

Automated visual grading of vegetative cuttings

Qiang Ji

Dept. of Electrical Engineering
University of Washington

Sanjiv Singh

The Robotics Institute
Carnegie Mellon University

ABSTRACT

Commercial vegetative propagation of floricultural crops requires the segregation of plant cuttings into categories based on size. The cuttings however must be graded when they are planted ("stuck"), at which time the grade of a cutting is not easy to determine. This paper reports on a system that learns to classify cuttings from being shown examples of images of cuttings that have been graded by a human expert. Based on the example set, the system learns to grade cuttings into categories. We report the results based on a set of 150 geranium plants that were graded by our system and compare the results to the performance of an expert grader.

Keywords: computer vision, image processing, classification, plant grading

1 INTRODUCTION

An important procedure in vegetative propagation of floriculture crops is to grade harvested vegetative cuttings into different categories according to market standards. Currently, the grading process is carried out manually by trained human graders. Human graders visually inspect each vegetative cutting to estimate its physical measurements like stem caliber, length, leaf surface area, weight, etc. Vegetative cuttings are subsequently classified into one of the three to five classes based on estimation of these physical parameters. Manual grading is labor intensive and expensive due to the cost of skilled labor required. In addition, precision varies greatly from one grader to another. Automation offers the opportunity to reduce production costs and improve quality.

Several commercial systems are now available that grade plant cuttings based on their weights. Also, a recent commercial system from Holland grades cuttings based on the leaf area as measured by a machine vision system. While the first approach is simple, it is not accurate since the weight of a vegetative cutting may not correlate well with its grade. The second approach also presents accuracy problem. Placed on a conveyor belt, the leaves may be folded and occluded. It would therefore be difficult to accurately estimate the leaf area from a single view.

Our project is developing a system that learns to grade vegetative cuttings based on features measured by computer vision and grades assigned by a human grader. Computer vision offers an effective means to automate the grading process since it can make full use of all the visual clues used by human operators only at much faster speed. Unlike the Dutch vision system, the system we propose will not only make use of all visual features employed by human graders but also exploit other available visual features. The computer vision system recognizes a vegetative cutting transported on a conveyor and then computes its properties. A classification system then classifies the cutting based on its property measurements from the vision system.

This paper describes the preliminary results of our efforts in developing an automated vision system for analyzing and grading stationary geranium cuttings. Efforts described in this paper are a preamble to a subsequent design and implementation of a commercial robotics vision system for automatically grading geranium cuttings transported on a conveying production line. The geranium cuttings described in this paper are limited to a variety known as veronica. The techniques however are also applicable to other varieties. Early results indicate that the classification performance of the machine classification system is comparable to that of human graders.

This paper divided into six sections. Section 2 discusses the process of acquiring images. Low level processing of the images is covered in Sections 3 and 4. Section 3 discusses the process of segmenting the images into background and plant. Section 4 discusses the techniques for segmenting the plant image into various parts. Extracting and computing the primary characterizing features of each component are dealt with in section 5. Section 6 deals with the issues of the classification. The paper concludes with a summary of the machine classification system and a discussion of the future directions.

2 HARDWARE SETUP FOR IMAGE ACQUISITION

The project described in this paper has concentrated on analyzing stationary images of manually handled vegetative cuttings. By stationary and manual handling, we mean vegetative cuttings to be studied are manually placed on a stationary platform as opposed to being automatically placed a moving conveyor. The stationary platform consists of a base table and a vertically movable camera holder. A high resolution Panasonic super VHS camera is mounted on the movable holder, aiming vertically downward at the vegetative cuttings. The distance from the camera to the table is adjustable, depending on plant size, magnification, and resolution required.

To enhance the contrast between vegetative cuttings and the background, a base table with uniform white surface finishing was selected. This is because the vegetative cuttings tend to look dark in the acquired gray scale image. The enhanced contrast can greatly facilitate subsequent image segmentation through thresholding. To minimize possible shadows and light reflection, additional lighting and photographic equipment were employed. Two photographic umbrellas were used to generate diffuse light to prevent presence of shadows. Caution was also exercised to experimentally adjust lighting strength, direction, and distance to reduce the light reflection from the white table onto the plants.

During image acquisition, vegetative cuttings were manually placed on the base table within the field of view of the camera one at a time. Appropriate camera settings were used. Our experiments reveal that to better distinguish different grades of plants, high resolution camera and digitizer should be employed to compensate for the quantization errors. The increased feature measurement discrimination among different grades of plants is essential for subsequent classification. The acquired images were subsequently digitized using the imaging utilities from the Silicon Graphic work station. The resulting digital images are 640x486, with 256 gray scales. Figure 1 shows images of two different categories (veronica and sincerity) of geranium cuttings.

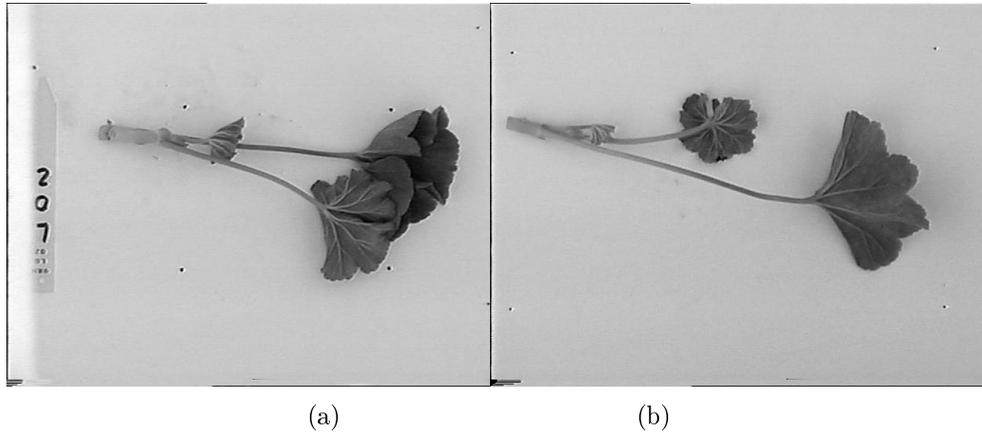


Figure 1: Digital images of two varieties of geranium cuttings (a) Veronica (b) Sincerity.

3 PRIMARY SEGMENTATION OF PLANT IMAGE

The very first task is to distinguish the plant from the background that is we would like to uniquely identify the parts of the image that correspond to the plant cutting. This section examines various issues involved in this process.

3.1 Global thresholding

To recognize the vegetative cuttings from their digital images, regions that contain vegetative cuttings need to be identified and separated from their background. This can usually be accomplished through thresholding if the premise that gray levels for pixels of vegetative cuttings are sufficiently different from those of background holds true. This assumption is basically true for this project since gray levels for vegetative cuttings tend to be darker than those of background. Thresholding converts an input gray scale image into a binary image, with white pixels representing vegetative cuttings and black pixels background.

To automatically threshold an input gray scale image requires automatic determination of an appropriate threshold value. An appropriate threshold value is crucial since a too high threshold may incorrectly label background pixels as vegetative cuttings while a too low thresholding may label part of vegetative cuttings as background. In either case, errors would be introduced to subsequent feature measurements. Based on the underlying intensities distributions, various techniques have been developed for automatic selection of an optimal threshold value. Employing different specialized principles, these algorithms usually require that the intensity distribution of the image be bimodal. The relative high contrast between vegetative cuttings and the background allows most of these thresholding algorithms applicable to our problem.

Among these algorithms, we implemented the algorithm developed by Otsui⁴ due to its simplicity in implementation and fast execution speed. According to this algorithm, given a threshold, it divides the intensity distribution into two groups, representing the object and background respectively. Means and variances are then computed for each group. The sum of the two variances, weighted by the probabilities of each group, is the criterion function that needs to be minimized in order to identify the best threshold value. So the best threshold yields the smallest weighted sum of within-group variance. Given this criterion function, this technique searches for all possible threshold values to locate the threshold that minimizes the sum of variance. In actual implementation, a recursive procedure was implemented to use results computed for the previous threshold value for next threshold

value. This can result in substantial saving in computation time. Figure 2. shows the binary image generated by Otsui's algorithm for the veronica cutting shown in figure 1(a)

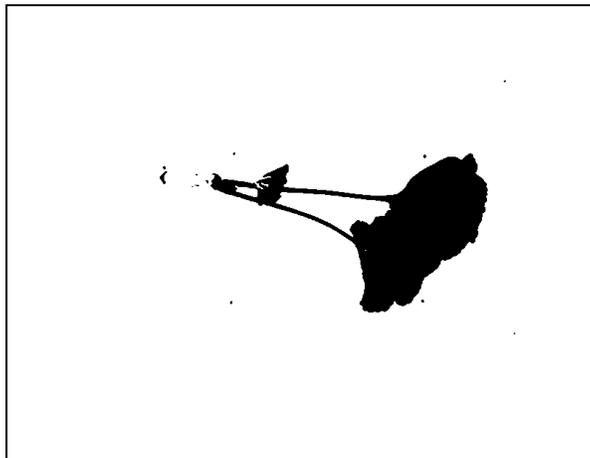


Figure 2: The binary image for the veronica cutting shown in Fig. 1(a) after global thresholding

3.2 Adaptive thresholding

Despite our efforts during image acquisition to increase the contrast between background and plants and to reduce the shadows, we still can not completely eliminate shadows and generate an uniformly high contrast throughout the vegetative cutting. This implies a global thresholding, which does not account for local contrast variation, may not always yield desired results as shown in figure 2. Compared with the original image, it is clear from figure 2 that most part of stem was mis-classified as background. On the other hand, some of the shadow near the leaves was mis-classified as leaves. This would result in incorrect identification of stem and leaves. In addition, the mis-classification of part of stem as background may in some cases result in tiny connections between different parts of stem. This could pose a serious problem for subsequent stem identification.

In view of above problem with global thresholding, an adaptive thresholding technique was implemented. The adaptive threshold can account for the spatial variations in image contrast and compensate for a variety of local spatial context effects by using a spatially varying threshold. The adaptive thresholding we implemented is a modified version of the technique developed by Chow and Kanade.² The technique starts with dividing an entire image into mutually exclusive blocks. Based on experimentation, the optimal block size for our project was found to be 80x54. It assumes that the intensity distributions for block regions consisting of only the object or background are unimodal and that distributions for regions containing both object and background are bimodal. Furthermore, the bimodal distributions can be modeled as the sum of two Gaussian distributions. The variances and means are subsequently computed for each block. The histograms of blocks with large variance are then computed.

For each selected histogram, a bimodal test is conducted to determine whether the distribution is bimodal or unimodal. The criteria for bimodality test include difference in means and ratio of variance. For every histogram with appreciable bimodal distribution, the Otsu technique is used to compute an optimal threshold. For regions that are not assigned thresholds, their thresholds are subsequently estimated as weighted averages of the computed thresholds of their neighboring regions. Finally, a binary decision is performed for each image block using the thresholds obtained. Figure 3(a) shows the binary image resulting from Chow's technique.

While adaptive thresholding can take the local contrast into consideration, it is so sensitive to small local spatial variations in intensities that it may classify certain details of plants as background, especially in leaf area

as shown in Figure 3(a). This is undesirable for we do not want to split a plant into different parts. To avoid this, the algorithm was modified as follows. A global threshold is obtained first. For blocks with mean intensities lower than global threshold (this is mostly in leaf area), the global threshold is used for thresholding. For blocks with mean intensities higher than global intensities (mostly this is in stem region), adaptive thresholding is used. Figure 3(b) is the binary image obtained using the modified adaptive thresholding technique.



Figure 3: Binary images resulting from adaptive thresholding (a) and modified adaptive thresholding (b)

4 IMAGE SEGMENTATION

Once we recognize a vegetative cutting and separate it from its background, our next task is to identify the primary components of a vegetative cutting. Generally speaking, a geranium plant consists of a stem, leaves, and leaf petioles as indicated in figure 1. The task of segmentation is therefore to segment the plant blob we obtained from thresholding into three distinct smaller blobs representing leaves, leaf petioles, and stem respectively.

Various techniques have been developed to perform segmentation. For example, Peleg⁵ suggested use of curvature of boundary edge pixels for segmenting a plantlet image into distinct blobs. The basic premise of this method is that each blob can be delimited by two edge points with maximal curvatures. For example, the two intersection points between a stem and leaf petioles have curvatures higher than those of the rest of edges points on the stem and therefore can be used to delimit the stem from leaf petioles. The problem with this technique is that it is time consuming (requiring to identify the boundary of the plant and compute the curvature of each boundary point) and sensitive to noises and local irregularities.

Instead, we observed that there exist substantial morphological differences between leaves, leaf petioles, and those of stems. For example, the stem caliber is usually three to five times larger than those of leaf petioles while leaf widths are normally at least twice as big as those of stems. Based on these observations, a mathematical morphological algorithm was developed for segmentation. The algorithm consists of three steps. First, it separates stem and leaves from leaf petioles. This can be achieved by removing leaf petioles. Based on their differences in size, the leaf petioles can be removed through a morphological opening operation. A square disk structure element was designed. The width of the structure element was selected such that it is slightly larger than the diameters of leaf petioles while much smaller than those of stem and leaves. As a result, a single opening operation removes leaf petioles from the original binary plant image. The resulting blobs consisting of leaves and stem tend to swell in perimeter due to the opening operation. This is because an opening operation consists of an erosion operation, followed by a dilation operation. While the erosion operation removes the leaf petioles, it has minimal impacts on the stem and leaves blobs. The subsequent dilation operation, however, smooths the contour, breaks narrow isthmuses, and eliminates small islands of stem and leaves. The expansion in blob perimeter must be minimized

so that stem and leaves can retain their original shape and contours. It ensures the accuracy of the measurements of the characteristic features of each blob. To do so, the resulting image from the opening was AND with the original binary image, generating a new image that contains only leaves and stem with their original shape and contours.

The second step is to identify the leaf petioles. This can be easily obtained by subtracting the above leaves and stem blobs from the original blob image. The last step is to separate the stem and leaves. Distinguishing between leaves and stems is made based on their inherent physical and geometric differences like shape and size, difference in mean intensity, and relative location in the image. Geometrically, stems tend to be of a rectangular shape while leaves can assume any shape, although most likely fan-shaped. Physically, stems are usually smaller in area than leaves. Intensity-wise, stems tend to be brighter than leaves. It is these differences plus any prior knowledge about the location of stem and leaves that are used for separating leaf blob from stem blob. Figure 4. shows the extracted plant components for the binary image shown in figure 3(b).

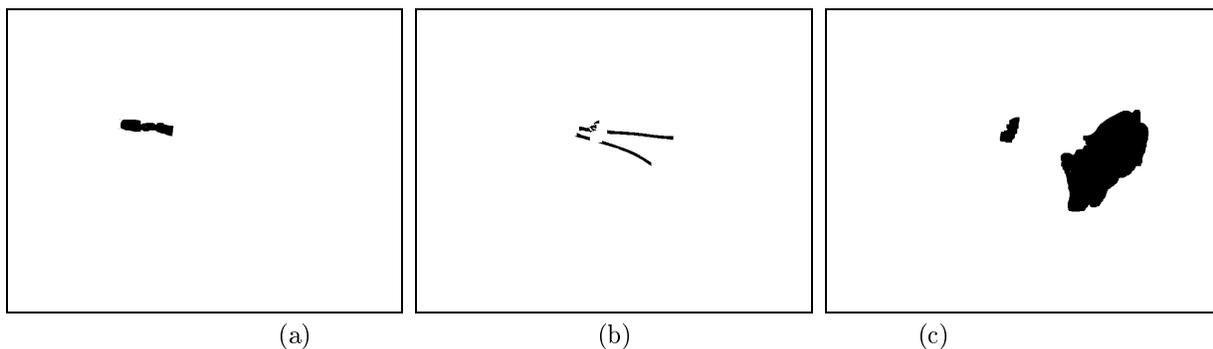


Figure 4: The extracted plant components: stem (a); leaf petioles (b); and leaves (c).

5 MEASUREMENT OF PLANT COMPONENTS

Now that we have successfully segmented original images into three meaningful connected components: stem, leaf petioles, and leaves, it is necessary to compute the primary characterizing features of each blob. The primary characterizing features of each blob will be used for classification. Based on what we learn from human graders, the primary characterizing features for each blob are defined as follows. For stem blob, the primary characterizing features are its length and its caliber since they are the most important criteria used by human operator. For leaf blob, the primary characterizing feature is its area. It approximates the leaf surface area. This feature is, however, ill-defined since we can not always accurately estimate the leaf area just from one view point. Leaves, when placed on a table, may be folded or occluded by each other. Therefore, even for the same plant, its leaf area may vary considerably, depending on the camera angle.

Finally, for leaf petiole blob, the primary characterizing features can be the number of leaf petioles, the diameter of each leaf petiole, and the length of each leaf petiole. The number of leaf petioles is important since it indirectly gives us the number of leaves. During human grading, the number of leaves is often used as a tie-breaking factor when stem length and caliber contradict each other. The diameter of each leaf petiole is a feature not used by human grader. But we feel that leaf diameter, especially the mean diameter of all petioles, may somehow correlate with leaf area. In fact, intuitively, there seems to be a positive correlation between them, i.e. the larger leaf surface, the larger the mean leaf petiole diameter. The potential use of this correlation is that we may use mean leaf petiole diameter as a replacement for leaf area since mean diameter of leaf petiole can be estimated much more accurately than leaf area. We will examine the correlation during the classification phase. Finally, length of leaf petioles do not seem to play any role in human judgment of plant grade. It is therefore not included as the primary characterizing feature of leaf petioles.

5.1 Computation of stem length and caliber

Once the stem has been isolated from the rest of a plant, we need to estimate its width (caliber) and its length. Since the shape of a stem roughly approximates a rectangle with many local irregularities, it is difficult to estimate its length and width directly from the stem blob. To have a better estimate of the length and width of the stem, we decided to fit a rectangle to the stem blob. Our assumption is that the length and width of the best fitting rectangle should constitute an optimal estimate of the length and width of the stem.

To fit a rectangle to a stem blob, we developed an iterative gradient-descent rectangle-fitting method. Like the popular ellipse¹ fitting algorithm, this algorithm is based on the premise that the spatial moments of a mass can fully characterize its orientation, location, its shape, and its contour. Though having been successfully applied to ellipse fitting, this method for rectangle fitting has not been found in literature.

Unlike the ellipse fitting, which may yield a closed-form solution, our algorithm is an iterative procedure. The algorithm starts with computing the first and second order moments for a rectangle. Moments include first and second row and column moments, and their mixed moments. It then computes the same set of moments for the stem blob. If the shape of the stem were a perfect rectangle, we could simply equate above two sets of moments to have a closed form solution for the length, width, and orientation of the rectangle. Since above assumption does not hold, we can not solve the stem length and caliber in a closed form. Instead, we try to identify unknowns (length, width, and orientation) that minimize the differences between above two sets of moments. The error function can be formulated as follows:

$$\epsilon^2 = (M_{cc}^b - M_{cc}^r)^2 + (M_{rr}^b - M_{rr}^r)^2 + (M_{rc}^b - M_{rc}^r)^2 + (A_b - A_r) \quad (1)$$

where M_{rr}^b , M_{cc}^b , M_{rc}^b , and A_b are the second moments and area of the stem blob and M_{rr}^r , M_{cc}^r , M_{rc}^r and A_r are the corresponding moments and area of the fitting rectangle.

Standard methods exist for computing the second moments of a blob. Readers may refer to Haralick¹ for details. The moments for the fitting rectangle can be easily derived as follows. Let W and L be the width and length of the fitting rectangle, we have

$$\begin{aligned} M_{cc}^r &= \frac{L^2}{3} \\ M_{rr}^r &= \frac{W^2}{3} \\ A^r &= WL \end{aligned} \quad (2)$$

the cross mement M_{rc}^r is zero since moments are calculated in the new coordinate frame defined by the fitting rectangle. With the error function in eq(1), following iterative equations can be obtained using the gradient descent method to update the length and width of the fitting rectangle. The initial estimate of the length and width may be obtained from a fitting ellipse.

$$L = L - \alpha \frac{\partial \epsilon^2}{\partial L} \quad (3)$$

$$W = W - \alpha \frac{\partial \epsilon^2}{\partial W} \quad (4)$$

where α is the convergence rate.

Figure 5 illustrates the image of another cutting, with its components identified (shaded with different gray levels). The stem part of the cutting is fit with a rectangle.

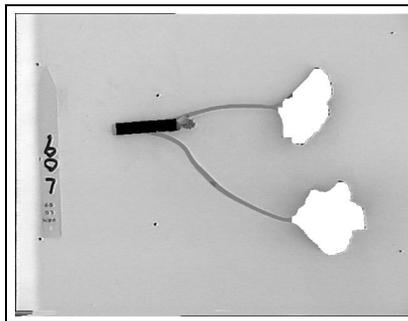


Figure 5: A rectangle (blacken part) that has been fitted to a stem

5.2 Computation of other properties

If we assume the leaf petioles are cylindrical, the two dimensional projection of leaf petioles on the image plane can be assumed to be something of ribbon shape with uniform width (or diameter). To estimate the diameter of each leaf petiole from its image, we estimate its area and length. The diameter of the petiole is therefore the division of its area by its length. The petiole area is simply equal to the number of pixels in the blob. The length of petiole can not be estimated directly. Instead, a skeletonizing operation is performed on the leaf petiole blobs to find their median axes. It is assumed the length of median axis equals to the length of the petiole. With the length and area known, the diameter of the petiole can be easily derived.

To compute the number of leaves, we use leaf petiole blobs instead of the leaf blob since leaf may be folded or hidden. It is difficult to estimate the leaf number from leaf blob. On the other hand, it is relatively easy to estimate the leaf number from leaf petioles. The number of leaves is simply equal to the number of leaf petioles. Finally, the leaf area is simply equal to the number of pixels occupied by the leaf components.

6 CLASSIFICATION

The outputs from the vision algorithm for a plant cutting are the measurements of each feature as discussed in section 5. These measurements do not directly indicate the membership of the plant, i.e., whether a plant belongs to small, medium, or large class. This is the task of classification. Classification involves a training process, during which the learning software is presented with a set of data points along with their output values. The learning method derives underlying relationships between the features and the classification and generalizes the relationship to predict the output of an unseen input data. For this project, we used a software package based on Generalized Memory Based Learning (GMBL)³ for classification. One of the most important features of this package is that given a dataset and original feature set, it can automatically rank each feature according to its contribution to the overall classification performance. This is important in that it allows to reduce original feature set to a subset of optimal features, therefore reducing the learning time with minimum loss in classification accuracy.

6.1 Data Acquisition And Data Preparation

To train the classification software, feature data were collected for 150 veronica cuttings. Data include measurements for five features-stem caliber, stem length, leaf number, leaf area, and mean leaf petiole diameter. Plants can be classified into three classes; large, medium, and small. Three classes of cuttings were uniformly represented in the dataset, with 50 cuttings for each class. Also included in the data are the label for each class. The input vector for the GMBL software therefore consists of five feature measurements while the output includes three nodes representing the three classes respectively.

For training, the class membership of each plant was determined manually by an experienced human grader through visual inspection. The accuracy of manually determining the class membership of each class is critical since subsequent machine training is based on this classification. To improve the accuracy of the training data and to study the machine performance relative to the manual performance, two sets of training data were collected. First, a very good human grader was employed to carry out the classification, yielding the first set of training data; second, after initial classification and image acquisition, the 150 veronica cuttings were grown for four weeks for subsequent reclassification. The grown cuttings were reclassified manually. Based on the reclassification, the class membership for previous 150 cuttings were modified, resulting in the second dataset. The premise here is that since the second set of data involves classifying grown vegetative cuttings, it should be more objective and accurate than the first grading even though domain experts still perform the actual classification. In the subsequent sections, we describe the training process using both sets of data. The two dataset are referred to as uncorrected dataset and corrected dataset respectively.

6.2 Experimentation

This section will discuss the performance evaluation of the machine inspection system using the acquired data. Since we do not have sufficient data to allow us to have a separate training and testing set, leave-one-out (LOO) cross-validation was used for training and testing. The results are shown in a confusion matrix, where the horizontal classes are the true classes while the vertical classes represent predicted classes.

6.2.1 Classification using uncorrected data

In this section, we will perform a classification on the uncorrected dataset. A total of 144 feature vectors were included in the dataset (the vision algorithm failed to compute the measurements for six of the 150 plants). The experiment starts with feeding the dataset to the GMBL software to allow it to search for the best GMBL features. For the given dataset, the best features were found to be feature number 1 (stem caliber), feature number 3 (leaf number), and number 5 (leaf petiole diameter). Using the best string, LOO classification was performed. The classification results are shown in confusion matrix in Table 1.

CLASS	1	2	3	%Correct	% Error
1	36	4	8	75	25
2	11	21	17	42.86	57.14
3	11	13	23	48.94	51.06
Total				55.31	44.67

The average classification accuracy is about 56%. Breaking down into each class, it is noted that the small class has the highest classification accuracy of 75% while the medium has the lowest of 42.86%, followed by the large of 48.9%. We may infer that there are more ambiguities involved with medium classes than small and large classes. This seems to be intuitively correct. Further analysis shows that 33% of the misclassifications for small

class is due to classifying small as medium while the remaining misclassification (66%) is due to classifying small as large. This counter-intuitive phenomenon may either be caused by the imprecision in measurements from the vision algorithm or great inconsistency in human grading. For median and large cuttings, most misclassifications occurred in the neighboring classes as expected. Data from all three categories seem to suggest that human graders tend to be more confused between large and medium than between medium and small.

To summarize, the best performance for the software using the uncorrected is a classification accuracy of 56%. As discussed above, the low classification accuracy primarily results from random errors from inconsistencies in human grading, the ambiguities between the neighboring classes, the lack of quantitative measurements for each feature (the measurement of each feature is very fuzzy), and the lack of systematic protocol for human grader while visually inspecting each cutting.

6.2.2 Classification using corrected data

In pervious data set, the plant membership was manually determined by the best human grader. As a result, our vision system can at best do as well as the human grader. Also, we have realized that due to random errors from inconsistencies in human grading, substantial overlapping and bias exist in the previous dataset. In this experiment, we obtained the corrected data in the hope to greatly reduce such ambiguities and overlapping. To do so, the pervious dataset are reclassified manually using the grown vegetative cuttings. The premise here is that it is much easier to distinguish the three classes once they grow up, and less subjectivity (more objectivity) involved in the classification, therefore bringing the class membership of each cutting closer to its true class.

Based on the reclassification, the original data set needs to be redistributed. Memberships of some plants are changed according to the new classification.

Considering the fact that ambiguities still exist between neighbor classes (this is observed during reclassification), we divide the cuttings into five classes; large, large-medium, medium, medium-small, and small to minimize overlapping and bias. The premise here is that most ambiguities will be concentrated in the two in-between classes (medium-large and medium-small), therefore minimizing the ambiguities associated with large, medium, and small. During the learning process, only the data from three distinct classes (large, medium, and small) are used for training. The advantage with above training method is that with reduced random noises and ambiguities associated with each class, the classification software may find a better and finer decision boundaries among three classes. The testing data however will involve data points from all five classes.

In this experiment, we study the classification performance of the vision system using the corrected data. In the meantime, we will also compare the machine classification accuracy against that of human. This experiment starts with identifying the best feature set. The best features are found to be the stem caliber and the mean leaf petiole diameter. LOO cross-validation was then performed in two cases as shown below.

Case 1: Of the five classes, data points belonging to the in-between classes (like medium-large class) were removed. LOO was performed on the dataset without data from in-between classes. The confusion matrix for both the human and machine classification are shown in tables 2 and 3:

Table 2 Human Classification						Table 3 Machine Classification					
CLASS	1	2	3	%Correct	% Error	CLASS	1	2	3	% Correct	% Error
1	42	4	0	91	9	1	33	13	4	66	34
2	5	29	1	82	18	2	12	25	7	57	43
3	2	14	29	65	35	3	5	6	18	62	38
Total				79	21	Total				62	38

The total average classification rate for machine and human is 61.79% and 79% respectively. The average

machine classification accuracy improves by 7% from previous 55%. Breaking down to each class, the small has a classification rate of 66% (down 9% from previous 75%), the medium 57% (up 15% from previous 42%) and finally large 63% (up 14% from previous 49%). The improvement for large and medium classes is substantial. But, there is a drop for small class. One possible explanation for the drop is that the three classes are not equally represented in the data set, with 50 small cuttings, 44 mediums, and only 27 large cuttings.

Looking at the data from machine classification, it shows that most of misclassification occurs in neighboring classes and that the largest confusion occurs between small and medium classes. This echoes the results from human classification. Based on above analysis, it can be seen that the performance of machine classification is roughly comparable to that of human.

Case 2: In this experiment, we will include all data points. For human grader, the data points classified into one of the in-between classes are counted as correct classification. For machine, a data point belonging to the in-between class is correctly classified if it is classified as one of the two classes involved. For example, if a point belonging to medium-large and is classified as either medium or large, then the classification is correct. However, if it is classified as small, then it is a mis-classification. For machine classification, only data belonging to the three distinct classes are used for training and all data points are used for tests. The results are shown in Tables 4 and 5.

CLASS	% Correct
1	92
2	88
3	68
Average	83

CLASS	% Correct
1	68
2	68
3	63
Average	69

The result from the second experiment echoes the conclusion we draw from the first one that the performance of machine classification is comparable to that of human.

7 CONCLUSIONS

This paper has described an automated visual inspection system for grading vegetative cuttings. Our preliminary performance evaluation of the system revealed that its performance at present stage is approximately comparable to that of the best human grader. Further improvement may be accomplished in several ways. First, the resolution of the images can be increased, allowing for better discrimination of features between different grades of cuttings. Second, segmentation algorithms can be made more robust and accurate. Finally, different types of classifiers may be investigated. Our preliminary investigation of neural networks showed a 7% performance improvement over the GMBL software.

8 REFERENCES

- [1] Robert M. Haralick and Linda G. Shapiro. *Computer and Robot Vision*, volume 2. Addison-Wesley Publishing Company, 1993.
- [2] T. Kanade and K. C. Chow. Boundary detection of radiographic images by a threshold method. *Frontiers Of Pattern Recognition*, 0(0):61–82, 1982.
- [3] A. Moore. General memory based learning, a brief explanation of the concept, the software and the c interface. Technical report, Carnegie Mellon University, 5000 Forbes Ave., Pittsburgh, PA 15217, 1994.
- [4] N. Otsui. A threshold selection method from gray-scale histograms. *IEEE Transaction on System, Man, and*

Cybernetics, 9(0):62–66, 1979.

- [5] K. Peleg, O. Cohen, M. Ziv, and E. Kimmel. Machine identification of buds in images of plant shoots. *Machine Vision And Applications*, 6(0):624, 1993.