October 2006

Classifying surface roughness with CTFM ultrasonic sensing

Phillip J. McKerrow
University of Wollongong, phillip@uow.edu.au

B. J. Kristiansen
University of Wollongong

Classifying surface roughness with CTFM ultrasonic sensing

Abstract

Roughness is a characteristic of a surface that is a function of its geometry. Ultrasonic sensing in air provides range, area, and angle information because the surface geometry determines the characteristics of the echo. The authors introduce the “spatial-angle-filter model” to explain the impact of surface roughness on the echo. On the basis of this model, they design a set of features for use in classifying surfaces. The quality of the features and the classification is measured with the Mahalanobis distance. The resultant system is able to achieve 99.73% classification of a set of 12 surfaces using five features.

Disciplines

Physical Sciences and Mathematics

Publication Details

This article was originally published as McKerrow, PJ and Kristiansen, BE, This Classifying surface roughness with CTFM ultrasonic sensing, IEEE Sensors Journal, 6(5), October 2006, 1267-1279. Copyright IEEE. Original journal available here.

Classifying Surface Roughness With CTFM Ultrasonic Sensing

Phillip J. McKerrow, *Member, IEEE*, and Bjorn E. Kristiansen

***Abstract*—Roughness is a characteristic of a surface that is a function of its geometry. Ultrasonic sensing in air provides range, area, and angle information because the surface geometry determines the characteristics of the echo. The authors introduce the "spatial-angle-filter model" to explain the impact of surface roughness on the echo. On the basis of this model, they design a set of features for use in classifying surfaces. The quality of the features and the classification is measured with the Mahalanobis distance. The resultant system is able to achieve 99.73% classification of a set of 12 surfaces using five features.**

***Index Terms*—Continuous transmission frequency modulated (CTFM) ultrasonic sensor, Mahalanobis classifier, spatial-angle filter, surface roughness, ultrasonic sensing, wheeled mobile robot.**

I. Introduction

WHEELED mobile robots often travel along paths. Some are paths built for human traversal, some for machines to follow, and others are simply the shortest way through an open space. Whether these paths are indoors or outdoors, one of their characteristics is the roughness of their surface. Surface roughness can be used to distinguish between a path and its surrounds, for example a brick path bordered by a garden bed.

If we can measure the roughness of a surface and distinguish between surfaces using measurements of roughness then we may be able to navigate a mobile robot to follow a path by sensing surface roughness. The robot could use roughness as one of a set of landmarks both to determine where it is and to determine the direction of the path [1], [2].

Roughness is a geometric characteristic of a surface caused by small variations in surface depth. A surface is uneven from projections, irregularities, pits, and breaks. In contrast, texture is the quality of a fabric conveyed by touch due to the interweaving of threads. Visual texture is a pattern of lightness and darkness in an image, which may be the result of physical texture or roughness.

As ultrasonic sensing in air is a range-area sensor it should be suitable for measuring surface roughness. Studies of bat echolocation [3] have concluded that bats can discriminate between surfaces with different roughness. Blind people using continuous transmission frequency modulated (CTFM) mobility aids report that they can distinguish between surfaces with different roughness. Our previous research [4] demonstrated that CTFM ultrasonic sensors in air can be used to discriminate between leafy plants.

Other researchers have reported good results when measuring roughness with ultrasonic sensors. Bozma and Kuc [5] showed that measurements of surface roughness can be used to map an environment by classifying surfaces as smooth, moderately rough, and rough. They identified two measurable echo features that varied with surface roughness: echo energy and echo duration.

Politis and Probert Smith [6], [7] used CTFM sensors angled to the surface to measure roughness. Using the distribution of energy between the specular and diffuse components of the echo [two-dimensional (2-D) feature vector], they were able to distinguish between six indoor surfaces typical of pathways.

Probert Smith and Zagrafos [8] extended this paper to achieve 99.2% classification of five surfaces from individual echoes recorded from a moving vehicle with a sensor angled at 40° to the surface. When they averaged the features over five echoes the classification improved to 99.5% for a sensor at 40° and 100% for a sensor at 25°.

Classification of the seafloor is of considerable interest to researchers in underwater sonar [9], [10] and the development of theoretical models of echoes from rough surfaces continues to receive attention [11], [12]. Many underwater surfaces have a similar impedance to water so the problems are different to those in air where there is usually a large impedance difference across the air surface interface.

Rough surfaces can be periodically rough or randomly rough. Periodically rough surfaces have a strong repetitive pattern in the distribution of surface heights from the mean level and can be described with a geometric model alone. Random rough surfaces can only be described with statistical models [12].

A fundamental principle of ultrasonic sensing is that the components of an echo reflect the geometry of the insonified object. In this paper, we devise a theoretical basis for measuring surface roughness by developing a model that combines a transducer model, a model of acoustic reflection, and a model of surface geometry.

From this theoretical model, we identify which geometric features create an echo component and how. Then, we predict how that echo component appears in the echo and design a feature extractor to extract it from the echo. The result is a "spatial-angle-filter model" of CTFM ultrasonic sensing of random rough surfaces.

Thus, our aim is to develop a mapping from geometry to acoustic feature to signal feature that we can use to identify

Manuscript received December 20, 2004; revised May 19, 2006. The associate editor coordinating the review of this paper and approving it for publication was Dr. Bahram Kermani.

The authors are with School of Information Technology and Computer Science, University of Wollongong, Wollongong 2522, Australia (e-mail: phillip@uow.edu.au; keng4you@yahoo.com).

Color versions of Figs. 7, 9, and 11–13 are available online at http://ieeexplore.ieee.org.

Digital Object Identifier 10.1109/JSEN.2006.881419

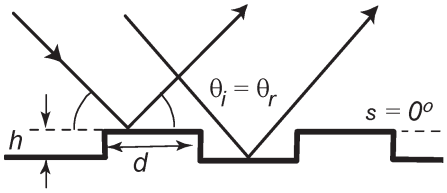

Fig. 1. Rayleigh model of scattering from a rough surface.

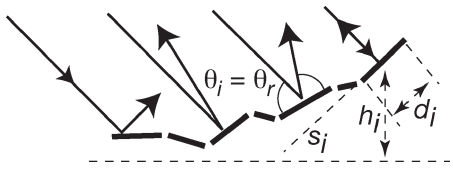

Fig. 2. Realistic model of a rough surface.

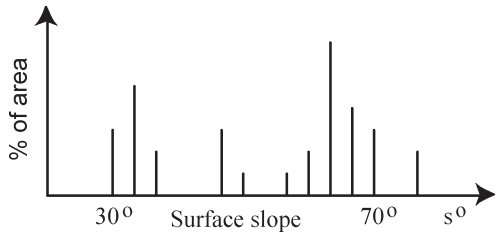

Fig. 3. Statistical model of surface slope—histogram of the percentage of surface area at each slope versus the angle (slope) of the facets.

suitable signal features. We experimentally measured the quality of each extracted feature with the statistics associated with calculating Mahalanobis distance. Finally, we iterated through the above process until we had a set of good features for measuring surface roughness that are grounded in a theoretical model.

In the process of this paper, we validated a theoretical basis for measuring surface roughness using ultrasonics in air. We can use the theory to explain the results of other researchers given their choice of features. Also, we can use the theory to choose and evaluate appropriate features for measuring surface roughness.

II. Models of Acoustic Reflection

When an ultrasonic sensor is pointed toward a surface (within $\pm\beta$ from the surface normal, where β is the beam angle) energy is reflected back to the sensor from the surface. The range measured is the orthogonal distance from the sensor to the surface. When the surface is smooth, energy is reflected from a single range giving a short duration high-amplitude echo to an impulse impingent on the surface. As the surface roughness increases, the range variation increases causing the echo duration to increase and the echo amplitude to drop. This is the basis of the ENDURA method [5] where high-speed sampling is used to measure the energy and duration.

When the sensor is at an angle greater than the beam angle $(\alpha > \beta)$ to the surface normal, a smooth surface will reflect most of the energy away from the sensor. As the surface roughness increases more energy is reflected back to the sensor. Echoes from both specular and diffuse regions contain information about surface roughness.

A. Rayleigh Model

The earliest and simplest model of surface roughness is the Rayleigh model [12] (Fig. 1). It models the surface as a set of elemental surfaces or facets, where a facet is a small flat surface region, and makes the following assumptions:

1) parallel incident rays and parallel reflected rays;
2) parallel facets with slope $s_i = 0^\circ$;
3) facets have equal areas: $a_1 = a_2$;
4) specular reflection: $\theta_\mathrm{r} = \theta_i$ is the same angle for all facets;
5) energy per unit area is the same for all facets.

The only effect modeled by the Rayleigh model is the phase shift caused by the varying depth of the facets.

B. Realistic Model

To produce a more realistic model of a rough surface, we relax the geometric assumptions except that the facets are assumed to be flat. First, we divide the surface up into a set of facets [13], each with its own height (h_i), length (d_i), and slope (s_i), as shown in two dimensions in Fig. 2. The surface spacing can vary from regular to random with the result that while the incident rays are still assumed to be parallel, the reflected rays are not. The reflection angle is a function of the surface slope s_i. In a three-dimensional (3-D) model, the surface also has orientation o_i and width w_i, and therefore, rectangular facets have area $d_i * w_i$.

Specular reflection occurs on each elemental surface with $\theta_\mathrm{r} = \theta_i$ relative to the surface but the absolute value of each angle of reflection depends on the surface slope s_i and orientation o_i. With this model, the effects modeled are phase shift and direction change.

C. Statistical Model

The more realistic model can provide an accurate model of reflection only when we know the parameters of each facet. Like most ultrasonic measurement problems, we do not know the geometry, but rather, we want to infer the characteristics of the geometry from the echo. In the echo, the reflections from the facets superimpose in all directions. Therefore, the echo reaching the receiver is a superposition of echo components from facets nearly normal to the axis of the receiver.

However, we can model the surface statistically. For each facet slope, sum the surface area, and normalize it as a percentage of the total surface area (Fig. 3). From this histogram, we can see how random the slopes of the facets are and whether any slopes dominate. To do this accurately, we need a way of measuring the geometric parameters of the facets. Alternately, we can use this model with a sensor that returns values proportional to the sum of the facet areas for each slope.

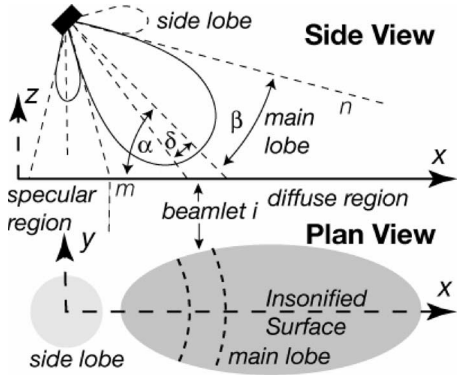

Fig. 4. Transducer model showing beam pattern and a single beamlet from a circular transducer. Dashed arcs are loci of constant range.

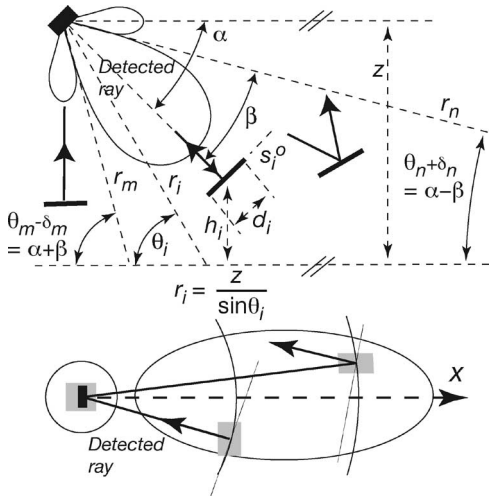

Fig. 5. Reception model showing that only energy reflected from near orthogonal surfaces returns to the receiver.

III. Transducer Model

A. Acoustic Transmission

We have stated that an ultrasonic sensor is a range area sensor; therefore, to determine whether an ultrasonic sensor returns a value for area suitable for use in the statistical model of surface slope, we need to combine it with a model of the transducer to obtain angle information. This combination further reduces the assumptions made with the Rayleigh model. The rays that model the ultrasonic waves are no longer parallel but emanate from a point on the surface of the transducer. Also, the beam pattern of the transducer (Fig. 4) results in the energy per unit area varying across the surface. This variation is accentuated along the surface (x-direction) by the slope of the sensor (α).

To model the transducer, we divide the beam into $n-m$ beamlets, from the bottom edge (m) to the top edge (n) of the main lobe, each with angle δ_i. Each beamlet i falls on a region i of the surface with an incident angle θ_i that varies from beamlet to beamlet. The area of surface covered by each beamlet is a sector of a circular shell (Fig. 4). The beamlet at the center of the main lobe is at angle α, the angle of the beam axis to the surface. Also, the amplitude of the energy in each beamlet A_i varies as the energy/unit area is a function of the transducer beam pattern.

The transmitter is at an angle α to the surface, so that the reflection from the main lobe that returns to a near coincident receiver is considered to be diffuse. Reflection from the first side lobe may be specular, depending on the angle of the axis of the side lobe to the surface. The surface pattern in Fig. 4 is for a side lobe near normal and a sloping main lobe. The actual pattern depends on the directivity of the transducer at the frequencies transmitted and the slope of the transducer.

B. Acoustic Reception

To model the energy returning to the transducer, we combine the realistic model with the transducer model to obtain an acoustic reception model (Fig. 5). The receiver only detects rays reflected at approximately 180° to their transmitted angle. Therefore, it only detects rays reflected by nearly orthogonal facets plus multipath echoes and diffraction. A rough surface will cause some rays to follow multiple reflections paths. Modeling them requires an accurate knowledge of the surface geometry [14], so in this paper we treat them as if they are direct echoes. For direct echoes, as θ_i varies from $\alpha+\beta^\circ$ to $\alpha-\beta^\circ$, the slope s_i° of the facets detected varies.

If we divide the main lobe $\alpha+\beta^\circ$ to $\alpha-\beta^\circ$ into $n-m$ equal angle beamlets of angle δ_i° then $\theta_i = \alpha+\beta-i*\delta_i$ for beamlet i. Therefore, reflections from facets in beamlet i with incident angle θ_i (where $\alpha-\beta<\theta_i<\alpha+\beta$) that are detected by the receiver are from surfaces with slope s_i that are nearly perpendicular to θ_i (within $\pm\delta_i/2^\circ$).

From this analysis, we conclude that the system is a "spatial-angle filter." At each range the amplitude of the echo is a function of the area of the facets that are nearly orthogonal to a ray from the transducer to the surfaces at that range. In the next section, we see how to extracts the spatial-angle filter values from a CTFM echo.

IV. Spatial-Angle-Filter Model

CTFM sonar periodically (every 100 ms) transmits a downward swept sine wave (100 to 50 kHz) [4], [15], [16]. The ultrasound reflects from objects and returns to the receiver as an echo. A demodulation sweep, derived from the transmitted sweep, is multiplied with the received echo in the time domain. The outputs of this multiplication are sum and difference frequencies.

The distance of flight information is contained in the difference frequencies (0 to 5 kHz), where frequency is proportional to range and amplitude is proportional to surface area. With a sweep time of 100 ms to sweep over 50 kHz, a maximum demodulation frequency of 5 kHz is equivalent to a range of 1.718 m when the speed of sound is 343.6 m/s at 20 °C.

This time-domain signal is converted to a power spectrum with a fast Fourier transform (FFT) to give a range-energy echo where the amplitude in frequency-bin i is the energy per unit time reflected at range r_i (Fig. 6). As a 1024 point FFT produces 512 range bins, each bin represents a 3.355-mm increment in range.

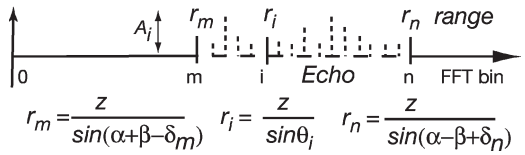

Fig. 6. CTFM echo is converted to a frequency spectrum where frequency is proportional to ranges to reflecting objects.

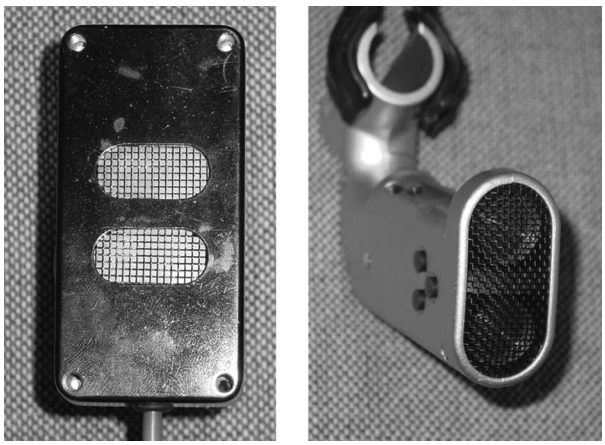

Fig. 7. Left: Prototype CTFM sonar sensor used to collect data showing oval transmitter and receiver. Right: Recently, released commercial CTFM mobility aid for blind people (Ksonar) with circular transducers.

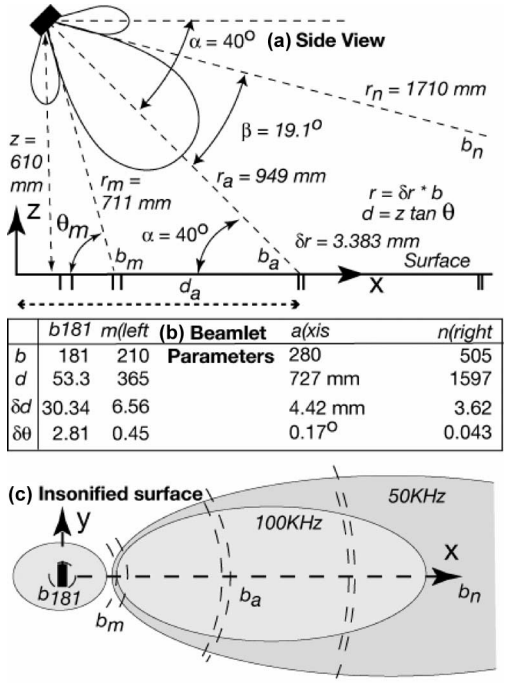

	b181	m(left	(b) Beamlet Parameters	a(xis	n(right
b	181	210		280	505
d	53.3	365		727 mm	1597
δd	30.34	6.56		4.42 mm	3.62
$\delta\theta$	2.81	0.45		0.17^o	0.043

Fig. 8. Geometry of sensor system used in experiments for transmission. (a) Side view at mid frequency of 75 kHz. (b) Ranges and angles of beamlets at mid-frequency. (c) Ovals show area of ensonification at ends of sweep.

Therefore, amplitude A_i in bin i is a measure of the energy from surfaces in beamlet i at range r_i whose slope $s_i \pm \delta_i/2$ is nearly orthogonal to θ_i. Hence, CTFM gives the output of a spatial-angle filter, where the amplitudes of bins b_m to b_n for ranges r_m to r_n are proportional to the areas of surfaces at angles s_m to s_n in the regions insonified by $n - m$ associated beamlets at angles θ_m to θ_n. The output of the spatial-angle filter produced by CTFM (Fig. 6) is analogous to the statistical model (Fig. 3).

There are some minor differences to the theoretical model in the previous section. In the theoretical model, the beamlets are all of equal angle. As the power spectrum bins are based on range, each bin represents an equal range increment. Therefore, the beamlet angles decrease as you move away from the sensor ($\delta\theta$ in Fig. 8). As a result, the range of slope angles and the area of the surface (δd in Fig. 8) in each beamlet decreases with distance from the transducer.

Consequently, the angular resolution of the spatial filter increases with distance from the sensor. As only a few bins (b in Fig. 8) fall in the specular region (about 30) compared to the diffuse region (about 300) the filter contains an order of magnitude less information in the specular region.

The sensor used in these experiments consists of two transducers (Fig. 7, left) arranged as transmitter and receiver. A recently released mobility aid for blind people has a similar transducer layout. As transducers have surface area, we model them with the radiating plane piston model and not with point models. The surface of the sensor has a large target area for the returning echo to hit. As we model a facet as a point, the sensor detects echoes over a wider range of facet angles blurring the sharpness of the spatial-angle filter.

V. Sensing Strategies

Sensing strategies include the design of the geometric layout of the transducers and the control of sensor motion. The transducers used to collect the data (Fig. 7, left) are 30-mm wide by 15-mm high with 22 mm between their centers. At 75 kHz, they have a vertical beam angle of 19.1° and a horizontal beam angle of 8.8° where the beam angle is from the axis of the main lobe to the first minima.

In our experiments, we placed the sensor at the front of a mobile platform with its axis sloping down at 40° to the horizontal. As the transducer sweeps from 100 down to 50 kHz the beam angle (along the x-axis away from the sensor) varies from 14.2° to 29.4°. At the left hand end (Fig. 8) θ_m varies from 54.2° to 69.4° with an angle of 59.1° at 75 kHz and at the right hand end θ_n varies from 25.8° to 10.6° with an angle of 20.9° at 75 kHz.

The sensor is placed to give good coverage of the surface in front of it. The varying beam angle that occurs with frequency sweep fills the minima between the main lobe and side lobe with energy and in effect smears out the side lobes. As the power spectrum is integrated over most of the frequency sweep, the echo includes significant information from the region where the minima are located. As an approximation, we use the beam angle at the middle frequency (75 kHz) to define a transition point between the main and side lobe.

The table in Fig. 8 presents the variation in the geometry of the beamlets at a center frequency of 75 kHz as we move along the surface away from the normal to the transducer. The range r_i from the transducer to the surface changes from 610 mm on the normal to 711 mm at the left edge of the main lobe to 1710 mm at the right edge. The corresponding frequency bin numbers

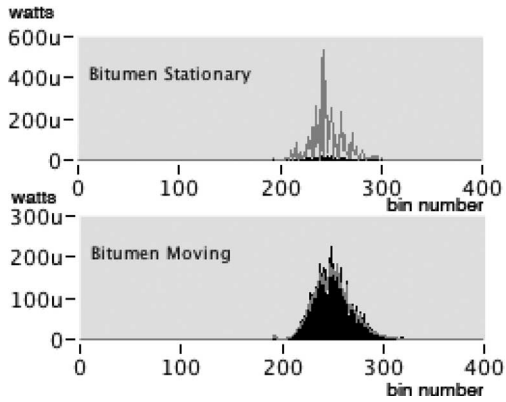

Fig. 9. Mean (gray lines) and standard deviation (black bars) of power spectrum of 64 echoes from bitumen surface for a stationary (top) and a moving (bottom) sensor. Values in microwatts. Only 400 of the 512 bins are shown.

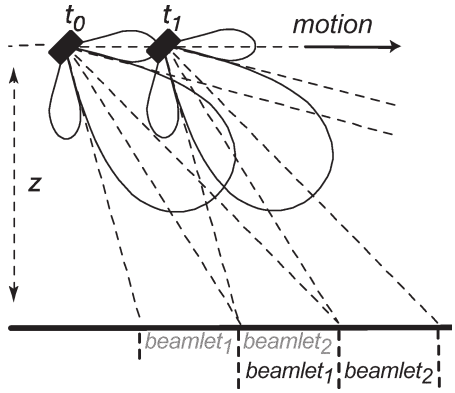

Fig. 10. Moving the sensor moves the beamlets to different regions of the surface.

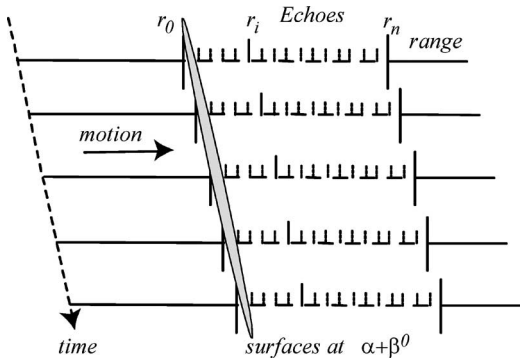

Fig. 11. Averaging echoes from beamlets defined by FFT bins as the sensor moves to obtain a better image of surface roughness.

b range from 181 to 505. Bins 181 to 210 are considered to be in the specular region and bins 211 to 505 in the diffuse region.

When the receiver is 610 mm above the surface, bin 180 includes the transition from the normal to the surface with a distance d of 53.3 mm along the surface from the normal to the first bin on the surface. The angle of the beamlets $\delta\theta$ in the diffuse region varies from 0.45° to 0.043°. The areas on the surface covered by each bin are sectors of circular shells of width δd_i, as shown at the bottom of Fig. 8.

A. Stationary Sensor

The amplitudes of the power spectrum bins vary from one sensor sweep to the next due to electronic noise, acoustic noise, surface roughness, and sensor motion that is not parallel to the surface. When the sensor is stationary, the sensing geometry is fixed so the echoes of repeated sweeps are from the same surface region. Hence, each bin contains a measure of the acoustic reflectors at a given angle for a fixed, small area of the surface. Averaging over multiple echoes will reduce the noise to reveal information buried in the noise for that surface region only. Fig. 9 compares the mean (gray curve) and standard deviation (black bars) of 64 echoes from a bitumen surface measured with a stationary and a moving sensor. The top trace shows that the standard deviation for a stationary sensor is small, demonstrating the repeatability of the echo.

Any variation in the amplitude of the echoes from the surface from one sweep to the next is due to electronic noise or to changes in air conditions, such as microconvection currents. We can consider these to be acoustic noise. We can minimize the effect of acoustic noise by averaging over p echoes.

B. Moving Sensor

When the sensor is translated across a surface, each beamlet senses a different region of the surface on each sweep. Fig. 10 shows the situation where the sensor translates by the width of a beamlet on each sweep. By moving the sensor and averaging the readings over several scans (Fig. 11), we can build up an image of a larger region of the surface. In this way, the spatial-angle filter can be moved over a surface to produce an echo representative of that surface.

As for a stationary sensor, averaging reduces acoustic and other forms of noise so that the variations in the echo represent variations in surface geometry (or roughness). The bottom trace of Fig. 9 shows that the standard deviation for a moving sensor is large, demonstrating that the geometry of the surface changes significantly in each beamlet as it moves over the surface.

However, averaging can also result in the sensor being slow to pick up genuine changes in roughness making it less useful when classifying surfaces for landmark navigation. On a very rough surface, the transducer can move in the vertical direction as the platform moves across the bumps changing both the sensor's height and orientation. In that case, the averaged signal contains information about the sensor motion (due to surface roughness) as well as surface geometry. This will blur the sharpness of the spatial-angle filter.

C. Range Normalization

Surface classification may be made more robust to signal amplitude variations when comparing measurements taken at different ranges by normalizing the echo before extracting features. An obvious way to normalize the echo is to convert

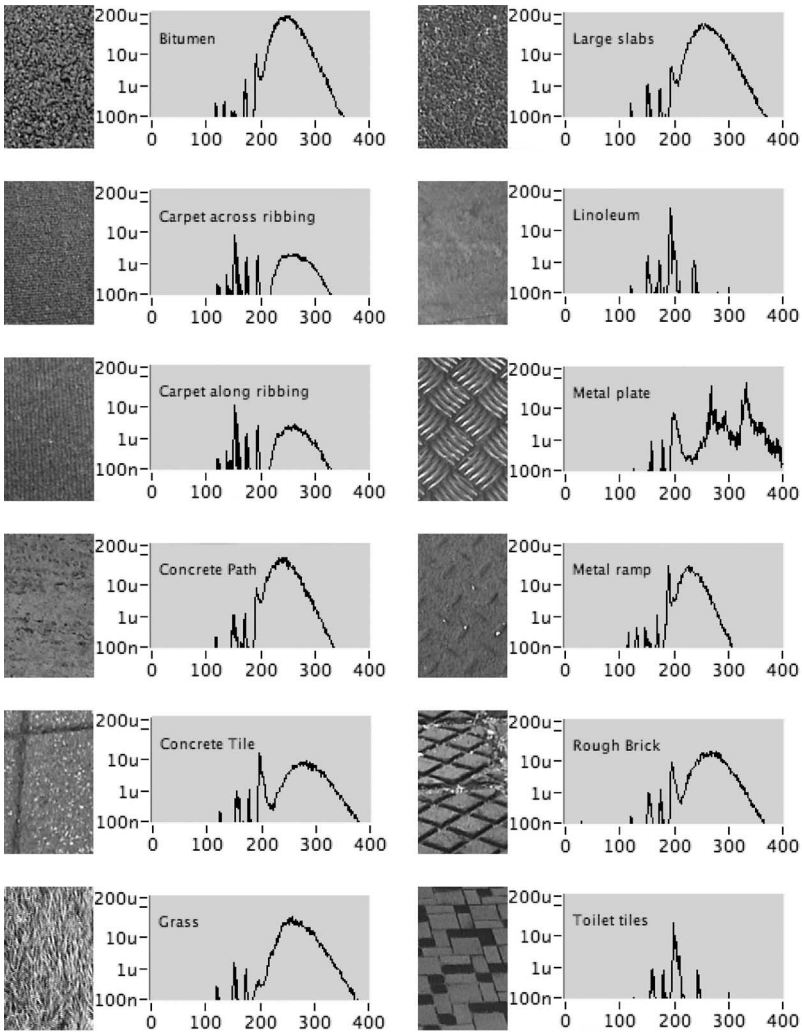

Fig. 12. Images and means of power spectrum of 64 echoes from 12 surfaces. Vertical axes in microwatts and horizontal axes in bin numbers.

the individual bin amplitudes to percentages of the total energy (Fig. 3). Normalization involves three steps for any surface.

1) Consider only the region of the surface producing diffuse reflection by determining the length of the surface (b_m to b_n) over which the spatial-angle filter detects echoes (length of acoustic density profile [4]—feature 3 in Section VII).
2) Sum the amplitudes of the bins in this window

$$\text{Total energy } T = \sum_{i=m}^{i=n} A_i. \tag{1}$$

3) Normalize spectrum

$$a_i = \frac{A_i}{T} \times 100\%. \tag{2}$$

However, normalization has the problem that changes in total energy (T) due to changes in roughness as the sensor moves over the surface can be lost. A solution to the problem of data loss with both averaging and normalization is to calculate a running average of both the bin values and the total energy and track their trends. The number of sweeps over which the averages are calculated is a compromise between speed of response and robustness of measurement.

Calculating a running average of the normalized echoes has the following benefits:

1) reduced noise when the sensor is stationary;
2) angle/area data are averaged over a larger surface area when the sensor is moving and is a more accurate measure of surface roughness;
3) integration of roughness change from one surface type to the next with motion smoothes out the noise caused by spurious changes but still detects real changes;
4) reduced effect of small surface changes due to joins.

Normalization was not used in the experiments discussed later in this paper, because the range was the same for all measurements. Therefore, we do not present quantitative results to validate the claim that it will increase robustness when measuring roughness. In previous research [4], we showed that normalization increased the robustness of plant sensing to variations in range to the plant.

VI. WINDOWING

The first step in feature extraction is to select a region of the echo that is of interest: either the specular region (under a side lobe) or the diffuse region (under the main lobe). Selecting the window of interest is usually done with a threshold function.

We used a two-stage windowing function: geometry followed by threshold. In the first stage, the diffuse region of the power spectrum (b_m to b_n is selected using the geometry of the sensing system (Fig. 8). As the transmitted signal sweeps down, the bin number of the left edge of the main lobe reduces from $b_m = 222$ to 192. At the center frequency of 75 kHz, the left edge is at bin $b_m = 210$ and the right hand edge at bin $b_n = 505$.

We do not expect any reflection at all in bins 0 to 180, because they represent the space between the transducer and the ground. We expect reflection in bins 181 to 192 to be from the first side lobe. We expect reflection in bins 192 to 222 to be a mixture of main lobe and side lobe reflections. Above bin 222, all reflection is from the main lobe.

In the experimental data (Fig. 12), we observe that linoleum (a very smooth surface) has a high percentage of energy in bins lower than $b_m = 200$ (specular), and a very rough surface (large slabs and grass) has a high percentage of energy in bins higher than $b_m = 220$ (diffuse). This is to be expected from the model. Significant energy appears in bins 200 to 220 for surfaces that contain joins (tiles).

The first windowing stage removes the bins to the left of $b_m = 210$. The second stage of windowing applies a threshold to the bins that remain in an attempt to reduce the window to the set of bins that represent the slopes of the facets that make up the surface.

A. Thresholding and Noise

The maximum amplitude of individual echoes ranges from 1 μW for smooth surfaces to over 1 mW for very rough surfaces. For several surfaces, this maximum occurs in the specular region, so when these bins are removed in the first stage of windowing, only a very small amount of energy is left. For example, the echo from linoleum in the diffuse region has a maximum of 2 μW. To detect this energy the system requires a threshold less than 1 μW.

The best way to determine a suitable threshold value is to measure the noise in the system [4]. Setting the threshold at the mean plus five standard deviations $(\mu + 5\sigma)$ will remove 99.999942% of the noise (that is it will let through 1 in 1 000 000 noise bins). Whether this level of separation between signal and noise can be achieved depends on the level and structure of noise in the system.

Measurement noise has five potential sources: electronic noise, acoustic noise, direct coupling of transmitter output to receiver input (crosstalk), faint echoes from other objects in the environment, and aliasing of echoes from far away objects because their demodulated frequency is greater than the Nyquist frequency of the analog to digital converter. To measure the noise we point the sensor into open space in an indoor environment and record 64 echoes. The mean and standard deviation of these echoes is shown in Fig. 13.

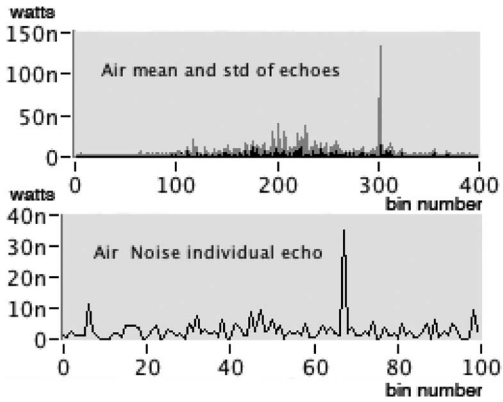

Fig. 13. Echo of noise recorded with sensor aimed horizontally into the air in an indoor environment; Top trace—mean (gray) and std dev (black) of 64 echoes; Bottom trace—power spectrum of lines 50 to 150 for an individual echo. Values in nanowatts.

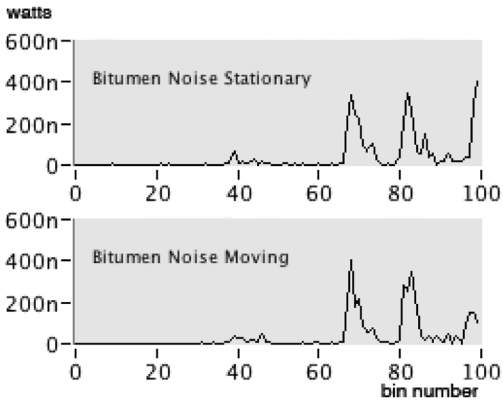

Fig. 14. Power spectrum lines 50 to 150 of Echo signals, representing the free space between the transducer and the surface, showing the noise in the echoes for both stationary and moving sensors. All values in nanowatts.

From this figure, we can see that there is no cross coupling, there is a faint echo component from something at bin 300, and the electronic and acoustic noise is low. This noise has a maximum mean of 9 nW and a standard deviation of 13 nW. If we consider the free space region (bottom of Fig. 13), we get a mean of 6.25 nW with a standard deviation of 6.75 nW. These values result in potential thresholds of 74 nW for the whole echo and 40 nW for the echo in the free space region.

However, when the sensor is in use, the geometry of the sensor system relative the environment is different, so we see some noise in the free space bins (50 to 150) from echoes from objects other than the surfaces of interest. Compared to Fig. 13, the free space portion of the echoes when sensing bitumen (Fig. 14) show considerably more noise. For bitumen, the mean is 36.5 nV, and the std is 77.3 nV. These values result in a $\mu + 5\sigma$ threshold of 423 nV. Based on these noise calculations, we chose a threshold of 400 nV for our experiments.

A second problem is that low-energy reflections from the surface may be below the threshold. For example, the metal plate has a lot of bins with values between 500 nV and 1 μV (Fig. 15), while the linoleum has a lot of bins with values between 100 and 200 nV. If these are between bins with amplitudes larger than the threshold, they will be included in

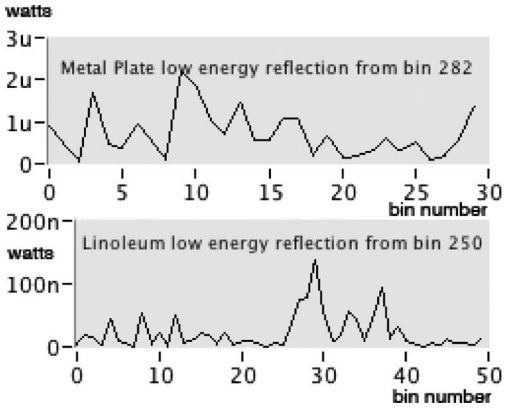

Fig. 15. Large differences between low-energy components of the echoes from different surfaces make threshold setting difficult. Top: echo component between bins 273 and 303 from metal plate. Bottom: echo component between bins 231 and 282 from linoleum. Values in nanowatts.

TABLE I
MEANS (TOP ROWS) AND STANDARD DEVIATIONS (BOTTOM ROWS) FOR FIVE FEATURES OF 64 ECHOES FROM A MOVING SENSOR FOR BITUMEN AT FOUR THRESHOLDS

Threshold	F1	F2	F3	F4	F5
100n	59.16	30.05	29.12	7.25m	2.75m
	49.7f	705m	705m	1.19m	318m
400n	59.16	32.03	27.12	7.25m	3.09m
	56m	695m	684m	1.19m	344m
1u mean	59.09	33.07	26.02	7.25m	3.29m
std	198m	610m	619m	1.19m	370m
10u	58.11	36m	22.11	7.24m	4.05m
	961m	779m	1.16m	1.19m	482m

the window, but if they are at the start or end of the window, they may be removed by the threshold.

The values of the features (Table II) and their standard deviations change with the threshold (Table I). For example, feature 3 (range of angles) varies from 22.11 bins at a threshold of 10 μW to 29.12 bins at a threshold of 100 nW. When no values in the echo are above the threshold, a window for feature extraction is not found, and the echo has to be removed from the set. At a threshold of 400 nV, one echo is discarded from the set of 64 echoes recorded for linoleum from a moving sensor.

VII. FEATURE EXTRACTION

In order to classify and recognize surfaces based on roughness, we have to extract features from the echo: features whose parameter values vary from surface type to surface type and have low standard deviation [4]. Before discussing features derived from the spatial-angle filter we will look at those used by other researchers.

In the ENDURA method, Bozma and Kuc [5] used a rotating sensor to measure two features: echo energy and echo duration. Both these features are a function of surface roughness, range, and orientation and are useful for measuring surface roughness for mapping the environment. Their measurements included all the energy in the echo (both specular and diffuse).

TABLE II
ELEVEN FEATURES DERIVED FROM THE SPATIAL-ANGLE-FILTER MODEL FOR EXTRACTION FROM THE DIFFUSE REGION OF THE ECHO

1	start angle θ_m (bin number b_m)
2	end angle θ_n
3	l = range of angles ($\theta_n - \theta_m$) = length of echo window
4	T = sum of the amplitudes A_{tot} of the power spectrum lines in window = total energy – Eqn.1.
5	D = energy distribution across window – Eqn. 3.
6	Position of peak energy – angle with most facets
7	Amplitude of peak – area of facets - energy reflected from those surfaces
8	Percentage of energy in the peak – significance of angle
9	Position of second peak
10	Amplitude of second peak
11	Distance between the two peaks – indication of presence of repetitive pattern

Probert Smith and co-workers [6]–[8] also used a CTFM system transmitting at an angle to the surface. They used a set of three features: energy in the specular region, energy in the diffuse region, and range of bins in the diffuse region above a threshold. The third feature was found to be sensitive to the value of the threshold so they replaced it with the standard deviation of the intensity in the diffuse part divided by the mean (coefficient of variation).

The purpose of the spatial-angle-filter model is to provide a theoretical basis for feature extraction that enables us to predict the geometric meaning of potential features. In this section, we derive 11 features from the model that we test empirically in Section IX.

From the spatial-angle-filter model, energy aimed directly at the surface will detect facets that are nearly parallel to the surface and corner reflectors produced by the intersection of facets. The FFT used to calculate the power spectrum has a bin resolution of 3.383 mm, so the surface has to be very rough for the duration to be greater than one bin. For this reason, we developed a set of features for the diffuse region only and windowed out the specular region.

The spatial-angle filter gives a range r_i, facet angle 90-θ_i, and energy a_i (proportional to area) for each bin of the FFT output. After the echo is windowed, the window-start and -end positions measure the range of angles for the slopes of the facets in the surface. Therefore, the three features are 1) start angle θ_m (bin number b_m), 2) end angle θ_n, and 3) range of angles $(\theta_n - \theta_m)$ or length l of the echo window (acoustic-density profile [4]). One of these three features is redundant therefore all three cannot be used in classification. A set of features is given in Table II.

The sum of the amplitudes $A_{\rm tot}$ of the power spectrum lines in the window (1) is the total energy T reflected by the diffuse region of the surface to the receiver (feature 4) and is proportional to the area of the facets reflecting that energy. The numeric value of this feature should increase with roughness.

If we divide the total energy T (4) by the length of the window l (3), we get the average energy per bin, which can be thought of as the width of an approximating rectangle. However, in doing this, we lose the information about the

TABLE III
NORMALIZED REFERENCE VECTORS—VALUES OF 11 FEATURES FOR 12 SURFACES CALCULATED FROM 64 ECHOES FROM A MOVING SENSOR FOR EACH SURFACE. EACH VALUE IS THE MEAN OF THE FEATURE NORMALIZED BY DIVIDING IT BY THE STANDARD DEVIATION, SO ALL VALUES HAVE A STANDARD DEVIATION OF 1

Feature / Surface	1	2	3	4	5	6	7	8	9	10	11
Bitumen	1055.57	46.116	39.627	6.087	8.986	17.836	18.519	2.681	3.174	4.273	22.746
Carpet Cross	41.648	49.272	13.581	9.408	10.395	16.709	13.954	3.658	5.325	5.908	21.477
Carpet Long	53.651	40.601	14.399	5.227	7.105	17.937	15.521	3.291	3.342	4.05	21.332
Concrete Path	1.2E+15	51.04	38.804	2.999	8.204	18.833	16.745	2.012	2.202	4.652	23.644
Concrete Tiles	43.751	37.627	18.508	6.959	9.872	18.644	17.022	3.258	4.006	4.668	24.018
Grass	53.706	48.556	23.273	2.396	7.279	14.333	16.197	1.411	1.626	4.412	18.797
Large Slabs	314.977	56.048	51.139	7.8	9.286	19.36	20.21	3.192	3.709	4.701	24.837
Linoleum	30.624	116.483	0.813	3.684	1.079	342.455	24.153	2.157	2.725	5.645	47.009
Metal Plate	58.845	10.693	11.048	0.619	2.195	5.159	5.755	0.586	0.589	1.716	7.44
Ramp	1055.57	57.855	34.695	4.081	8.837	23.19	19.804	2.413	2.986	4.316	30.376
Rough Brick	73.332	43.037	24.998	3.916	6.662	14.143	15.368	1.836	2.209	3.435	17.849
Toilet Tiles	29.676	14.365	2.957	4.944	0.4	21.081	10.205	2.287	3.083	4.909	20.546

distribution of energy in the echo. Two surfaces can have a similar average but a very different energy distribution

$$D = (T/l)/\sigma_a. \tag{3}$$

If we divide the width by the standard deviation of the amplitudes (3), we get a value representing the distribution of energy (feature 5). If the energy profile has a single large peak such that most of the energy is in a few lines, due to the facets all having a similar slope, then the standard deviation will be high and the value for feature 5 will be low. When the energy is spread across the window, by facets with many angles, the standard deviation will be low and the value for feature 5 will be high.

In the echo, a bin with a large amplitude has a large percentage of the surfaces at that angle. Two features that may be useful are the position of the peak (6) (proportional to the slope of the facets) and the amplitude of the peak (7) (proportional to the area of the facets at that slope). More useful than the amplitude may be the percentage of the energy in the peak (8).

The next set of features will discriminate between surfaces that have two regions where the energy peaks. Feature 9 is the position of peak 2 and feature 10 is the amplitude of peak 2. The features of interest for discrimination are the differences between the peaks. The distance between the peaks (feature 11) shows how far apart the two angles are. When it is small the peaks are probably from the same surface (could be FFT leakage) and when it is large the reflections are from facets with different angles. In the latter case, the difference between the percentages of the energy in these two peaks may be a useful feature. For example, the metal plate in Fig. 12 has two distinct similar sized peaks due to its repetitive structure.

Table III shows the values for 11 features for all 12 surfaces as 12 reference vectors. Each feature value was computed by extracting feature values from 64 echoes, calculating their means and standard deviations, and then normalizing the means by dividing them by the standard deviations to produce a set of feature vectors where every value has a standard deviation of 1 (Section IX). The results presented in the following sections for a moving sensor are derived from these 12 reference vectors.

VIII. CLASSIFICATION

From the spatial-angle-filter model, we defined the 11 features in Table II. Each feature has a physical meaning corresponding to a characteristic of the roughness of a surface. Most are different to those we derived for the acoustic-density-profile model for plant sensing [4]. How many of these features do we need to correctly recognize these surfaces with a classifier? To answer this question, we will use the Mahalanobis distance to develop a measure of the quality of the classification as well as to do the classification.

A. Pattern Classification

Pattern classification is the task of giving names to surfaces based on measurements of ultrasonic echoes from those surfaces. A set of n features that have been extracted from an echo are combined into a feature vector

$$\boldsymbol{v} = [v_1 \ldots\ldots v_n]^{\mathbf{T}}. \tag{4}$$

The feature vector $\boldsymbol{v}$ points to a point in the n-dimensional feature space $\mathbf{V}$. Every point in feature space $\mathbf{V}$ corresponds to one of the possible constellations of the measurement data. We seek to select the set of features such that vectors of feature

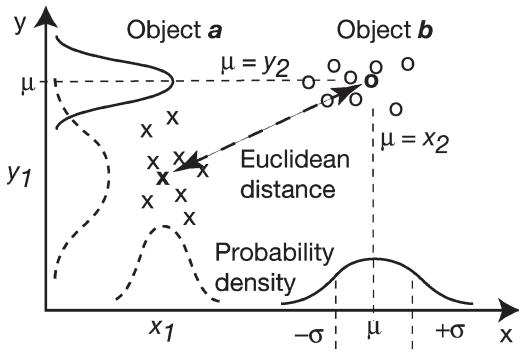

Fig. 16. Euclidian distance (7) between mean vectors composed of values for two features for clusters of measurements from two surfaces.

data for each surface cluster together with minimum overlap with clusters for other surfaces.

To uniquely identify a surface, we calculate a reference vector with the cluster of echoes from that surface. We model the measured features for a surface with a probability density function. Vector $\boldsymbol{v}$ has mean vector $\boldsymbol{\mu}$ of the features that represents the cluster, a standard deviation vector $\boldsymbol{\sigma}$ that represents the spread of the cluster in feature space and a covariance matrix $\mathbf{K}$. For feature vectors containing two features (a 2-D feature (x, y) space), the vectors are

$$v_i = \begin{bmatrix} x_i \\ y_i \end{bmatrix} = [x_i \quad y_i]^{\mathrm{T}} \tag{5}$$

and the mean vector (x_μ, y_μ) (or first-order moment) is

$$\mu = \begin{bmatrix} x_\mu \\ y_\mu \end{bmatrix} \qquad \text{where } x_\mu = \frac{\sum_{i=1}^{n} x_i}{n}. \tag{6}$$

In our experiments, we collected 64 echoes from each surface (Fig. 12). From each echo, we extracted values for 11 features (Table II) and calculated the means and standard deviations for the cluster of 64 echoes from each surface (Table III). Thus, we have potential for an 11-dimensional space representing 12 surfaces, which is large.

There is a tradeoff between vector dimensionality and computational complexity. We wish to reduce the set of features to the minimum required to separate the clusters. To reduce the feature set, we need a way of determining the quality of features for classification. A common way is to calculate the Euclidian distance between reference vectors and then choose the set of features that give the largest distances.

The Euclidian distance is the physical distance between the mean vectors. For the surfaces $\mathbf{a}$ and $\mathbf{b}$ in 2-D space (Fig. 16) the mean vectors are $[x_1, y_1]^{\mathbf{T}}$ and $[x_2, y_2]^{\mathbf{T}}$ and the Euclidean distance is

$$d_{\mathrm{e}} = \sqrt{(x_1 - x_2)^2 + (y_1 - y_2)^2}. \tag{7}$$

The disadvantage of the Euclidean distance is that it is a linear classifier that assumes that the standard deviations of each feature are the same. Therefore, the variations between the measured vectors are represented with an error circle around the mean (for vectors of two features).

The standard deviation measures the spread of the cluster and, hence, gives a measure of the quality of a feature (Fig. 16). Assuming that the variation in a feature fits a Gaussian distribution (bell curve), the area under the curve between points p_1 and p_2 is the probability that a random value p is between p_1 and p_2, where the total area under the curve from $-\infty$ to $+\infty$ converges to 1.0. As the standard deviation σ increases, the curve gets wider and flatter. The slope of the curve is steepest at σ either side of mean (Fig. 16).

The variations of the measurements from the mean are represented as error ellipses (in 2-D space). Thus, including the standard deviations in the distance calculations will result in more accurate linear classification. When there is no correlation between the features, the axes of the error ellipse are parallel to the axes of the graph. When there is cross correlation the axes are at an angle ($45^\circ = 0.95$ cross correlation).

B. Mahalanobis Distance

A measure that does include the standard deviations of the features is the Mahalanobis distance (8). It includes the probability density functions of the features through use of the covariance matrix $\mathbf{K}$ (9). It measures the distance between feature vectors in the cluster and the mean vector μ, for which the covariance has been measured, in units of standard deviations not in physical distance.

$$d_m = (v - \mu_n)^{\mathrm{T}} \mathbf{K}^{-1} (v - \mu_n). \tag{8}$$

The covariance matrix has some useful properties for analyzing features. The diagonal terms in the covariance matrix are the variances of the features (9) and (10), therefore they show the spread of the feature values. The off-diagonal terms are the covariances between the features. When the off-diagonal terms are zero (10) the features are uncorrelated. Equation (9) gives the general covariance matrix $\mathbf{K}$ (or central-second-order moment) for a cluster with two features, where n is the number of samples

$$\mathbf{K} = \frac{1}{n} \sum_{i-1}^{n} (v_i - \mu)(v_i - \mu)^{\mathrm{T}} = \begin{bmatrix} \sigma_x^2 & \sigma_{xy} \\ \sigma_{yx} & \sigma_y^2 \end{bmatrix}. \tag{9}$$

Equation (10) gives the Mahalanobis distance for a cluster with three features where there is no correlation between the features. When the cross correlation between two features is negative the feature values increase together. When the cross correlation between two features is positive, the first feature decreases when the second increases

$$d_m = [x \quad y \quad z] \frac{1}{\begin{bmatrix} \sigma_x^2 & 0 & 0 \\ 0 & \sigma_y^2 & 0 \\ 0 & 0 & \sigma_z^2 \end{bmatrix}} \begin{bmatrix} x \\ y \\ z \end{bmatrix}. \tag{10}$$

If the covariance matrix cannot be inverted, then we have a singularity. One cause of a singularity is redundant features. A redundant feature will not improve the classification and can be

removed. Another error that can cause the distance calculation to blow up is a zero standard deviation.

The Mahalanobis distance is measured in units of the standard deviation of the reference cluster. Each surface has a different covariance matrix; therefore, the distance from the reference vector of surface A to the reference vector of surface B is different to the distance from B to A.

When the variances are 1, and the covariances are 0, the covariance matrix becomes the identity matrix and the Mahalanobis distance becomes the Euclidian distance. We can reduce the Mahalanobis distance to the Euclidean distance by normalizing the feature values by dividing them by their standard deviations. The division results in a new set of feature vectors all with a standard deviation of $\sigma = 1$.

One result is a common covariance matrix for all surfaces where all features have the same size error circles. A second result is a Euclidean distance measure in units of standard deviation. A third result is that the normalization of the features maps them all into the same dimensions, making features easier to compare.

The Mahalanobis distance (or Euclidean distance of the normalized vectors) gives us a measure of the quality of the feature set for classification. As the unit of distance is the standard deviation of the set of measurements in the reference cluster, the distance tells us the probability of a new measurement being from that surface. 68.27% of values are within 1σ of the mean vector, 95.45% within 2σ, 99.73% within 3σ, 99.9937% within 4σ, and 99.999942% within 5σ. Therefore, when comparing a measurement vector to a set of mean vectors, the Mahalanobis distance tells us which surface the measurement is closest to and the probability that the echo is from that surface.

IX. Experimental Results

We recorded echoes from 12 objects (Fig. 12—two are carpet tiles: one across the ribbing and one along it). For each object, we recorded 64 echoes from a stationary sensor and 64 echoes from a moving sensor. The velocity of motion was not recorded and it would have varied over the 6.4 s taken to record a set of echoes (one echo was captured every 100 ms). For these 12 surfaces, we calculated normalized reference vectors (Table III).

To measure the quality of a feature, we calculated the Euclidean distance in a one-dimensional (1-D) space for that feature between all pairs of the 12 surfaces and recorded the minimum distance In Table IV, the features are ranked in descending order of minimum distance and, hence, descending order of quality.

The start angle (feature 1) is a poor feature because it hardly varies, due to the geometry of the sensing configuration (Fig. 8). This same geometry results in large variations in the end angle (feature 2) with roughness. The range of angles (feature 3) is the best feature as it measures the range of the orientations of the surface facets, which we expect to vary with roughness.

The energy reflected back toward the transducer (feature 4) is expected to increase as the roughness increases, because the sensor is angled so that most of the reflection from a smooth surface will reflect away from the sensor. Three other features (5, 7, and 11) are of similar quality. Feature 5 (energy distribution) measures the distribution of the orientations of the facets. Feature 7 (amplitude of peak 1) measures the energy reflected from the most common facet orientation.

TABLE IV
Eleven Features in Decreasing Order of Minimum Distance Between Surfaces for 1-D Classification, Also Showing the Number of Distances Less Than 6σ of a Total of 66 Distances

Feature Number	Minimum distance	Count $<6\sigma$	Feature
3	0.818σ	11	range of angles $(\theta_n - \theta_m)$ l
2	0.716σ	13	end angle θ_n
4	0.165σ	60	sum of the amplitudes T
7	0.152σ	41	Amplitude of peak 1
5	0.149σ	44	D = energy distribution
11	0.144σ	34	Distance between peaks
6	0.101σ	40	Position of peak 1
9	0.091σ	66	Position of peak 2
1	0.055σ	5	start angle θ_m
8	0.034σ	66	% of energy in peak 1
10	0.016σ	66	Amplitude of peak 2

Feature 11 (distance between peaks) measures the distance between the bins with the two highest amplitudes. These peaks are single bin values; therefore, they represent the two most common facet orientations. Features 9 and 10 (position and amplitude of peak 2) are of lower quality. In future research, we will trial alternate algorithms for peak detection, which define peaks with groups of bins, to see if we can improve the quality of these features. The percentage of energy in peak 1 (feature 8) is a poor feature indicating that several surfaces have similar peak reflection values (Fig. 12).

To achieve a recognition rate of 99.73%, each cluster must include a region of 3σ around the reference vector. This maps to a distance of 6σ between reference vectors. It is obvious from Table IV that a single feature cannot discriminate between the surfaces. Using features 2 and 3 as 2-D reference vectors (best quality in Table IV), three of the 66 distances are less than 6σ with a minimum distance of 4.99.

With 3-D reference vectors consisting of features 2, 3, and 4, there are still three distances below, 6 but they have increased in value (5.07 to 5.36, 4.99 to 5.87, and 5.78 to 5.98), as shown in Table V. Moving to a four-dimensional (4-D) space with either feature 5 or 7 reduces the number of distances less than 6 to 1 (5.92 and 5.57, respectively). A 4-D space with feature 11 has two distances under 6 but the minimum is higher (5.94).

To achieve all distances greater than 6σ and hence a recognition rate of 99.73%, we need to use a five-dimensional (5D) vector composed of features 2, 3, 4, 5, and 6, or 7 or 11. These results verify that the quality measure in Table IV gives an excellent guide to feature selection. The highest quality features are based on range (angle) values not on amplitude (energy) values, which is consistent with our understanding of ultrasonic sensing.

Table V shows all distances between 12 surfaces (66 distances) for both moving and stationary sensors for vectors composed of the three highest quality features (2–4). The infinite distances for toilet tiles with a stationary sensor is a result of

TABLE V
EUCLIDEAN DISTANCES FOR MOVING (TOP RIGHT), AND STATIONARY (BOTTOM LEFT) SENSORS IN THE 3-D SPACE DEFINED BY FEATURES 2 (END ANGLE), 3 (LENGTH OF PROFILE) AND 4 (SUM OF ACOUSTIC DENSITY PROFILE). THE BOLD DIAGONAL VALUES ARE THE DISTANCES BETWEEN MOVING AND STATIONARY. ALL VALUES IN UNITS OF σ

Moving / Stationary	Bitumen	Carpet Cross	Carpet Long	Concrete Path	Concrete Tiles	Grass	Large Slabs	Linoleum	Metal Plate	Ramp	Rough Brick	Toilet Tiles
Bitumen	**2287.6**	26.446	25.838	5.871	22.778	16.942	15.3	80.398	45.842	12.89	15.106	48.519
Carpet Cross	2267.89	**44.3**	9.661	26.084	12.879	11.984	38.198	68.652	39.648	23.407	14.12	36.759
Carpet Long	2285.78	37.38	**69.6**	26.637	5.359	12.249	39.938	77.104	30.446	26.664	10.954	28.624
Concrete Path	2151.29	117.78	142.59	**132.3**	24.647	15.739	14.152	75.673	49.03	8.031	15.984	51.321
Concrete Tiles	1893.96	431.72	451.08	326.65	**463.2**	12.765	37.48	80.883	28.658	26.067	8.98	28.054
Grass	2257.14	30.6	66.97	107.92	425.52	**40**	29.357	71.556	39.827	14.825	5.978	39.853
Large Slabs	2275.19	22.76	55.197	124.68	439.28	19.69	**14.8**	78.752	60.958	16.955	29.457	63.774
Linoleum	2288.54	50.457	70.615	140.65	433.64	56.97	44.92	**42.85**	106.328	67.715	77.326	102.148
Metal Plate	1816.36	461.91	483.02	345.1	135.38	451.53	468.57	475.67	**523.14**	52.872	35.378	9.881
Ramp	2270.86	3.737	36.07	120.64	433.41	32.62	22.66	47.69	464.49	**026.56**	17.71	53.847
Rough Brick	2104.74	208.58	224.19	129.41	231.68	209.08	219.04	209.24	294.22	209.67	**236.1**	36.18
Toilet Tiles	Inf	Inf	Inf	Inf	Inf	Inf	Inf	Inf	Inf	Inf	Inf	**inf**

the 64 values for feature 2 being identical, and hence, it has $\sigma = 0$. The distances for stationary measurements are much higher than those made during motion, as expected from the discussion in Section V.

The diagonal values in bold text in Table V are the distances between the moving and stationary sensor for the same object. These are quite large, reflecting the fact that the standard deviations of the measurements made by the moving sensor are much larger than those made by the stationary sensors, as shown in Fig. 9.

X. CONCLUSION

The experimental results confirm that the spatial-angle-filter model provides a physical basis for selecting features for measuring surface roughness with CTFM ultrasonic sensing. The highest quality features (the range of reflection angles and end angle) are a direct function of the roughness. At a lower quality, a group of four features are similar. Three are based on amplitude values and one is the distance between two peaks. Most of the features relating to the location and amplitude of two peaks are poor. In future research, we will study alternate feature detection algorithms to see if we can get values that relate better to the model.

The Mahalanobis distance (or Euclidean distance with normalized vectors) provides both a measure of the quality of the features and a robust method for classifying surfaces with those features. We have demonstrated better than 99.73% classification of 12 surfaces using five features measured with a moving sensor.

ACKNOWLEDGMENT

P. J. McKerrow would like to thank Dr. P. Probert Smith of the Department of Engineering Science, University of Oxford, Oxford, U.K., for a wonderful summer in Oxford, exploring and discussing the measurement of surface roughness. He would also like to thank the students of his "Perception and Planning" subject who implemented and tested the algorithms with the experimental data. The sensing system was purchased from Bay Advanced Technologies, Ltd., Auckland, New Zealand.

REFERENCES

[1] N. Harper and P. J. McKerrow, "Recognising plants with ultrasonic sensing for mobile robot navigation, Special Issue: Best Papers from EURobot'99," *J. Robot. Auton. Syst.*, vol. 34, no. 2–4, pp. 71–82, 2001.
[2] D. Ratner and P. J. McKerrow, "Navigating an outdoor robot along continuous landmarks with ultrasonic sensing," *J. Robot. Auton. Syst.*, vol. 45, no. 2, pp. 73–82, 2003.
[3] S. Schmidt, "Evidence for a spectral basis of texture perception in bat sonar," *Nature*, vol. 331, no. 6157, pp. 617–619, Feb. 18, 1988.
[4] P. J. McKerrow and N. Harper, "Plant acoustic density profile model of CTFM ultrasonic sensing," *IEEE Sensors J.*, vol. 1, no. 4, pp. 245–255, Dec. 2001.
[5] O. Bozma and R. Kuc, "A physical model based analysis of heterogeneous environments using sonar—ENDURA method," *IEEE Trans. Pattern Anal. Mach. Intell.*, vol. 16, no. 5, pp. 497–506, May 1994.
[6] Z. Politis and P. J. Probert Smith, "Modelling and classifying rough surfaces using CTFM sonar imaging," in *Proc. ICRS*, 1999, pp. 298–2993.
[7] ——, "Classification of textured surfaces for robot navigation using continuous transmission frequency modulated sonar signatures," *Int. J. Rob. Res.*, vol. 20, no. 2, pp. 107–128, Feb. 2001.
[8] P. J. Probert Smith and K. Zografos, "Sonar for recognizing the texture of pathways," *J. Robot. Auton. Syst.*, vol. 51, no. 1, pp. 17–28, 2005.

[9] O. Bergem, E. Pouliquen, G. Canepa, and N. G. Pace, "Time-evolution modeling of seafloor scatter II. Numerical and experimental evaluation," *J. Acoust. Soc. Amer.*, vol. 105, no. 6, pp. 3142–3150, Jun. 1999.
[10] H. J. Kim *et al.*, "Seabed classification from acoustic property data using the similarity index," *J. Acoust. Soc. Amer.*, vol. 111, no. 2, pp. 794–799, Feb. 2002.
[11] A. Lhemery, "Impulse response method to predict echo-responses from targets of complex geometry, Part 1: Theory," *J. Acoust. Soc. Amer.*, vol. 90, no. 5, pp. 2799–2807, Nov. 1991.
[12] J. A. Ogilvy, *Theory of Wave Scattering From Random Rough Surfaces*. Bristol, U.K.: IOP, 1991.
[13] Y. Y. Dorfman and I. Dyer, "Monostatic and bistatic reverberation statistics west of the Mid-Atlantic Ridge," *J. Acoust. Soc. Amer.*, vol. 106, no. 4, pp. 1755–1764, Oct. 1999.
[14] P. J. McKerrow and S. M. Zhu, "Modelling multiple reflection paths in ultrasonic sensing," in *Proc. IROS*, Osaka, Japan, Nov. 1996, pp. 284–291.
[15] L. Kay, "A sonar aid to enhance spatial perception of the blind: Engineering design and evaluation," *Radio Electron. Eng.*, vol. 44, no. 11, pp. 605–627, 1974.
[16] ——, "A CTFM acoustic spatial sensing technology: Its use by blind persons and robots," *Sens. Rev.*, vol. 19, no. 3, pp. 195–201, 1999.

Phillip J. McKerrow (M'73) received the B.E. (Hons.) degree in electrical engineering from University of New South Wales, Australia, and the M.E. degree in electrical engineering and the Ph.D. degree in computer science, both from University of Wollongong, Wollongong, Australia.

He is an Associate Professor in the School of Information Technology and Computer Science at the University of Wollongong. He is a Coordinator of the "Intelligent Systems Research Group." His research interests include machine perception with ultrasonic sensing, sensor-based control of indoor aerial robots, landmark navigation of mobile robots, and multimedia programming.

Bjorn E. Kristiansen, photograph and biography not available at the time of publication.