Classifying plants with ultrasonic sensing

Neil Lindsay Harper and **Phillip John M^cKerrow**
School of Information Technology and Computer Science
University of Wollongong
Wollongong
NSW, 2522
AUSTRALIA

Email: phillip@uow.edu.au
Phone: +61 (02) 4221 3771
Fax: +61 (02) 4221 4170

Abstract

When a surface is insonified, the resultant echo contains information about the geometric structure of the surface. Plants have complex geometry and occur in many potential applications of robotics. 100 plants of different species were sensed with a CTFM ultrasonic sensor. These plants were classified into groups using both the acoustic density profile produced from the echo and features extracted from the acoustic density profile

1. Introduction

The recognition and classification of plants and fruit will enable agricultural robotics to address many new applications. Vision research has succeeded in classifying plants only in highly constrained situations where the environment is controlled. Examples of environment control include placing all of the samples at a known distance and against plain backgrounds, very specific lighting conditions, and even placing leaves to be recognised in certain positions within the frame.

Some of the problems encountered with vision systems can be solved with ultrasonic sensing. The rest of this paper discusses the classification of plants using the geometric information in ultrasonic echoes. In previous research, we demonstrated that there is sufficient information in the echo for a neural network to recognise 1 of 4 plants over a wide range of orientations [Harper and M^cKerrow, 1997].

To investigate the problems of classifying plants, we scaled up to measuring 100 plants. The plants were young trees supplied by the Wollongong Botanic Gardens in pots. They were chosen in accordance with plant taxonomy to give a variety of leaf shapes and foliage structures. A database was built for each plant containing echoes from 360° of rotation in 1° steps.

Two problems were immediately obvious: plant similarity, and asymmetry with rotation. The echo received by the sensor varies depending on the angle from which the plant is insonified. However, there is information in the signal which is relatively invariant with orientation and which characterises the structure of the plant. We can capture this information by extracting 19 features from the echo.

With these features plants can be classified into groups and, often, individual plants can be recognised in a group. The results of these experiments confirm that there is sufficient information in an ultrasonic echo to classify plants into groups.

2. Machine perception of plants

Some agricultural applications of live plant classification have achieved limited success. Plant sensing is difficult because plants change considerably with orientation, and their shape changes as they grow. Vision has been applied under significant constraints. It has problems with consistency of illumination. Prior to the research project reported in this paper, plant classification has not been attempted with ultrasonics due to the lack of an appropriate sensing device.

Guyer *et al* [1993] attempted to identify plants based on the shape of their leaves. Their research was limited to images of leaves in isolation against a uniform background of soil. They found that the errors in classification were due to poor images, poor segmentation, and natural scene variation. They reported that the major problem for vision systems in agricultural applications is biological variability within plant species.

Classification of plant cuttings was attempted by Singh & Montemerlo [1997]. They graded three different Cultivars of *geraniums* as either small, medium, or large to high accuracy (90%) using computer vision. They found that the appearance of the cutting is changed significantly by the change in viewing angle of the sample but good feature calculation can reduce the effects of this.

Nabout *et al* [1994] identified plants so that they could separate weeds from crop in order to apply herbicides. They found that there are many different kinds of plants with complex forms which cannot be described using simple geometric models. They claim that they could recognise 17 different weed species to 82% accuracy. Rigney & Kranzler, [1988] claim that they can classify seedlings into one of two classes "acceptable" or "cull" with a low classification error. They use a vision system to capture a silhouette of the seedling in order to measure several physical properties of the seedlings.

3. CTFM

When a plant is insonified the resultant echo contains information about the geometry of the plant, in particular the structure of the foliage. The echo data was produced with a Continuously Transmitted Frequency Modulated (CTFM) ultrasonic system developed by Leslie Kay [Kay, 1974]. A CTFM system transmits a sine wave signal that is repeatedly frequency swept over a one-octave

range (typically 100 to 50 kHz with a sweep period of 102.4 ms). The echo is a filtered version of the transmitted signal, offset in time.

The echo is demodulated with the transmitted signal to obtain a set of audio tones (0..5KHz) proportional to range (Figure 1). The audio tone for a target is continuous from the time at which the echo arrives until the end of the sweep. Thus, the sweep consists of two time periods: the first is the time to the arrival of an echo from maximum range and the second is the time to capture the samples for the Fast Fourier Transform (FFT).

In the time domain, the complexity of audio signal is proportional to the geometric complexity of the target. In the frequency domain, a spectral line occurs for each range where energy is reflected. A 1024 point FFT produces positive amplitude values for 512 * 9.77 Hz frequency bands. The amplitude of the spectral line is proportional to the intensity of the echo and hence to the area of the reflecting surface normal to the receiver.

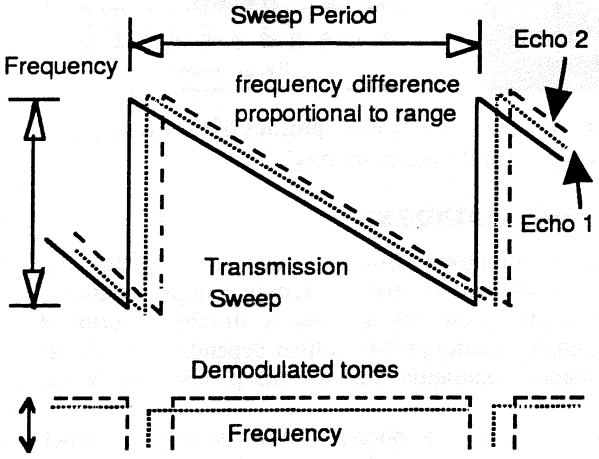

Figure 1 CTFM Signals

4. Acoustic density profile

The audio tones are range measurements to every near normal surface in the region of insonification. A small amount of the echo results from multiple reflections and interference, and some of the echo is lost due to shadowing. The frequency spectrum includes information about the range, size and orientation of surfaces within the field of audition and can be modelled as an acoustic density profile. In this model, the sensor measures the acoustic area at each range. The acoustic density profile for a specimen of *Leptospermum laevigatum* is shown in Figure 2. The sum of the values in the range cells is a measure of the acoustic area of the plant.

The acoustic density profile is a measure of the acoustic density in each range cell. Figure 2 shows a significant acoustic density between the ranges of 200 mm and 540 mm and this corresponds to the position of the plant in space. The amplitudes of the individual range lines are a function of the properties of the surfaces at that range (Figure 3):

the area of the leaves in the range cell;
the orientation of the leaves;
the texture of the leaves; and
the amount of occlusion that affects the leaves.

In Figure 3, the boundary of the plant is represented with a circle and leaves are represented as filled ellipses. Waves reflect from the surfaces of the leaves and are detected by the receiver R_x. Each range cell is 3.4 mm deep. We developed the acoustic density profile model in order to interpret the information in the echo reflected from the plant [Harper, 1998]. We can interpret the frequency spectra of the received echo, and match the features used for classification to geometric properties of the plant, using this acoustic density profile model.

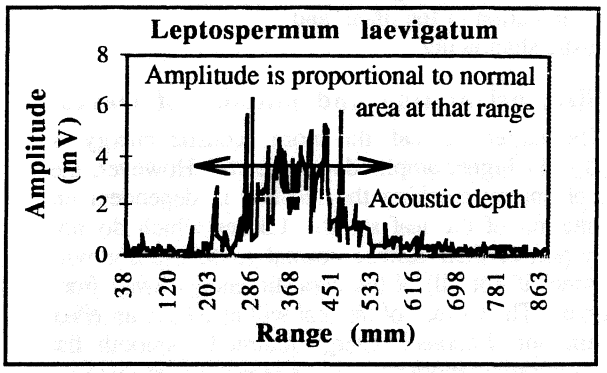

Figure 2 The acoustic density profile model of the spectrum of the demodulated CTFM echo.

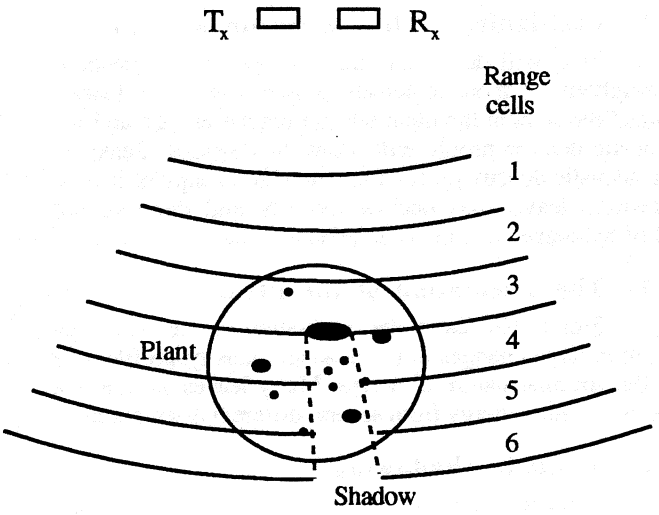

Figure 3 Plant insonification showing distribution of reflectors through range cells.

5. Data characteristics

Specular surfaces have a narrow-band signal which often consist of a very large magnitude at a single range. Diffuse scatterers result in much smaller magnitudes. Objects with multiple surfaces generate echoes from each surface and the resulting acoustic density profile contains information about all of the surfaces. It includes the absolute range to the object and the geometric structure of the object. The absolute and relative range information can be decoupled.

Each spectral line has two parameters: frequency and amplitude. The frequency represents the absolute range to the surface and the amplitude is a function of the size, specularity and orientation of the surface Since there are many reflective surfaces in a plant, the return signal is

complex (Figure 2). Also, the maximum amplitude (6 mV) is significantly smaller than that for a float wall (100 mV) or a 10 mm metal rod (20 mV).

A classification system can be developed and that is almost independent of the distance to the plant, the height of the plant, and the width of the plant, provided that enough of the plant is in the field of audition. The following properties of plants affect the echo:

Size, orientation and number of leaves,
The spatial positioning of leaves within the plant,
The orientation of the plant, and
Acoustic shadowing.

5.1 Size, orientation and number of leaves

The larger the leaf, the more acoustic energy is reflected and a higher amplitude is received. However, the amount of energy reaching the receiver is dependent on the orientation of the leaf surface. Leaves which do not have any part of their surface normal to the sensor will reflect some (if not all) of the acoustic energy away from the receiver. The surface of the leaf will also have an effect on the amount of acoustic energy returned - smooth flat surfaces will reflect more back to the receiver than textured surfaces. In general, the more leaves that are on a plant, the larger the percentage of acoustic energy that will be reflected.

5.2 Positioning of leaves within a plant

This will determine the distribution of reflections throughout the acoustic density profile (Figure 3). Leaves spread throughout the plant will reflect the energy and the acoustic density profile will show this spread. Peaks in the acoustic density profile may indicate groups of leaves. However, leaves may occlude those behind them so not all of the leaves will be detected at any one time.

5.3 The orientation of the Plant

Small changes in the orientation of the plant can result in large changes in the acoustic density profile, due to the specular nature of leaves. Many leaves are not flat and will return energy from several different orientations.

5.4 Acoustic shadowing

Acoustic shadowing occurs when the ultrasonic wave either does not reach a reflector or is reflected but does not reach the receiver due to another reflector being in the way. Figure 3 shows some acoustic shadowing on the bottom right hand corner. The large leaf prevents the bulk of the acoustic energy from penetrating into the range cells behind it. Some refraction may occur in this situation. Refraction allows a small amount of the acoustic energy to penetrate but this will result in a very small contribution to the amplitude in the range cell.

Dense plants have more shadowing and less of the acoustic density profile is a measure of the leaves beyond the front surface of the plant. In general, the acoustic density profile of dense plants has one region with large amplitude which corresponds to the front of the plant (e.g. *Polyscias murrayi* in Figure 4b). Very little acoustic energy penetrates the front layer of foliage, and leaves at the back are shadowed so they are not present in the acoustic density profile.

In contrast, the acoustic density profile of *Eucalyptus maculata* includes echoes from the leaves at the back of the plant. The acoustic density profiles of sparse plants have lower amplitudes but reveal more about the structure of the plant. Note that the background in the images is not present when the acoustic density profile is measured.

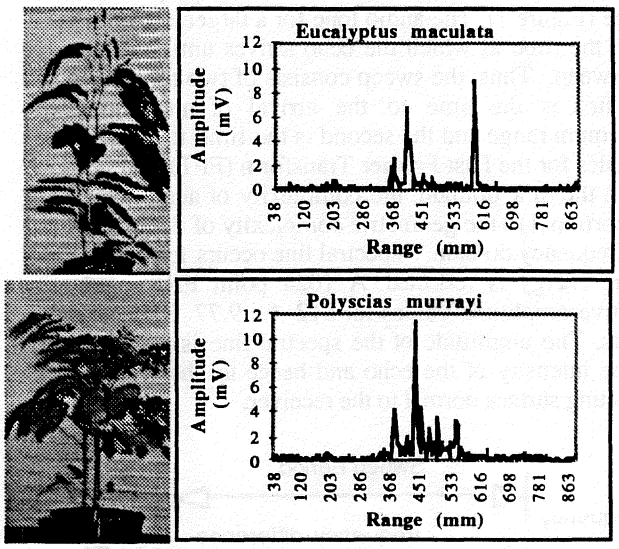

Figure 4 Acoustic density profiles of (a) Eucalyptus maculata and (b) Polyscias murrayi.

6. Plant database

The Wollongong Botanic Gardens supplied 100 plants for these experiments, in family groups. Plants of the same family are not necessarily similar in terms of their acoustic density profile, which depends more on the size, shape, orientation and overall positioning of the leaves.

For each plant, a portfolio was established which contains data about the plant, including a photograph. The portfolio contains the information in Table 1, the conditions under which the plant was processed ie. date, time, temperature and humidity; and a sample acoustic density profile for the plant. In addition, a database was built for each plant containing echoes from 360° of rotation in 1° steps

Table 1 Characteristics of plant species [Roger & Carolin, 1990].

The scientific name of the plant species.
The assigned common name.
The family of plant to which this species belongs.
The total height of the specimen from the soil to the top
The total with of the plant at its widest point
The average leaf length from the base to the tip.
The average width of the leaves at their widest point.
Count of the total number of leaves on this specimen
Estimate of plant density - scale of high, medium or low
Leaf shape of plant [Roger & Carolin, 1990].
Arrangement of the leaves in relation to each other.
Leaves composed of several parts are compound.
The shape of the leaf at its apex
The shape of the leaf at its base
Pattern of the leaf around its outer edge.
Manner in which the veins of leaves are arranged.
The hairy or scaly surface of the leaf
Notes about the physical characteristics of the plants

7. Feature extraction

There is sufficient information in the echo for a neural network to recognise 1 of 4 plants over a wide range of orientations [Harper and M^{c}Kerrow, 1997]. However, the echo can vary considerably with rotation. To achieve more robust classification we sought to find a set of features that: a) could be easily extracted from the echo, b) were invarient with orientation, and c) represent defined geometric characteristics of a plant.

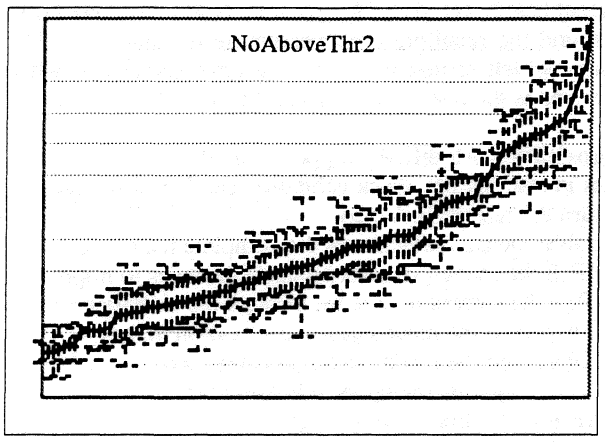

Figure 5 Distributions of values for one feature for all 100 plants.

Table 2 The features which characterise the plant acoustic density profiles

Feature	Description
no_above_threshold 1-9	Counts of the number of range cells at a specified threshold
sum_of_density_profile	Sum of all of the range cells
variance_range, stdev_range, mean_abs_dev_range coeff_of_var_range	The variation of the detected reflections from the central point of the acoustic density profile.
front_to_peak_dist	Distance from the first detectable surface to the surface with the highest amplitude.
length_of_density_profile	The range over which reflections are detected.
freq_75_acoustic_area	The range from the first detected reflecting surface to the cell where 75% of the sum is accumulated.
no_of_major_peaks1 , no_of_major_peaks2	Count of range cells which have reflections significantly stronger than those around it.

Consider the plants shown in Figure 4. The acoustic density profile of the *Polyscias murrayi* has a higher peak than the *Eucalyptus maculata*. So a feature which may distinguish these plants is the maximum amplitude of the acoustic density profile. If this feature is consistent through rotation then it is a good feature for classification.

To establish candidate features, we studied the acoustic density profile of different objects and defined measures to characterise the shape of the patterns. We developed a set of 67 features. Then we reduced this set down to a manageable set of 19 features (Table 2) using mathematics of the visualisation of the distributions.

In Figure 5, a feature and its standard deviation is plotted for all 100 plants. Features were selected based on the slope of the line and the standard deviation. The standard deviation is a measure of how much the feature changes with rotation. The slope of the line is a measure of how easy it is to separate two plants. This feature is a good feature for classification.

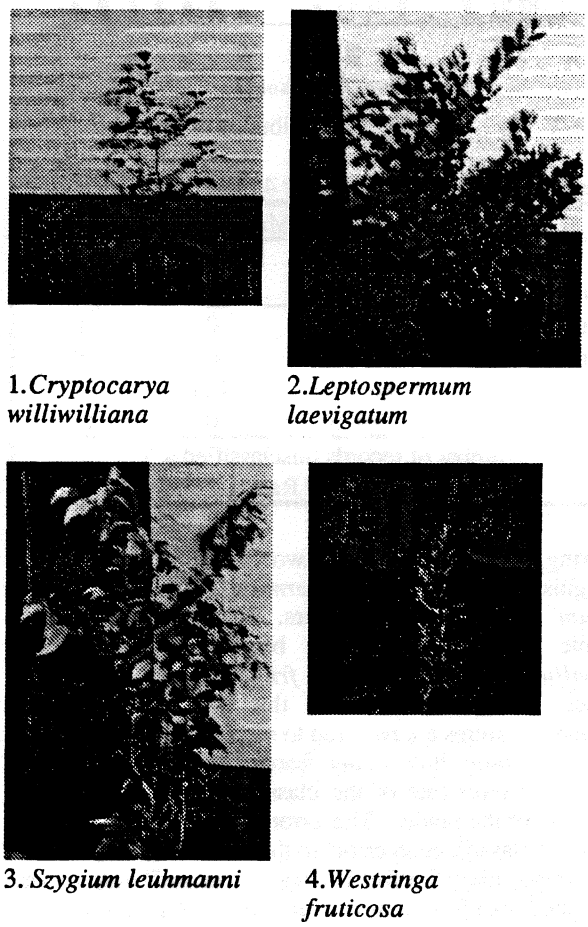

1. *Cryptocarya williwilliana*
2. *Leptospermum laevigatum*
3. *Szygium leuhmanni*
4. *Westringa fruticosa*

Figure 6 Images of four plants

8. Classification of plants

A small number of plants can often be discriminated with a simple statistical classifier using a small set of features. Images of four plants are shown in Figure 6: *Cryptocarya williwilliana* (Small-leaved Laurel); *Leptospermum laevigatum* (Coast Tea-tree); *Szygium leuhmanni*; and *Westringa fruticosa* (Rosemary).

Feature 1 (no_above_threshold1) was calculated for all 360 echo readings for each of these plants. This feature is plotted in a frequency histogram in Figure 7. The x-axis represents the value of no_above_threshold1 and shows values in the range of 0 to 105. The frequency bins are 5 wide and each is a count of the number of records in the dataset which fall within that particular range of values.

The data in Figure 7 falls into three distinct groups. Two of the plants fall in their own separate positions in the histogram (*Leptospermum laevigatum* - around 90 and *Szygium leuhmanni* - around 50). The other two plants (*Cryptocarya williwilliana* and *Westringa fruticosa*) overlap in. bins 15, 25, 30 and 35.

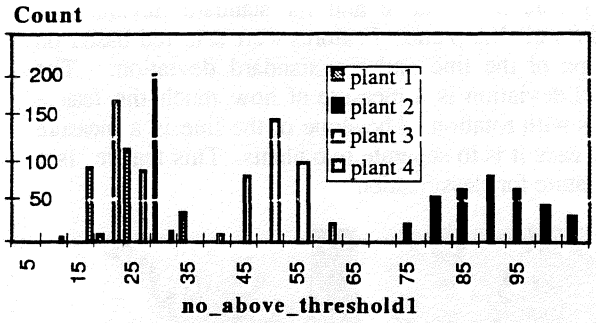

Figure 7 Frequency distribution of feature 1

Table 3 Classification Matrix for 4 plants with 3 features

Classification Matrix				
	Classified			
Actual	Plant 1	Plant 2	Plant 3	Plant 4
Plant 1	88	0	0	6
Plant 2	0	87	0	0
Plant 3	0	0	106	0
Plant 4	1	0	0	80
Number of records misclassified = 7 Classification Error Rate = 1.90%				

Using this single feature, two of the plants can be distinguished from (*Leptospermum laevigatum* and *Szygium leuhmanni*) the others. However, it is not possible to differentiate between *Cryptocarya williwilliana* and *Westringa fruticosa* as they both produce similar values for this particular feature. Additional features are required to separate them..

To determine how many features we need we must measure the error rate of the classifier to find the linear separation of the plants. The error rate is the ratio of the number of classification errors to the number of cases.

A simple classifier for the four plants in Figure 6 was established using template matching. When only the feature in Figure 7 was used the classification error rate was 3.61% and 14 records were missclassified. When 3 features were used (Table 3) the error rate reduced to 1.9% and the number of miss classified records was halved.

9. Feature selection

The performance of a classifier can be improved by removing poor, noisy and redundant features before the data is presented to the classifier. Stepwise forward selection starts with the best single feature. At each step, it tests every one of the remaining features with the current set of features and adds the best one. It continues until it reaches a specified minimum error rate, or until all features have been tried.

We modified this algorithm to arrange the set of features in the order that achieves the best recognition percentage of the cross validation set. The improvement this modification achieved is that it avoids local minima but it has the drawback of testing more feature combinations.

The stepwise forward algorithm provides a strategy to guide the search but we also need a way of evaluating alternative subsets of features. The wrapper method runs the classifier on the candidate features and uses the resulting classification error to select the combination of features with the lowest error.

9.1 The Wrapper Method

We developed a new wrapper method algorithm (Algorithm 1) which combines modified stepwise forward selection and stepwise elimination functions over multiple iterations to calculate a more robust set of features for the particular combination of plants. *Stepwise feature elimination* removes one feature at a time from the current feature set. At each step, each feature is removed in turn and the resulting error percentage is calculated. The feature which results in the lowest error percentage when removed is the one which is taken from the feature set.

Algorithm 1 Modified wrapper method .
Pass in training set, cross validation set, list of features.
Return the features selected.
Initialise the candidate set as the complete feature set
WHILE the recognition percentage continues to improve
  Run the current feature set through the classifier
  Displayed the results as a classification matrix
  Displayed the results as a percentage when tested against the cross validation set.
  IF it is the first time through
    Use the empty set as input to the stepwise feature selection only
      (leave complete set for elimination).
  ELSE
    Perform Stepwise forward feature selection
  END
  Perform Stepwise feature elimination
  IF feature elimination produces better results
    Select that subset as the current set
  ELSE
    Select the subset resulting from feature elimination
  END
END

In Algorithm 1, the initial candidate set of features is a set of all possible features. The modified *stepwise forward feature selection* adds features one at a time based on how good they are at classifying the cross validation set in combination with the features already selected. The stepwise feature elimination algorithm tries to minimise the number of features selected. Whichever of the two feature sets produces the best result are used as input to the next iteration of the loop.

10. Statistical classifier results

When we apply Algorithm 1 to the set of 4 plants in Figure 6, it selects 7 of the 19 features to get a classification percentage of 98.4%. The best features for separating these 4 plants are: mean_abs_dev_range, no_above_threshold6, freq_75_pow, no_major_peaks, coefficient_of_variation_frequency, no_major_peaks2, and no_above_threshold7.

The correlation matrix for these four plants when rotated against a template (Table 4) shows that they are all well correlated through rotations, and so are good candidates for a classifier [Harper, 1998]. The correlation between plants show that *Cryptocarya williwilliana* is the most highly correlated with *Westringa fruticosa* with a correlation level of 0.6. The plants with the highest correlation are most likely to be misclassified.

Table 4 Correlation of the four plants in Figure 6

Plant	1	2	3	4
1	0.9	0.2	0.5	0.6
2	0.2	0.92	0.5	0.2
3	0.5	0.5	0.92	0.3
4	0.6	0.2	0.4	0.79

11. Pairwise classification results

The statistical classifier was run on all pairs of plants to determine which plants are the most similar. Pairwise classification also calculates the average classification of a plant against all of the 99 other plants using the Sonar data. The results of pairwise classification for all 100 plants was an average classification of 90.5% with a standard deviation of 9.6.

Most pairs produced good results and many pairs correctly classify 100% of the test set. But, there are pairs of plants where only a small percentage of the plants are correctly classified. For example, the plants *Ayrtera distylis* and *Cupaniopsis parvifolia* result in a correct classification of only 50.6%, which is a similar probability to a random guess, so it can be said that the two specimens cannot be distinguished from each other. This indicates either that the acoustic density profile changes a lot through rotation or that the plants are similar to each other.

Both of these plants can be easily distinguished from other plants in the population. *Ayrtera distylis* can be classified up to 100% of the test set against other plants in the population (average 85.8%). Similarly, *Cupaniopsis parvifolia* can be classified up to 100% (average 89.79%). This means that the acoustic density profiles of these two plants are very similar. These plants have similar foliage structure: medium to small sized leaves, most of which are horizontal to the ground, and the leaves are relatively spread through the plant

Table 5 Comparison of classifiers with different input data.

	Acoustic Density Profile		Features	
Plant Combination	**TM**	**NN**	**M**	**NN**
Cryptocarya williwilliana *Leptospermum laevigatum* *Szygium leuhmanni* *Westringa fruticosa*	94.96	91.42	98.38	99.73
Ayrtera distylis *Cupaniopsis parvifolia*	68.18	74.29	50.56	77.71
Acacia binervata *Acacia cultriformis* *Azalea splenda* *Polyscias murrayi*	82.24	86.98	83.0	87.95

12. Comparing classifiers

Plants can be classified using linear classifiers (statistical classifier) and non-linear classifiers (neural network). Both classes can use the acoustic density profile or the features as input data. The classification results for these schemes are tabulated for selected sets of plants in Table 5. Row 1 in the table is a comparison for the four plants in Figure 6. Using features results in a higher classification percentage than using the acoustic density profile. From Table 5, we see that features provide better results in all cases. In 5 of the 6 examples, the non-linear classifier produced better results. Also, the results vary widely between different plant combinations.

13. Conclusion

Analysis of the acoustic density profiles of 100 different plants shows that they can be separated using pattern classification when a small sample is considered. As the number of plants increases the classification percentage generally decreases. The result of classifying a single return as one of 100 plants is low. When there is a large number of plants, there is a high probability of plants with similar acoustic density profiles.

Plants can be clustered either on the basis of botanical classification or on the basis of similar acoustic templates. The result is very different clusters. Clustering on the basis of acoustic templates gives similar clusters to clustering on the basis of physical characteristics, due to the fact that the acoustic features can be mapped to geometric features of the plant [Harper, 1998]. Some plants show much more variation with rotation than others. The features show less variation with rotation than the acoustic density profiles.

The statistical classifier has the advantage over the neural net classifier that it provides a list of features that are ordered in terms of their discriminatory ability. This provides important information about the differences between different combinations of plants. A neural net classifier has the advantage that it is a non-linear classifier and generally achieves a better classification result.

14. Acknowledgment

This research was supported by Thomson Marconi Sonar, and the sensing system was purchased from Bay Advanced Technologies in New Zealand.

15. References

Guyer, D.E., Miles, G.E., Gaultney, L.D., Schreiber, M.M., 1993, Application of Machine Vision to Shape Analysis in Leaf and Plant Identification, Transactions of the American Society of Agricultural Engineers, Vol 36(1) January - February, 1993, pp163 - 171.

Harper, N. and McKerrow, P., 1997, *Recognition of Plants with CTFM Ultrasonic Range Data using a Neural Network.*, The 1997 IEEE International Conference on Robotics and Automation

Harper, N. 1998. Classification of plants by the interpretation of CTFM sonar data, PhD Thesis, School of Information Technology and Computer Science, University of Wollongong.

Kay, L.,1974,A sonar aid to enhance spatial perception of the blind : engineering design and evaluation, The Radio and Electronic Engineer, Vol. 44 No. 11, pp605-626.

Nabout, A., Gerhards, R., Su, B., Nour Eldin, A., Kuhbauch, W., 1994, Plant Species Identification using Fuzzy Set Theory, Proceedings of the IEEE Southwest Symposium on Image Analysis and Interpretation, 1994, pp48 - 53.

Singh, S., Montemerlo, M., 1997, Recent Results in the Grading of Vegetative Cuttings Using Computer Vision, Innovative Robotics for Real-World Applications (IROS-97). Grenoble, France, Sep, 1997.