Object Recognition with FM Sonar; An Assistive Device for Blind and Visually-Impaired People*

Gordon Kao, Penny Probert, and David Lee
Robotics Research Group
Engineering Science Department
19 Parks Road, Oxford OX1 3PJ
Tel: +44-1865-273926, Fax: +44-1865-273908
Email: [kao,pjp,david]@robots.ox.ac.uk

Abstract

FM sonar sensors have been used in mobility aids for the visually-impaired. However, previous FM sonar systems have generated continuous audio signals and rely on the user interpreting them. Our research work is carried out to solve the problem of overloading users of FM sonar system with excessive information by machine interpreting the audio signal. The signal is sampled and Fourier transformed to generate an FM sonar image. Automatic computer analysis of the FM sonar image is carried out to compress and extract information for the purpose of object recognition. A method is developed to classify an object into one of the three groups: smooth surfaces, repetitive objects and textured surfaces. This method is based on the evaluation of the autocorrelation function of a single raw FM sonar image. A second method is also developed to reliably distinguish surfaces with varying degrees of roughness. An FM sonar model is constructed to predict FM sonar images of a rough surface at different sensor orientations. Templates are generated from the model and matched against the real images. Surfaces with varying degrees of roughness can therefore be identified.

1 Introduction

Ultrasonic rangefinders have been widely used for the guidance of mobile robots (Borenstein & Koren 1989; Bozma & Kuc 1994; Brown 1985; Elfes 1987). The most commonly used type of ultrasonic ('sonar') device is the time-of-flight sensor, which emits a brief pulse of ultrasound and measures the time until an echo is detected at the transducer. This delay is proportional to the distance travelled and can therefore be used to measure the range of the nearest reflecting object. Another type of sonar sensor, frequency-modulated (FM) sonar, has not been used extensively in robotics but has found some acceptance in mobility aids for the visually-impaired, such as the Sonic Torch and the Sonic-guide. An FM sonar sensor emits a succession of 'chirps' in which the frequency of the ultrasound changes between lower and upper limits. The distance to a reflecting object is found by blending the echo with the transmitted signal and analysing the resulting 'beat' frequencies. This technique will be described in more detail later in this paper.

A potential advantage of FM sonar over simple time-of-flight sonar is that the 'beat' signal contains information about all of the reflecting objects, not just the nearest one. In fact, the Sonic Torch and the Sonic-guide simply present this signal to the user through an earpiece or headphones. Some trained users can then determine not just the distance to a number of objects but also their texture. One difficulty with this direct presentation of the beat signal is that it monopolises the user's sense of hearing, a vital source of information for a blind person. It also requires a lengthy training period before the signal can be well understood.

In this paper we describe an investigation into the automatic analysis of the output from an FM sonar sensor. It is hoped that a blind user could be given the results of this analysis, possibly through synthetic speech. This will decrease the amount of audible information given to the user and will eliminate the need for training.

Section 2 explains the principle of FM sonar in more detail and describes the sensor which was used in this investigation. Section 3 investigates the properties of the 'beat signal'. The beat signal can be transformed to the frequency domain, allowing the reader to see the type of information which is present in the signal. Throughout this paper, we will refer this frequency spectrum of the beat signal as the *FM Sonar Image*. Two methods are developed in order to automatically analyse the FM sonar images and to tackle the problem of object recognition. Section 4 describes the method based on analysing a single FM sonar image. Section 5 describes the alternative method based on analysing multiple FM sonar images obtained at the same distance but with different sensor orientations.

The first method is fast and effective in recognising specific spatial geometric configurations, such as walls and stairs. The second method can be used to obtain robust surface texture information, such as distinguishing glass windows from rough brick walls. The two methods can be combined together to achieve satisfactory performance, depending on the user's requirements. Section 6 summarises the research and presents conclusions and finally closes the paper by considering suitable topics for further research.

2 The Sensor

The FM rangefinder uses two ultrasonic transducers: one transmitter and one receiver. In contrast to simple time-of-flight sensors, which transmit a single short pulse of ultrasound, the FM sensor transmits a repeating pattern of varying frequency as shown in Figure 1.

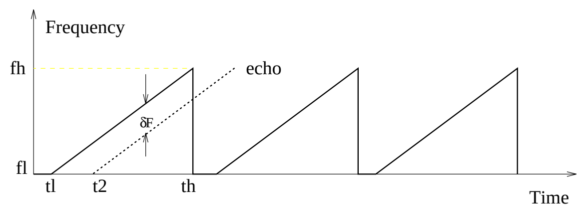

Figure 1: Frequency pattern of the FM sonar
The dotted line shows the frequency of the received echo if an object is placed in front of the sonar.

Imagine a single object which reflects this signal back to the receiving transducer. The received signal (echo) will then have the same sawtooth frequency pattern but it will be 'out of step' with the transmitted signal. After a delay during which the sound travels to and from the object, there is a period of overlap during which the frequency-ramped signal is being transmitted and received simultaneously. The difference between the frequencies of the two signals, δF, is constant during the overlap period. The value of δF is proportional to the delay between transmitting and receiving the signal, a delay which is in turn proportional to the distance to the reflecting object. The range can therefore be deduced from the value of δF.

To find δF, a new signal is created which is the product of the transmitted and received signals. This product can be shown to have two additive components: one has a frequency equal to the sum of the original two frequencies; the other has a frequency equal to the difference between the two frequencies, δF. The value of δF can therefore be found by applying a Fourier transform to the product signal. (In practice the other, high frequency, component can not be detected at the sampling rates used in this sensor.)

In experimental environments the sensor received echoes from a number of environmental features, giving rise to multiple peak frequencies in the Fourier transform. Therefore, as will be demonstrated later, the sensor is able to detect several objects simultaneously.

The sensor that was used in this investigation has the characteristics as shown in Table 1. The sensor is connected through an earphone socket to a laptop PC. The PC has a PCMCIA-compatible card which includes an 8-channel analogue-to-digital converter, 3 digital inputs and 3 digital outputs. The maximum total rate at which the PC can drive the ADC is 25 kHz. One of the channels has been used to sample the 'beat signal'. The third digital input channel has been used as an external trigger for the sampling. (The first stage in the frequency generation process in the sensor is to create a square wave. This square wave has been connected to the digital input and acts as a trigger.)

To make it easier to collect experimental data, two buttons have been mounted on the sensor to provide input to the remaining two digital input channels on the PC. The state of these buttons can be read by the PC, enabling the system to be controlled without using the keyboard.

Name	Meaning	Value
fl	Lowest transmitted frequency	45 kHz
fh	Highest transmitted frequency	90 kHz
tl	Time for which signal stays at fl	24 ms
tr	Time to ramp from fl to fh	160 ms

Table 1: Sensor Characteristics

3 The Properties of the 'Beat Signal'

In the original guidance system, the FM sonar output, i.e the beat signal, was directly presented to the user through an earpiece or headphone. Users were required to interpret information contained within the beat signal themselves. In our work, the beat signal is sampled and transformed for the purpose of automatic analysis. Before describing the automatic analysis, we first examine the properties of this beat signal.

3.1 The Sampling Effect

The discrete Fourier transform determines the frequencies present in the beat signal by taking a set of samples. The sampling process itself can lead the Fourier transform to detect component frequencies which were not present in the original signal. The following two examples will make this clear.

Figure 2(a) shows a computer-generated example in which the wavelength of the signal was chosen to be an exact fraction of the total number of sample points. The sample therefore includes a whole number of cycles of the signal and the frequency distribution shows a single sharp peak. In contrast, Figure 2(b) shows the results when the wavelength is not an exact fraction of the number of sample points. In this case the sample ends part-way through one of the cycles of the signal. The Fourier transform then generates additional frequency components to describe the discontinuity. The peak in the frequency distribution is then less sharp.

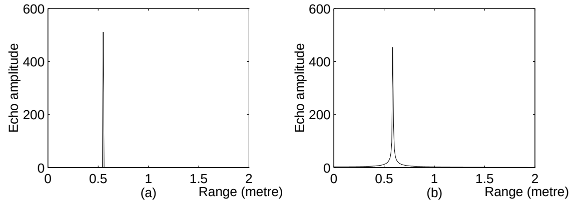

Figure 2: Illustration of the sampling effect
The beat signal is sampled and transformed to the frequency domain. The horizontal axis is the frequency δF which is proportional to the range. The vertical axis shows the strength of the beat signal which is called 'Echo Amplitude'. We call this frequency representation an FM sonar image.

This characteristic shape will be seen in the experimental results presented in the remainder of this paper.

3.2 Compensation for Loss of Signal Power

The further the sound has to travel before it reaches the receiving transducer, the weaker the echo will be. Two factors contribute to the loss of power:

- **Dispersion of the signal.** The transmitted signal forms part of a spherical wavefront. The signal power therefore decreases with the square of the distance from the transducer. If the signal is reflected from a flat specular surface the signal power continues to decrease as the square of the total round-trip distance. If, on the other hand, the signal is reflected as a diffuse echo from, for example, a convex corner, then the signal power decreases as the fourth power of the distance to the object.
- **Attenuation of the signal.** As the sound travels through the air, it loses power because of factors such as the viscosity of the air. This attenuation takes the form of exponential decay over distance. The rate of decay is higher for higher-frequency sound.

Some time-of-flight ultrasonic rangefinders include a variable-gain amplifier in the receiving electronics to compensate for this loss of signal power; the longer the echo takes to arrive, the higher the gain of the amplifier. Such an approach is reasonable for time-of-flight sensors because the transmitted signal is short and one is usually only interested in the first echo. However, the situation is different with the FM sensor. At any moment the received signal could be a mixture of multiple echoes, each with its own signal power and frequency. A superficially-tempting approach is to apply a frequency-dependent amplifier to the results of the FFT of the beat signal. The lower the frequency in the beat signal, the nearer the reflecting object and the less signal power will have been lost. Unfortunately this situation is complicated by the variation of attenuation rate with signal frequency. A single frequency in the beat signal will be caused by echoes at a range of frequencies, each with its own attenuation rate.

Given the difficulty of finding an effective amplification function, it was decided not to implement a time- or frequency-dependent amplifier. Instead the display software on the PC and the graphs in this paper simply scale the display of the beat signal and the frequency spectrum so that the largest measured value achieves a full-scale deflection.

3.3 Crosstalk and Minimum Range

The crosstalk problem is significant when the FM sonar is used to measure a distant object, as illustrated in Figure 3. Figure 3 shows the frequency spectrum for a distant object, a bookcase at a range of about 3.5 metres. Two distinct peaks can be seen, corresponding to the front of the bookcase and the wall behind it. But even more noticeable is the number of low-frequency components, which would be interpreted as range readings of less than 300 mm. However, there were no obstructions close to the sensor in practice. These low-frequency values are caused by direct transmission of the signal from the transmitter to the receiver within the sensor itself.

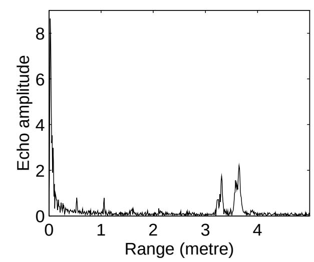

Figure 3: A bookcase at 3.5 metres

As the echo from the detected obstacle becomes weaker, the low-frequency components begin to dominate the frequency spectrum and restrict the auto-scaling process, making it harder to detect distant objects. To prevent this, it was decided to discount components of the frequency spectrum which correspond to a specified minimum range, which was taken to be 300 mm. The auto-scaling process can then magnify the weaker signals from more distant objects.

4 Object Recognition with a Single FM Sonar Image

The capability of FM sonar has already been demonstrated by users of guidance systems. Users of the Sonic Guide (which includes two sensors, one placed on each side of a pair of spectacles) report extraordinary ability to discriminate between environmental features by listening to the stereo audio output from the FM sonar sensor (Kay 1985): A user one of the authors met recently could detect the difference between conifers and deciduous trees. Recently, Harper

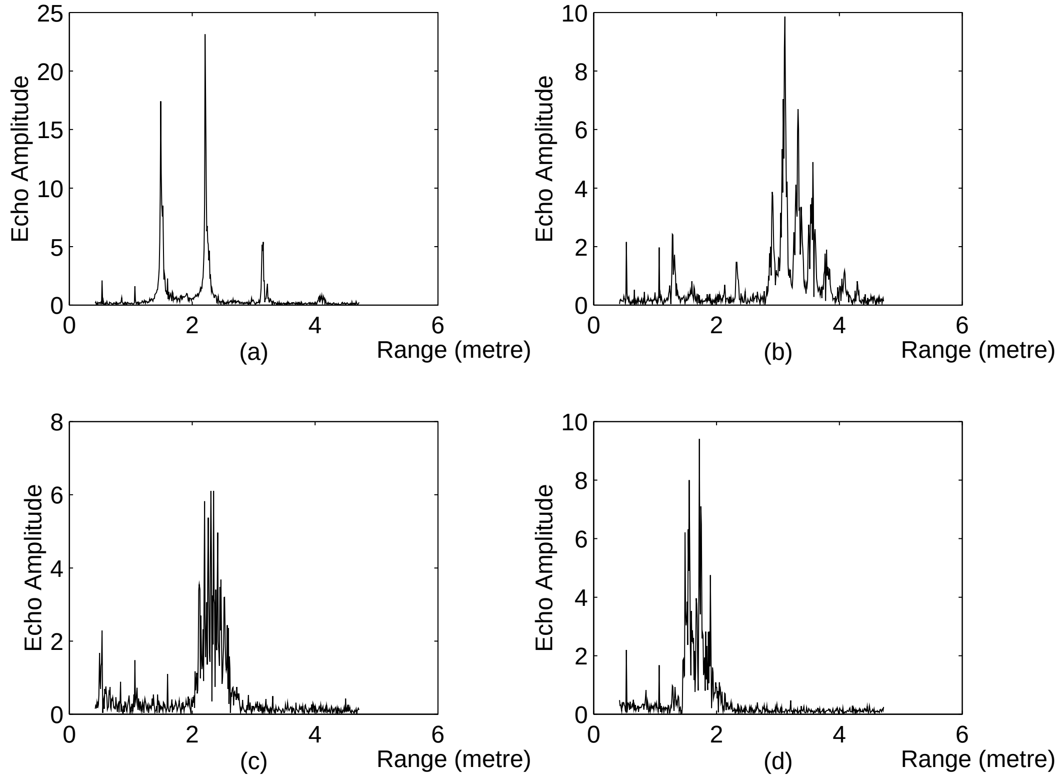

Figure 4: FM sonar images of common objects
Plot (a) shows the image of multiple targets including a mesh, a bar fence and a brick wall. Plot (b) shows the image of stone steps. Plot (c) shows the image of a rubbish bag. Plot (d) shows the image of an overhanging tree branch.

and McKerrow reported their experiment in classifying plant species with the aid of an FM sonar using a neural network (Harper & McKerrow 1995). However, despite its sensing capability, there are problems associated with using Sonic-guide. The problems are twofold:

- **Training Requirement** The user of Sonic-guide is normally required to be trained for a lengthy period before he(she) can well understand the audio signal.
- **Information Overloading** The continuous audio signal output can overload the user and sometimes become irritating.

In order to solve the problems mentioned above, machine interpretation of the FM sonar output is necessary. By automatically analysing the FM sonar output, we aim to compress the FM sonar output and extract some of the most important information such as the object types and texture from the raw FM sonar output. Fortunately, recent fast increases in computational power make it feasible to interpret FM sonar outputs and process them in real time.

Two methods are developed to automatically analyse the FM sonar image for the purpose of object recognition. The first method is described in this section and it is based on analysing a single FM sonar image. This method is fast and effective in recognising spatial correlation in a group objects (e.g walls and steps). However, this method is not robust when it is used to examine rough surfaces, because the FM sonar image also depends on the sensor orientation in this case. To solve the problem, a second method is developed and it is described in Section 5. This method is based on analysing multiple FM sonar images obtained by sweeping the sonar across the surface being examined. The second method is much more complex and time-consuming, but it gives reliable surface texture discrimination (e.g distinguishing a glass window from a rough brick wall). The two methods may be combined together to give satisfactory performance.

4.1 FM Sonar Images of Objects in an Urban Environment

Before we proceed with our machine interpretation of a FM sonar image, we examine some typical FM sonar images of a variety of objects which are likely to be encountered on town streets (Figure 4). It is shown in Figure 4 that the sensor is able to detect multiple reflecting surfaces simultaneously, and the image either shows a number of objects or gives an indication of the texture or shape of the object. The sensor can 'see through' certain objects in Figure 4(a), where the sensor was pointing at 2 cm square wire mesh at 1.5 metres, behind which was a metal bar fence at 2.2 metres and a brick wall of building at 3.2 metres. All three objects were clearly detected. On the other hand, Figure 4(b) shows the characteristic periodicity associated with steps. Finally in Figure 4(c) and (d) we show typical images resulted from clusters of objects. It is clear that the echo energy is spread over a large group of ranges. Figure 4 shows that each object generates its own specific pattern in the FM sonar image.

4.2 An Automatic Analysis System

An FM sonar image can be mostly characterised by its first two statistical moments, the mean and variance. The mean shows the average echo energy, i.e the average intensity of the FM sonar image. The variance described by the autocorrelation is more important, since it describes the spatial correlation of the FM sonar image. Based on the evaluation of the autocorrelation function (ACF) of the image, we have developed an automatic analysis system, which is able to categorise an object into one of the following three groups:

- Smooth surfaces
- Repetitive objects (stairs or step-like objects)
- Textured surfaces or a group of objects (clusters)

The system is illustrated in Figure 5.

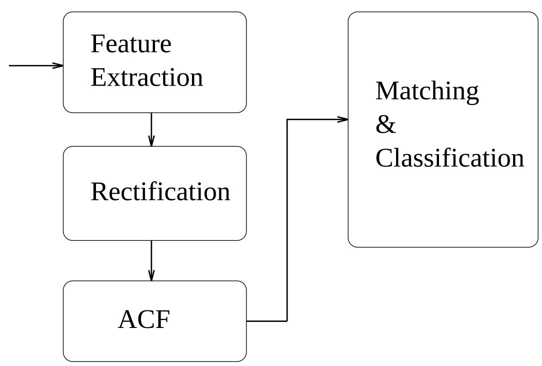

Figure 5: The automatic analysis system

Features are first extracted from the image by thresholding the whole image, and picking out those area(s) with amplitudes above an adaptive threshold. The adaptive threshold is chosen as a percentage of the mean intensity of the FM sonar image. The early returns from the sonar (corresponding apparently to short ranges) resulted from the crosstalk are not included in the processing. A minimum range for valid results is set to remove these components.

Each extracted feature is then enhanced through rectification. The mean amplitude of each feature is computed. Amplitudes exceeding the mean value are assigned a value of 1 and amplitudes below the mean value are assigned a value of 0. The resulting image is binary. Such a process can be regarded as a crude low-pass filtering process which not only provides a significant reduction in the effects of noise but also speeds up the evaluation of the ACF.

The ACF is evaluated by:

$$R_j = \frac{\Sigma_{i=0}^{N-j} \Im_i \Im_{i+j}}{\Sigma_{i=0}^{N} \Im_i \Im_i} \tag{1}$$

where $\Im_i$ is the echo amplitude at range i in the rectified FM sonar image, j is the spatial correlation distance, and N is the maximum range of the image. Performing the ACF evaluation on FM sonar images of wall, steps, rubbish bag and overhanging branches in Figure 4, we can obtain their corresponding ACFs as plotted in Figure 6.

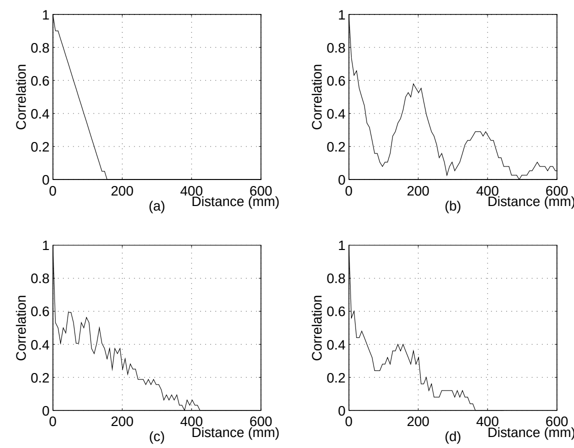

Figure 6: ACFs of various objects
Plot (a) shows the ACF of a brick wall. Plot (b) shows the ACF of stone step. Plot (c) shows the ACF of a rubbish bag. Plot (d) shows the ACF of overhanging tree branch.

It is shown that periodicity can be detected and that stairs can be distinguished as shown in Figure 6(b). Although at first sight the image for the branches looks as if there may be periodicity here (Figure 4(d)), it is clear from its autocorrelation function that such periodicity if any, is very weak (Figure 6(d)). The rate of decay of the ACF describes the characteristics of the surface, with smooth surfaces showing a fast rate of decay (Figure 6(a)), while the rate of decay of the ACF for other clusters is much slower (Figure 6(c) and (d)).

The result from the ACF evaluation is used for the purpose of matching and classification. The periodicity of the ACF is examined first. The first local minimum and local maximum are extracted from the ACF. A stair or step-like object is recognised if the minimum falls below its threshold and the maximum exceeds its threshold. Both threshold values are determined from experiments. If the periodicity of the ACF is not detected in the ACF, the decay rate of the ACF is then examined. The distance where the correlation falls below 10% is estimated. A smooth surface is recognised if the distance is small, otherwise a texture surface (or clusters) is recognised. The result after matching and classification is finally communicated to the user through a speech generator.

Experiments were repeatedly carried out to examine different types of object with different sensor position and orientation. It is found that the system works successfully and it is able to interpret and categorise most of the objects in a real world consistently. However errors were encountered while the FM sonar was used to examine rough surfaces. The system was not able to categorise a rough surface consistently. The system classifies a rough surface either as a smoother surface or a textured surface depending on the specific orientation of the sensor. This problem is further described and tackled in the next section.

5 Surface Recognition Using Multiple FM Sonar Images

In the previous section, we developed an object recognition system. The system is based on analysing a single FM sonar image, and is effective in categorising an object into one of the three groups. However the system is not able to discriminate different surface textures, since the FM sonar image of a rough surface changes with different sensor orientations. The problem is illustrated in Figures 7 and 8.

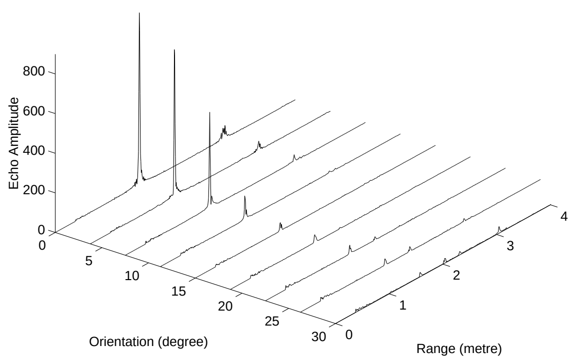

Figure 7: FM sonar images of a smooth wall
All images obtained at a normal distance of 1.55 metres to the wall, but with different orientation to the wall. A zero degree orientation means the FM sonar is directed normal to the wall.

In Figure 7, images of a smooth surface were obtained at the same distance but with different sensor orientations. It can be seen that the shape of the FM sonar image does not depend on the sensor orientation. The images shows a single peak at the range equal to the normal distance to the wall until it is lost in noise.

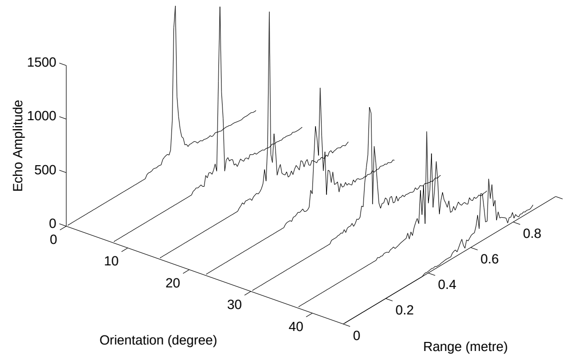

Figure 8: FM sonar images of a rough wall
All images obtained at a normal distance of 0.5 metres to the wall, but with different orientation to the wall. A zero degree orientation means the FM sonar is directed normal to the wall

The same experiment was carried out for a rough surface and the images are plotted in Figure 8. It can be seen that the basic shape of a rough surfaces image varies significantly with varying sensor orientation. As the orientation is large, echo energy is spread much wider. On the other hand, as the orientation is close to zero, echo energy is more concentrated on the specular direction, and the image has the same shape as that of the smooth surface. Because the image of a rough surface changes significantly with different sensor orientations, surface texture recognition based on a single FM sonar image is not reliable. Note that a rough surface is visible to the FM sonar for a much larger angular interval than a smooth surface.

The problem is tackled by analysing multiple images simultaneously. By performing scattering analysis and modelling a rough surface as a Gaussian distributed process, we can construct a model which is able to predict FM sonar images for surfaces with varying degrees of roughness. By matching the templates generated from the model with FM sonar images, we can discriminate different surface textures, and estimate the texture parameter in terms of the mean height deviation and the spatial correlation constant of the surface.

5.1 Model Construction

When the FM sonar is used to examine the texture of a rough surface, the transmitted wave of the FM sonar can be considered as a collection of rays as shown in Figure 9. Each ray is directed towards the surface and backscattered from it. The backscattered ray is received by the FM sonar and results in an echo amplitude value at the range $d = d_0 / \cos^{-1}\theta$ in the FM sonar image, where d_0 is the normal distance to the surface and θ is the incident angle of an individual ray. Therefore, we can describe our FM sonar image $\Im$ of a

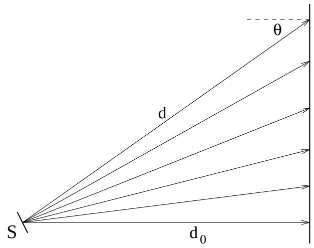

Figure 9: Examining a surface using the FM sonar

surface as an ensemble of echo components, where each echo component is associated with the ray of incident angle θ, and the amplitude of each is described by:

$$\Im_\theta = \wp \rho_{bs}$$

where ρ_{bs} is the scattering coefficient of the surface in the backward direction, and $\wp$ is the sonar beam strength defined from the sonar directivity. Once the scattering property of a rough surface and the directivity of the sonar are known, the FM sonar model can be constructed and the FM sonar image can be predicted.

The theoretical derivation of this model is cumbersome and it is not presented here. The reader is referred to (Kao 1996) for more details. Here, we only describe one of most important procedures during model construction: the statistical surface modelling.

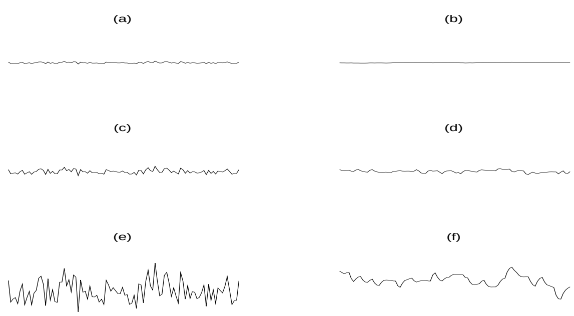

Figure 10: Simulated surfaces using a Gaussian model. The roughness of the surface is described by the height deviation σ. Plot (a), (c) and (e) show surfaces with increasing σ. The surface irregularity is described by the spatial correlation constant T. The heigh deviation of surfaces in plot (b), (d) and (f) are same as those in plot (a), (c) and (e) respectively, but their spatial correlation constant T are increased.

Our rough surface is modelled as a Gaussian distribution (Bozma & Kuc 1994; Beckmann & Spizzichino 1963). In this model, a Gaussian distribution of the surface height is assumed with zero mean and standard deviation σ. σ describes the surface roughness. If σ is small comparing with the wavelength, the surface is smooth. If σ is large, then the surface is rough. The distribution of the valleys and hills on the surface is modelled by a spatial correlation constant T. T is small if the surface is irregular and T is large if the surface is regular.

5.2 Template Matching

The sonar model is expressed in terms of the ratio of the surface correlation constant to the mean surface height deviation: T/σ (Kao 1996). The ratio T/σ describes the surface texture. T/σ is small if the surface is rough. T/σ becomes larger if the surface is smoother.

A template can be generated from the model to predict FM sonar images of a rough surface at different sensor orientations. Several templates can be generated from the model using different values of T/σ. Let us use $\Im^{M_j}$ to denote each of the templates. The template is matched against the true FM sonar images denoted by $\Im^T$ obtained from the experiment. We use the correlation coefficients to examine the matching (mismatching). The correlation coefficient for each template is given by

$$\nu^{M_j} = \frac{2\sum_\beta \sum_{i=1}^N \Im_{i,\beta}^{M_j} \Im_{i,\beta}^T}{\sum_\beta \sum_{i=1}^N \Im_{i,\beta}^{M_j} \Im_{i,\beta}^{M_j} + \sum_\beta \sum_{i=1}^N \Im_{i,\beta}^T \Im_{i,\beta}^T}$$

where element-to-element multiplication is performed between a true image and a predicted image from the template. The product is summed for the whole image size N. The summation is repeated for all images at different sensor orientations β. The denominator provides a normalising factor to generate a correlation coefficient ν^{M_j} between 0 and 1. A perfect matching is obtained if $\nu^{M_j} = 1$, a mismatching is obtained if ν^{M_j} is close to zero.

T/σ	ν
1.5	0.75
2.0	0.82
2.5	0.84
3.0	0.82
3.5	0.78
4.0	0.73

Table 2: Correlation coefficients

The correlation coefficient is evaluated for each template with different value of T/σ. The template which gives the maximum correlation coefficient is declared as the matched model. Table 2 shows the correlation coefficients for the rough surface plotted in Figure 8. The correlation coefficient is found to be a maximum of 84% where the ratio T/σ is 2.5. Therefore the template with T/σ is declared as the matched model. It can be seen from Figure 11 that the prediction from the sonar model fits the experimental results very well. As predicted, the maximum amplitude drops and the acoustic energy spread wider as the sensor orientation increases.

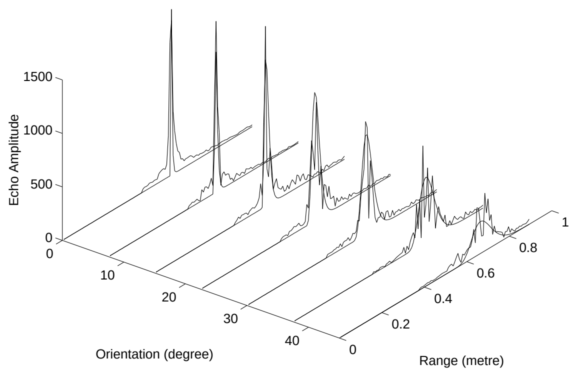

Figure 11: Matched FM sonar images of a rough wall The matched prediction obtained from the sonar model with $T/\sigma = 2.5$ is superimposed on the true images from Figure 8.

6 Conclusion

In this paper, we have described methods for automatic analysis and interpretation of FM sonar output to assist a visually impaired person. The basic method based on analysing a single image is very efficient and effective in identifying smooth surfaces and stairs, and this method can be further compensated by analysing multiple images to discriminate and identify different surface texture. The results in classification suggest that the FM sonar is a promising device for low cost, fast discrimination between a number of features useful for navigation in a street scene.

One of the ongoing themes in the future research is how far information should be abstracted before it is presented to the user. People are far better than computers at recognising patterns, but the presentation of large amount of raw data can overload the sensing systems available to the blind. The early work on the FM sonar is encouraging but it is too early still to know whether a sufficient range of features can be included, either with the sonar alone or in conjunction with another sensor such as vision.

Acknowledgements

The work was funded by the TIDE programme of the European Union. One of the authors is supported on an EPSRC grant. The original FM sonar unit was designed and built by David Witt, following an earlier design by Leslie Kay, and we are grateful for the excellent construction and ideas.

References

Beckmann, P., and Spizzichino, A. 1963. *Electromagnetic Waves.* MacMillan, New York.

Borenstein, J., and Koren, Y. 1989. Real-time obstacle avoidance for fast mobile robots. *IEEE Trans. Systems Man and Cybernetics* (19).

Bozma, O., and Kuc, R. 1994. A physical model-based analysis of heterogeneous environments using sonar - ENDURA method. *IEEE Transactions on pattern analysis and machine intelligence* 16(5).

Brown, M. 1985. Feature extraction techniques for recognising solid objects with an ultrasonic range sensor. *IEEE journal of Robotics and Automation* (4).

Elfes, A. 1987. Sonar-based real world mapping and navigation. *IEEE Trans. on Robotics and Automation* (3).

Harper, N. L., and McKerrow, P. J. 1995. Classification of plant species from CTFM ultrasonic range data using a neural network. In *IEEE International Conference on Neural Networks.*

Kao, G. 1996. FM sonar modelling for navigation. Technical report, Department of Engineering Science, University of Oxford.

Kay, L. 1985. Sensory aids to spatial perception for blind persons - their design and evaluation. In *Electronic spatial sensing for the blind.* Martinus Nijhoff Publishers.

Thorsos, E. I. 1988. The validity of the Kirchoff approximation for rough surface scattering using a Gaussian roughness spectrum. *J. Acoust. Soc. Am.*