empty. The function detects the wide pulse 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 function also detects that this is a wide pulse and assigns a weaker level of evidence for the cells farther away. The function detects two more peaks. Arguably the values of the evidence for those pulses should be weaker but there is the possibility that the pulse width estimation classifies the pulse incorrectly.

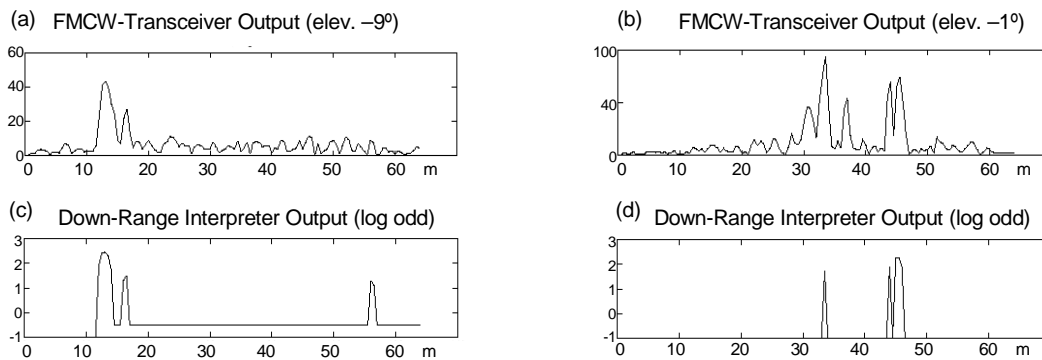


Figure 10 The downrange interpretation function applied to real data. Graph (a) shows the original data. The corresponding result appears in graph (b) where the units are in log odds to update the grid. Graphs (c) and (d) show a second example for observations at a larger depression.

The inset (b) shows a signal from an observation that is almost parallel to the ground. The pulses correspond to a fence and a tree and there is no ground return. The interpreter outputs strong evidence of emptiness in all those cells that do not belong to a pulse, as shown in inset (d). The outputs for these two observations correspond to the expected interpreter results. Not all observations followed the expected output. This study does not provide a estimation of the interpreter goodness.

4.3. Outdoor scene representation

The scene shown in Figure 8 (c) contains a variety of objects, some on the fairly flat surface covered with grass; trees, a metal fence, a 55-gallon drum, steel bars and rocks are some examples. The data set consists of observations taken over 15 elevations between -0.2° to -14.2° 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. The following figure shows the integration of a large number of the observations. The grid visualization consists in displaying a small sphere for each cell that contains odds of occupancy. The sphere radius is proportional to the grid value. The visualization shows that the map has a correct geometry and that with large RCS accumulate the most evidence. Holes in the ground plane suggest that the thresholds are eliminating the weak signals that result from grazing angles.

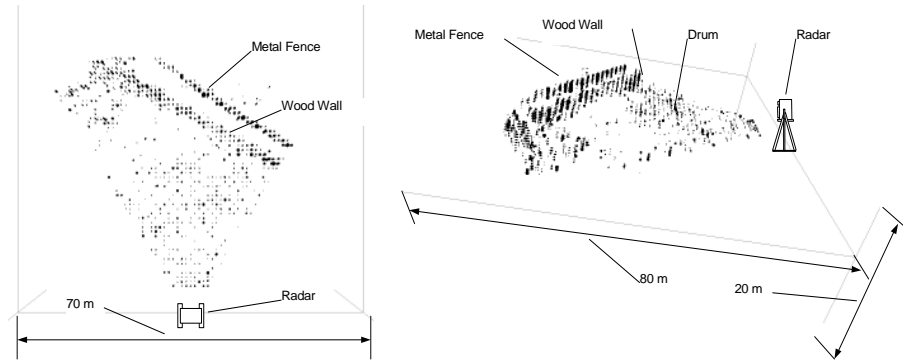


Figure 11 Visualization of the grid-based representation of an outdoor scene. The display technique thresholds the grid cells to positive values and represents the accumulated value with a sphere. The radius of the sphere is proportional to the log odds.

5. CONCLUSION

This work develops a radar-sensor model for evidence grids and demonstrates the possibilities of building three-dimensional maps of outdoors using radar sensing. The appreciation of the fidelity of the map is qualitative at this stage.

The paper investigates and presents FMCW-radar with the intent of understanding the phenomenology, which is relevant for the progress in the development of a radar-sensor model. The study of the radar output signal produced techniques to extract the best downrange resolution. Specifically zero padding and tapered windows show a downrange resolution that is near the theoretical one. This signal-processing technique improved the effective range resolution of 300-MHz frequency-sweep radar from 1.5 m to almost 0.5 m. The experiments show these results at target ranges of 8 m away from the radar antenna. A concatenation that is three times the original signal-vector length provides optimal visualization.

This research makes the assumptions that the radar antenna does not have strong sidelobes and that there are no specular surfaces or nearby radar sources. Under these assumptions, pulses above the noise level imply the presence of objects at the corresponding range with certainty. In addition, these assumptions suggest qualitative rules for a downrange signal interpretation that accounts for occlusions—that cause strong interdependencies along the radar signal.

The paper presents the Constant false-alarm rate technique for the estimation of the noise level and the setting of a detection threshold. This technique is utilized in long-range radar but also has applicability to the signal classification as pulse or noise. The experimental results indicate that the combination of qualitative rules and CFAR threshold correctly classify the signal and provide a means to transform the range-amplitude vector in evidence of occupancy or emptiness. In this work, the beam geometry captures the near-field distortions through the adoption of a cylindrical-conical model. A radar sensor model is complete when the downrange interpretation takes the beam geometry.

This work also indicates several issues worth of further investigation: The absence of sidelobes and specular surfaces is not realistic. A more rigorous method to set the thresholds is necessary. Estimations of the noise distributions can yield better pulse classification. Finally, a reference scene is essential to evaluate the goodness of new methods or new parameters.

ACKNOWLEDGMENTS

The author thanks Hans Moravec for useful discussions and for making his evidence-grids code available. Discussions with William Whittaker and John Bares contributed clarity and focus to the approach. The Telerobotics Program of NASA's Office of Space Science, Advanced Technology and Mission Studies Division funds the Autonomous Robotic Technologies Program at the Field Robotics Center, which partially sponsored this work. The Motion-Free Scanning Radar Alliance at the National Robotics Engineering Consortium also sponsored the study and provided the experimental unit. Both centers are part of The Robotics Institute, Carnegie Mellon University.

REFERENCES

1. 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.
2. 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, August 1999.
3. 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.
4. H. Durrant-Whyte, E. Bell and P. Avery, "The Design of a Radar-Based Navigation System for Large Outdoor Vehicles," *Proc. IEEE Int. Conf. on Robotics and Automation* (Nagoya, Japan 1995) pp. 764/769.
5. S. Clark and H. Durrant-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.
6. J. Bjornholt and T. Wilson, "FMCW Radar for Imaging Applications," *Proc. SPIE Radar Processing, Technology, and Applications III*, vol. 3462, p. 89-98, 1998.
7. M. Monod, P. Faure, J. Dusi and P. Lacchia, "Perception of the environment with a low cost radar sensor," *L. Radar 97 Conf. Publ. No.449*, 1997
8. X. Xu and Z. Fan, "High Resolution Imaging of Large Vehicle by MMW Scanning Sensors," *Proc. IEEE Int. Radar Conf.*, 1995, Arlington, VA.
9. M. Vossiek and P. Heide, "Signal Processing Methods for Millimeter Wave FMCW Radar With High Distance and Doppler Resolution," *European Microwave Conference 1997*; vol. 27.
10. J. Bound, "The Infrared Airborne Radar Sensor Suite," *RLE Technical Report No. 610*, MIT, 1996.
11. H. Moravec and A. Elfes, "High resolution maps from wide angle sonar," *Proc. IEEE Int. Conf. on Robotics and Automation*, Washington, DC, (1985) 116-121.
12. 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.
13. 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.
14. K. Konolige, "Improved Occupancy Grids for Map Building," 1997, *Autonomous Robots*, vol. 4, 351-367.
15. R. Takada, "Modeling Surface Orientation in 2D Evidence Grids," *Technical Report 93-026*, Mechanical Technology R&E Center, Nippon Steel Corporation, Japan, 1993.
16. S. Thrun and A. Bucken, "Integrating grid-based and topological maps for mobile robot navigation," *Proc. Thirteenth National Conference on Artificial Intelligence*, pages 944-950, 1996.
17. Ltd. Peters, W. Burgard, D. Fox, D. Hennig and T. Schmidt, "Estimating the absolute position of a mobile robot using position probability grids," *Proc. of the Fourteenth National Conf. on Artificial Intelligence*, 896-901. 1996.
18. M. Martin and H. Moravec, "Robot Evidence Grids," *Technical Report CMU-RI-TR-96-06*, Carnegie Mellon University RI, 1996.
19. E. Ifeachor and B. Jarvis, *Digital Signal Processing, A Practical Approach*, Addison-Wesley, 1993.