J. agric. Engng Res. (1996) **64,** 165–172

Location and Characterization of the Stem–Calyx Area on Oranges by Computer Vision

L.A. Ruiz;* E. Moltó;* F. Juste;* F. Plá;† R. Valiente‡

* Instituto Valenciano de Investigaciones Agrarias (I.V.I.A.), Apartado Oficial 46113 Moncada, Valencia, Spain
† Departament d'Informàtica, Universitat Jaume I, 12071 Castellón, Spain
‡ Centro de Cálculo, Universidad de Valencia, Burjassot, Valencia, Spain

(Received 6 February 1995; accepted in revised form 11 October 1995)

Three image analysis methods were studied and evaluated to solve the problem of removing long stems attached to mechanically harvested oranges: colour segmentation based on linear discriminant analysis, contour curvature analysis, and a thinning process which involves iterating until the stem becomes a skeleton. These techniques are able to determine the presence or absence of a stem with certainty, to locate the stems from random views with more than 90% accuracy and from profile images with an accuracy ranging from 92·4% to 100% depending on the method used. Finally, determination of the length and cutting point of the stem is achieved with only 3·8% of failures.

1. Introduction

Mechanical harvesting of citrus fruits brings some additional problems that have not been present in manual harvesting, such as the presence of fruits with long stems, with leaves or without calyx after detachment from the tree. Long stems and leaves can cause damage on adjacent fruits, while the absence of calyx opens a way for possible infections during transport and storage. This also means a loss of uniformity of the product which is not desirable for the fresh market. Therefore, a system for cutting long stems and for detecting the absence of calyx before the fruit arrives at the packing houses would be advantageous.

Mechanical destemming systems, such as the one reported by Chen,[1] are usually based on random rotation of the oranges against cutting surfaces. However, the contact with these surfaces may cause some damage to the fruits. Some researchers have worked on image analysis techniques to orientate fruits and vegetables or to locate stems or some protruding parts of the plant. Most of the algorithms, which use a profile view of the fruit, are based on contour analysis and detection of sharp changes in the contour itself. Thus, Wolfe and Sandler[2] developed a fruit stem detection algorithm using profile images of blueberries. Another computer vision algorithm was proposed by Langers[3] to locate the growing point of flower bulbs. Howarth *et al.*[4] studied the tip shape in carrots and analysed the curvature profile, classifying carrot tips into five classes. Yang[5] used structured light to distinguish between the stalk and calyx of apples from true blemishes.

The random orientation of fruit on conveyor belts is often a problem in stem cutting systems. The sphericity of oranges hinders the possibility of mechanical orientation, so the detection of the stem–calyx area by a camera seems to be an adequate way to orientate the fruit. Once the stem has been situated at the correct position, it can be characterized and measured so the decision to cut the stem off or not can be made. A system for orientation of oranges would be of interest for implementation on the CITRUS robot[6] in order to cut the stem off after the picking operation.

The objectives of this study were to design reliable image analysis methods (1) to locate, using colour vision, the stem–calyx area of oranges randomly presented to a camera, in order to be able to orientate the fruit, as well as to classify it on the basis of presence or absence of stem and leaves, and (2) to study the profiles of the fruit previously oriented and locate the stem to determine its length and cutting point.

2. Materials and methods

Colour images were acquired with a charge coupled device (CCD) video camera. Red, green and blue

Presented at AgEng 94, Milan, Italy, 29 August–1 September 1994.

0021-8634/96/070165 + 08 $18.00/0

(RGB) video signals were digitized by a colour frame grabber connected to a PC bus, to form colour images having $512 \times 512 \times 24$-bit pixel resolution. Grey level images were acquired with the CCD video camera and a video digitizer board captured the images of $512 \times 512 \times 8$-bit pixel resolution. The spatial resolution provided by the digitizing boards was excessively large for colour and profile analysis, so the images were subdivided into 256×256 pixel subsets.

The sample used for this study was composed of 120 oranges, variety "Salustiana", manually collected from the field, which kept the same proportion of long-stemmed oranges that were harvested with the CITRUS robot[6] and reported by Fornes *et al.*,[7] in such a way that the results could be evaluated and referenced with respect to the robot performance.

Since diffuse light is highly effective in eliminating shadows and specular reflection, and in preserving well-defined edges, an illumination chamber with indirect fluorescent light and diffusing material was built and used to take colour images, while illumination by contrast was employed to acquire profile images.

The following three working methods were defined to achieve the objectives. The first method is focused on the orientation problem, while the aim of the second and third method is the definition of the stem cutting point of pre-oriented oranges. (1) Colour segmentation of images by lineal discriminant analysis techniques to detect the presence or absence of stem and leaves, with estimation of spatial coordinates of the stem insertion point using information of size and centroid position. (2) Contour extraction from binary images and contour tracking to study local curvatures and several distances from a set of contour points to the centroid, to detect the insertion and end points of the stem. (3) Analysis of profile images by an iterative thinning process, until the skeleton of the stem is obtained, so that it can easily be located on the contour of the fruit.

2.1. *Colour segmentation procedure*

Colour machine vision has been widely studied for classification tasks of agricultural products in cases where grey level information was not sufficient to obtain accurate results. Slaughter and Harrell[8] developed a classification model to discriminate oranges from their natural background using digital colour images. Miller and Delwiche[9] found a high correlation between colour maturity classification of peaches with a computer vision system and manually. Good results were obtained by Alchanatis *et al.*[10] in the segmentation into distinct colour regions and subsequent classification of tissue culture segments of potato plantlets. Paulsen *et al.*[11] used colour average in RGB and HSI (hue, saturation and intensity) coordinates for specific regions of kernels in inspection and classification of maize.

The classification and stem location algorithm proposed is based on a colour segmentation of the images using the red, green and blue colour space. This colour space was chosen because it gave better results than the hue, saturation and intensity space in some preliminary trials.

Eighty six colour images were taken using diffuse light inside the illumination chamber as described before, with a dark background to provide a high contrast with the fruit. The harvesting performance of the CITRUS robot was simulated for an average of three different varieties of citrus. Based on the results obtained by Fornes *et al.*,[7] 10% of the chosen images presented fruits with some leaves attached to the stem.

The segmentation process was based on a linear discriminant analysis, involving the following steps. (1) Definition of colour classes in which the images are segmented. (2) Estimation of the covariance matrix of the red, green and blue variables, from an initial set of samples or *training set*. In order to simplify the process, equal covariance matrices were assumed for the four classes considered, so that the samples for each class were considered to be distributed in clusters of equal size and shape, and centered about their respective mean vectors. (3) Development of a classifier based on Bayesian decision rules, assuming that the independent variables (RGB) follow multivariate normal distribution functions. (4) A testing process to evaluate the performance of the model, using a set of samples, or *testing set*, independent of the *training set*, to obtain a classification matrix in order to compare the actual and the estimated classifications.

Four colour classes were defined to represent those parts of the images relevant to determination of the location of the stem: namely, *background, peel, stem-calyx* and *cut stem*. This last class, represents those parts of the stem that have been torn during the harvesting process, and is useful in some cases to better estimate the insertion point of the stem, since it is not considered to determine the centroid of the stem class [*Fig. 1*(*c*)]. A total of 15 877 pixels were sampled from ten images and used as the *training set*. These pixels from the four classes were manually selected to represent a wide range of RGB values present in each class.

The Bayesian classifier is based on the computation of the *a posteriori* probabilities that a sample belongs to each class (Duda and Hart[12]). In our case, these are the probabilities that each pixel from the training set belongs to *background, peel, stem* or *cut stem*, according to its values of red, green and blue components. From the results obtained from an estimation made over several images, the *a priori* probabilities entered for each class were 0·4, 0·4, 0·18 and 0·02 respectively.

The *testing* set was composed of 50 427 pixels randomly selected from 24 different images. The implemented classifier was tested over this set of data and its performance evaluated.

Four decision functions (one per class), each a linear combination of the independent variables were computed. Each of these functions is directly related to the probability that a pixel with a given set of RGB values belongs to the corresponding class. They were implemented in an algorithm to segment the colour images, so each pixel on the image was classified into the class for which the decision function value was higher, according to the respective RGB values.

The resulting class decision functions were

$$f_{\text{peel}} = 0{\cdot}7029R - 0{\cdot}0396G - 0{\cdot}0062B - 66{\cdot}69$$

$$f_{\text{stem/calyx}} = -0{\cdot}0178R + 0{\cdot}3629G - 0{\cdot}0696B - 23{\cdot}46$$

$$f_{\text{cut stem}} = 0{\cdot}2089R + 0{\cdot}3727G + 0{\cdot}0708B - 69{\cdot}22$$

$$f_{\text{background}} = 0{\cdot}1581R - 0{\cdot}0092G + 0{\cdot}1423B - 9{\cdot}91$$

where R, G and B represent the values of the three colour components, ranging from 0 to 255. The coefficients of the variables for these class related functions are greatest for those colour components that are dominant in the corresponding class, and were obtained using information of averages and covariances from the training set, as well as *a priori* probabilities of the classes. Thus, as was expected, the red component is higher for the class *peel*, the green for the *stem/calyx*, and the red and blue for the *background*, while the *cut stem* class presents relatively high coefficients for the green and red components.

The detection of de-stemmed oranges was achieved by computing the total area of the pixel classes *stem* and *cut stem*, applying the condition

If $(\text{Area}_{\text{stem}} + \text{Area}_{\text{cut stem}}) < l$, then neither stem nor calyx are present.

Here, "l" was an empirical limit of the *stem–calyx* area under which the orange was considered without stem. In general, this limit will depend on the resolution of the image.

Finally, and only for the fruits that were decided to have stem or calyx after applying this rule, the estimation of the insertion point of the stem, or point of attachment of the stem to the fruit, was made by computing the centroid coordinates of only those pixels belonging to the class *stem–calyx* that were inside the convex contour of the class *peel*. Thus, the possible errors in cases where stems projected out of the fruit surface were corrected by only considering those pixels belonging to the class stem inner to the fruit contour to determine the insertion point. Two examples of projected stems projecting out of the profile of the class peel are shown in *Fig. 1* (a) and (b), where the detection of the insertion point is made correctly.

3.2. *Contour tracking*

This method was employed to detect the stem and to measure its length by studying the profile of pre-oriented fruits. Images with 256 grey levels were taken using contrast illumination, having a bright area corresponding to the background and a dark area corresponding to the fruit, so they were easily thresholded to obtain binary images. The algorithm developed, first extracts the fruit contour, then divides it into segments to estimate local curvatures and finally calculates distances from some specific points to the fruit centroid. The processes are as follows.

Contour extraction: a sweeping of each binary profile image is made until the first pixel of the fruit is reached. Then, the next point of its neighbourhood is looked for following the most probable direction to find it. The most probable direction is the current direction, then the two adjacent directions forming a 45° angle to the current direction and so on. Small projections inward or outward of the contour are removed, with the subsequent smoothing effect over the boundary (*Fig. 2*), desirable for the next operations over the contour.

Segmentation of contour: an optimum and constant number of fragments (n) is chosen ranging between 80 and 100, with k pixels each, where $k = N/n$ and N is the total number of pixels of the contour.

Detection of *maximum curvature*: this is calculated as the maximum difference of slopes between adjacent segments. Then its neighbourhood is analysed to find the farthest pixel from the centroid of the fruit, that will be the end of the stem (point A, *Fig. 3*). From this point, the contour is followed to both sides until points B and C are found (*Fig. 3*), characterized by a strong descent of their distance to the centroid before them (zones D, E, *Fig. 3*) and a practically constant distance after that (zones F, G, *Fig. 3*). The length of

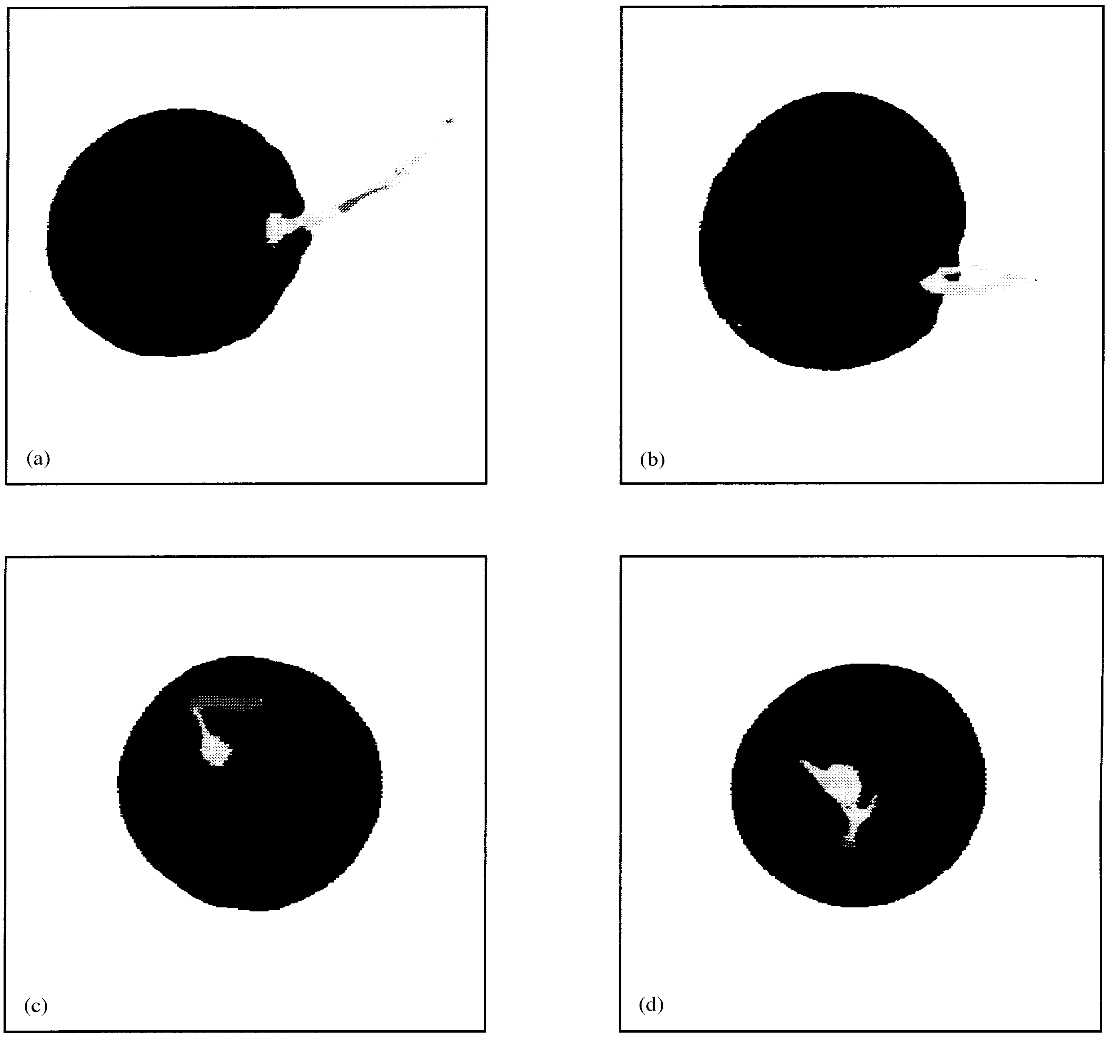

Fig. 1. Examples of segmented images of oranges with indication, by mean of a white cross, of the insertion point estimated by the location algorithm. Each class is represented by a different grey level: peel (black), stem–calyx (light grey), cut stem (dark grey) and background (white).(a) and (b) show the effect of using only the inner pixels of the stem–calyx class to compute the insertion point. (c) pixels of class cut stem are not considered to locate the centroid of the stem. (d) another example

the stem is then estimated as the distance between points A and B.

3.3. *Thinning approach*

In general, on citrus and other fruits, the contour of the fruit in a profile view has a smooth curvature, and the presence of stems produces salients or projecting parts in this smooth contour. Given a binary image of a fruit profile with a stem, working out the skeleton of the region corresponding to the fruit, the presence of a stem produces a branch in this skeleton, as any other salient in the region.

On the other hand, because salients owing to stems are thinner than the region of the body, which corresponds to the fruit itself, using a thinning algorithm to extract the skeleton of the region, the part of the region corresponding to the salient will become skeletal before the rest of the body, that is, because thinning is performed iteratively to the region contour, the skeleton of stems and other salients will arise within the first iterations of the thinning.

Therefore, from a binary image of a profile view of a fruit (*Fig. 4a*), a thinning algorithm is performed, stopping after a given number of iterations. After this partial thinning, salients have become skeletons (*Fig. 4b*). To identify the branches corresponding to salients, the thinning performed is undone, adding layers of pixels only to the body pixels; thus, parts of the region which have become completely skeleton do not dilate, and they remain after the thinning is undone.

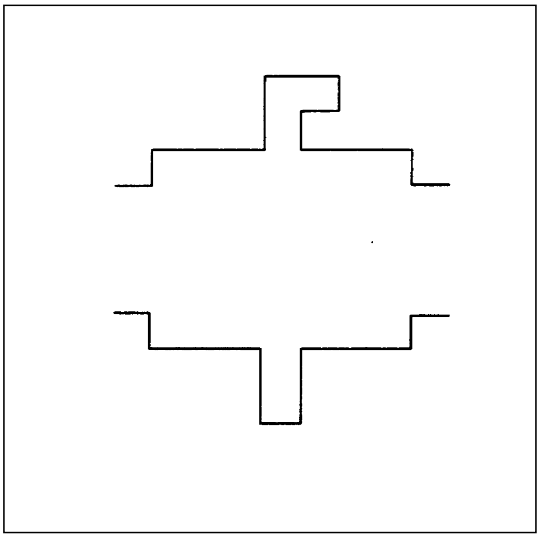

Fig. 2. Two types of small projections removed by the contour tracking algorithm, causing a smoothing effect over the contour

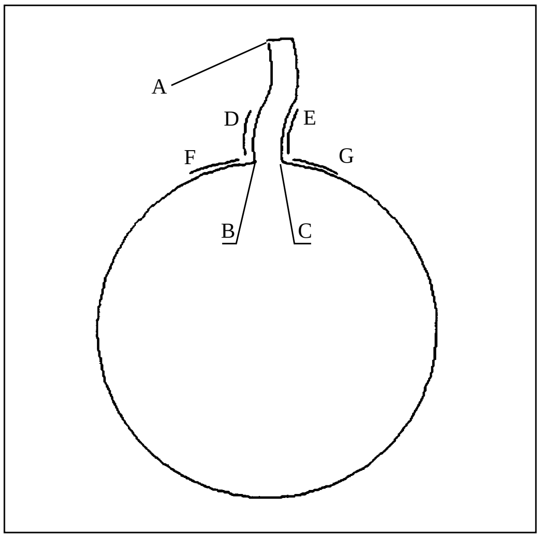

Fig. 3. Distances from the contour points to the centroid of the oranges change sharply on D and E, but are similar along F and G

Fig. 4c shows the two differentiated areas: the body of the fruit (grey) and the stem after location and reconstruction (black).

The procedure can be summarized as follows:

1. Threshold a fruit image to produce a binary image.
2. Perform a thinning, stopping after N iterations.
3. Undo the thinning, stopping after N iterations.
4. Identify salients looking for remaining skeleton branches.

During the N iterations of the partial thinning, the coordinates of points thinned in each iteration are saved, and used later to undo the thinning. A sequential thinning algorithm, developed by Kwok,[13] was used for the thinning process, and the connectivity properties of points defined for this thinning algorithm were used to label region pixels during the thinning, either as skeleton pixels or region body pixels. Number of iterations N is a constant set depending on the image resolution, to assure that after these iterations, any salient will be completely thinned.

To undo the thinning, pixels saved in the thinning step are considered, beginning with pixels of the last thinning iteration performed, and finishing with pixels saved during the first iteration. During each iteration of the undoing step, each pixel is added if and only if any of its eight adjacent neighbours is labelled as body pixel. If not, the pixel is definitely discarded. All added pixels during this step are labelled as region body pixels.

Skeleton branches resulting from this algorithm correspond with stems of other salients, for example leaves which could remain attached to the stem. Once these branches are located, the approximate point

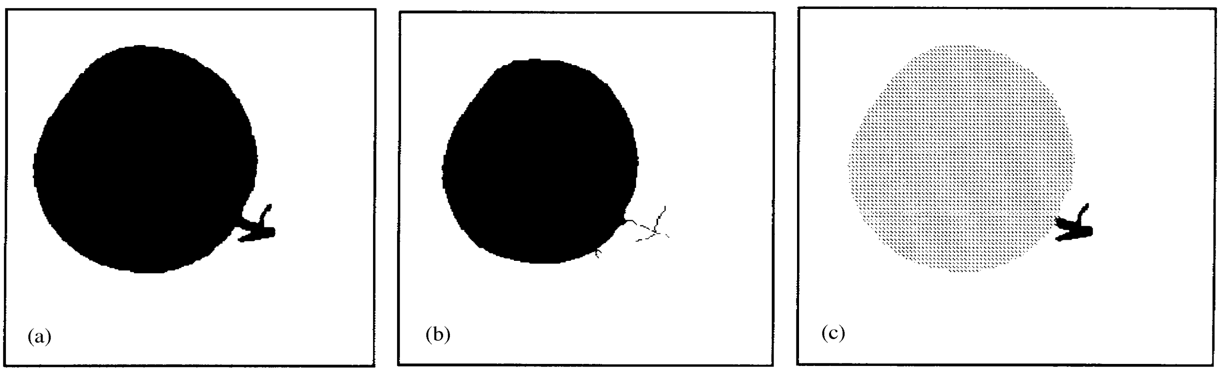

Fig. 4. Profile of an orange at three different phases of the thinning algorithm: (a) initial binary image, (b) after the thinning process, and (c) after undoing the thinning and reconstructing the stem (fruit body and stem are represented with different grey levels)

where the stem is attached to the fruit can be determined, and the stem length estimated by measuring the length of its corresponding skeleton branch.

4. Results

Discriminant functions produced with the colour segmentation procedure, based on discriminant analysis, were effective, dividing the colour space into the selected classes. *Fig. 5* shows the projection of some randomly chosen pixels from the testing set on a plane defined by two discriminant functions, which minimize the variability within the classes and maximize the variability between classes, with the assumptions of multivariate normal distributions for the red, green and blue variables, and equal covariance matrices of the four classes for the independent variables, as mentioned before. A good discrimination is achieved in this new working space, specially for classes *background* and *peel*, that form two perfectly separated clusters of pixels, while classes *stem–calyx* and *cut stem* are slightly overlapped.

Table 1 shows the classification matrix obtained after applying the colour segmentation model on the test set. The column under "% correct" represents percentages of correctly classified pixels for each of the four classes. Reading across each row, reveals how the classifier assigned pixels that actually belong to each of the classes. The diagonal from upper left to lower right gives correctly classified pixels, so that the sum of these values gives the total of correctly classified pixels. The bottom row of totals shows the numbers of pixels assigned to each class by the algorithm. The analysis of this table points to an excellent performance of the classifier for the background and the orange peel (100% and 99·97% of correctly classified pixels, respectively), a good performance for stem, calyx and leaves (93·37%), and acceptable for the *cut stem* (87·30%) class. It also shows that some pixels belonging to the *stem–calyx* class (most of them shadowed areas of leaves) are classified as background, and darkest pixels of the class *cut stem* are confused with the calyx.

The results of the stem location algorithm based on the segmentation process, for a total of 86 oranges tested, indicate 100% identification between destemmed and stemmed oranges, and image coordinates of the stem were correctly estimated in 90·3% of the cases, using a sample of 72 stemmed oranges. It is remarkable that most of the cases where the algorithm failed was for fruits with leaves appended to the stem.

The thinning algorithm showed 100% efficiency in distinguishing fruits with stem from a sample of 120 profile images, and 98·9% efficiency in locating the stem from 91 fruits tested.

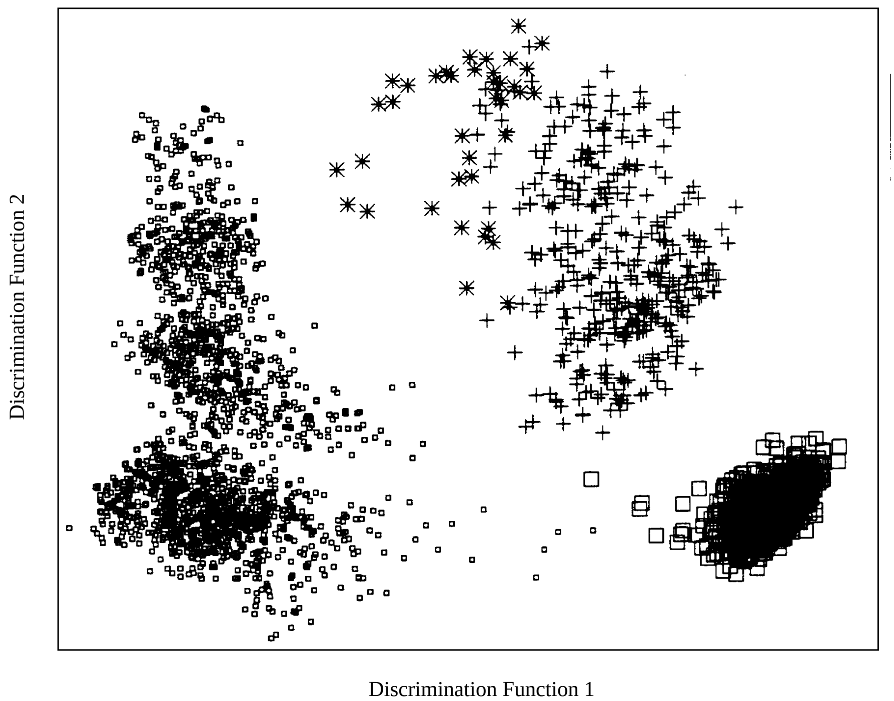

*Fig. 5. Projection of a group of pixels, randomly chosen from the testing set, on the plane defined by two discriminant functions. (□, background; ▫, peel; +, stem-calyx; *, cut stem)*

Table 1

Classification matrix of the segmentation process. (Rows:observedclassifications; columns: predicted classifications)

Class	*% Correct*	*Peel*	*Stem-calyx*	*Cut stem*	*Background*
Peel	99·97	21 395	1	0	5
Stem–calyx	93·37	6	7099	2	496
Cut stem	87·30	10	54	440	0
Background	100·00	0	0	0	20 919
Total	98·86	21 411	7154	442	21 420

Table 2 shows the results of the test for the contour tracking algorithm, for which 106 profile images were used. The rows indicate actual number of samples without stem and with stem, respectively, and the columns indicate estimations by the algorithm. The column under "stem" is divided into three columns that represent the results of location and measurement of the stem only for the stemmed samples. Only one orange with stem was misclassified, while all the de-stemmed oranges were well classified. From 79 fruits with stem, the stem was well located and measured in 88·6% of the cases, well located but incorrectly measured in 3·8% of cases, and failed to be located in 7·6% of cases. However, the performance was greatly improved when no leaves were attached to the stem. In all the cases, it was considered an error in the location of the stem when the deviation between the actual and the estimated insertion point was higher than 0·5 cm.

5. Conclusions

The colour segmentation procedure is a valuable method to distinguish between fruits with and without stem and to locate its insertion point on fruits randomly presented to the camera. The performance of the stem location algorithm decreases considerably for oranges with leaves and branches attached to the stem.

Contour tracking discriminates between stemmed and destemmed preoriented oranges, and is satisfactory to locate and measure the stem, though the performance is lower when leaves and branches are present.

The thinning algorithm detects successfully, with negligible mistakes, any salient from the fruit, working with profile images of preoriented oranges. This algorithm can also deal with most of the cases where leaves are appended from the stem.

The combination of these methods is useful to characterize the stem–calyx area of oranges, in the sense of presence or absence of stem and leaves, to locate the stem with a good accuracy, and to measure it in order to decide if it should be cut or not.

Fruit used for this research did not suffer serious damage or injuries but, for a destemming system before fruits reach packing houses, further research should be done in order to discriminate between the stem–calyx region and some external damage.

Table 2

Results of the contour tracking algorithm for detection, location and measurement of the stem. (Rows indicate observed classifications and columns algorithm estimations). Percentages results for location and measurement of the stem are shown in brackets

	No stem	*Stem*		
		Good location and measure	*Good location fail measure*	*Fail location*
No stem	27	–	–	–
Stem	1	70 (88·6%)	3 (3·8%)	5 (7·6%)

References

[1] **Chen P** Mechanical destemming of oranges. IV International Symposium on Fruit, Nut and Vegetable Production Engineering 1993, Valencia–Zaragoza, Spain

[2] **Wolfe R R; Sandler W E** An algorithm for stem detection using digital image analysis. Transactions of the ASAE 1985, 28: 641–644

[3] **Langers R A** Locating the growing-point of flower bulbs with computer vision. Proceedings of the Conference "Systemes de vision embarques pour l'agriculture et l'industrie agro-alimentaire" 1991, Montpellier, France

[4] **Howarth M S; Brandon J R; Searcy S W; Kehtarnavaz N** Estimation of tip shape for carrot classification by machine vision. Journal of Agricultural Engineering Research 1992, 53: 123–139

[5] **Yang Q** Finding stalk and calyx of apples using structured lighting. Computers and Electronics in Agriculture 1993, 8: 31–42

[6] **Juste F; Sevila F** Citrus: A European project to study the robotic harvesting of oranges. Proceedings of the 3rd

International Symposium on Fruit, Nut and Vegetable Harvesting Mechanization. 5–15 August 1991, Denmark–Sweden–Norway

[7] **Fornes I; Juste F; Santamarina C; Bimbo B** Potential damage to Citrus fruits in mechanical harvesting. IV International Symposium on Fruit, Nut and Vegetable Production Engineering 1993, Valencia–Zaragoza, Spain

[8] **Slaughter D C; Harrell R C** Discriminating Fruit for Robotic Harvesting Using Color in Natural Outdoor Scenes. Transactions of the ASAE 1989, 32: 757–763

[9] **Miller B K; Delwiche M J** A Color Vision System for Peach Grading. Transactions of the ASAE 1989, **32:** 1484–1490

[10] **Alchanatis V; Peleg K; Ziv M** Classification of tissue culture segments by colour machine. Journal of Agricultural Engineering Research 1993, **55:** 299–311

[11] **Paulsen M R; Liao K; Reid J F** Real-time detection of colour and surface defects of maize kernels using machine vision. AgEng 1992. Paper No.9206 17

[12] **Duda R O; Hart P E** Pattern classification and scene analysis, New York: John Wiley & Sons, Inc., 1973

[13] **Kwok P C K** A thinning algorithm by contour generation. Communications of the ACM 1988, 31(11)**:** 1314–1324