

Estimation of the Number of Apples in Color Images Recorded in Orchards

Oded Cohen¹, Raphael Linker¹, Amos Naor²

¹ Faculty of Civil and Environmental Engineering, Technion – Israel Institute of Technology, Haifa 32000, Israel

² Golan Research Institute, P. O. Box 97, Katzrin 12900, Israel
{oded@odedcohen.com, linkerr@tx.technion.ac.il, amosnaor@research.haifa}

Abstract. This work presents an algorithm for estimating the number of apples on trees using images acquired with a standard color CCD camera. The proposed system is capable of correctly identifying and localizing more than 85% of the apples in the images. To achieve this high detection rate, color and texture analyses are combined together with shape analysis. In the first step, pixels with a high probability of belonging to an "apple object" are detected according to their color and texture. In the second step, "seed areas" consisting of connected sets of pixels with a high probability of belonging to an apple object are detected. Each seed area is then extended to cover the entire visible area of the apple to which it belongs. Finally, each blob is segmented into simple components that can either be combined into circles or are discarded, so that each of the resulting circles corresponds to an apple

Keywords: artificial vision; image processing; yield estimation

1 Introduction

Early estimation of the future yield has always been a major challenge in agriculture. In orchards, such predictions are of interest to growers, packing and storage houses, and compensation funds. For growers, flowering intensity and yield forecast would be valuable at several stages of the season to optimize the intensity of chemical and hand thinning. Today, in the absence of accurate flowering and yield forecast, growers tend to perform insufficient thinning, preferring to be on the safe side rather than causing irreversibly low yield. This has significant repercussions later in the season when labor-consuming hand thinning is required to compensate for the insufficient chemical thinning. In addition, hand thinning is performed only after the process of natural thinning has ended and the fruits have reached a size where experienced growers can roughly estimate yield by visual inspection. This delay has two serious horticulturist consequences: (1) it results in smaller fruits at harvest, which, in addition to having a lower market value, require more time for harvesting, and (2) high crop load before the late thinning reduces flower bud initiation, which lowers the potential bloom in the subsequent season. For packing and storing houses, accurate yield forecast is needed toward the middle of the season. Such information would enable proper planning and adequate allocation of storage spaces. Finally, the compensation funds need a yield estimate early in the season when unexpected fruit drop or hail damage may occur.

Currently, yield forecasts are based on manual counting of the number of apples on selected trees. This method is extremely time-consuming and the small number of trees that can be inspected is insufficient due to the high variability of the yield that exists in apple orchards. More accurate forecast of flowering intensity and yield will increase the confidence of growers to perform early chemical thinning, which will minimize the needs for hand thinning, will increase the fruits' size and will increase potential flowering in subsequent seasons.

The present study focused on the development of a color-imaging system for automatic estimation of the number of apples present on a tree. The more general task of fruits localization on trees using artificial vision has been investigated in numerous studies, and most of these are summarized in the excellent review paper of Jimenez et al. [3]. These studies used either standard color cameras as in the present study (e.g. [1], [2], [5]), multispectral or hyperspectral imaging (e.g. [4],[6]), or thermal imaging (e.g. [7]). As noted in some of these studies, the task is especially challenging under natural illumination, and under such conditions shading and high contrast transitions are a main problem. For instance, sunlit fruits are ten times brighter than shadowed leaves, and sunlit leaves are

four times brighter than shadowed fruits [4]. When the fruits of interest are green as in the present study, the task is made more complex by the low color contrast between the background and the fruits.

2 Problem Description and Challenges

In ideal conditions, apples would appear as non-touching circular objects with smooth surface and distinctive color such as shown in Figure 1. In such a situation, finding and counting the apples within the image would be quite simple and would basically consist of finding blobs (clusters) of pixels within some predefined color range. However, such ideal situations are rarely found in images recorded in natural outdoor conditions. In such images, most apples are partially hidden by other apples and leaves, and their color is influenced by the environment, the light conditions and the photography parameters.



Fig. 1. Apples in ideal image

Light conditions have a major influence on detection feasibility. Figure 2 shows the two major effects of direct sunlight: (1) in some areas the CCD sensor is saturated and therefore any color information is lost, and (2) shading makes it impossible to recognize some of the apples as smooth surfaces. Having identified these negative factors, we tried to determine photography parameters that would eliminate or significantly reduce such problems. The best results were achieved by acquiring the images under diffusive sunlight conditions (pictures taken near sunset), and by manually underexposing the images compared to the camera's automatic setting (exposure reduction 0.7 units). Lowering the amount of light that reaches the sensor in such a manner brings the apples to the middle of the dynamic range of the sensor and avoids saturation.

Figure 3 shows a typical image obtained with such optimal photography parameters. Although to the human eye such an image may seem too dark and "flat", it is much more suitable for computer analysis than a bright and shiny image.



Fig. 2. Apples exposed to direct sunlight that causes saturation of the CCD and strong shadows

Visual inspection of the numerous images collected during this study led to identifying the following common situations and developing an algorithm that could handle them accordingly:



Fig. 3. Optimal image recorded under diffuse light (close to sunset) and with manually reduced exposure

2.1 Apple Clusters and Leaf Occultation

Most apples are partially hidden by other apples and/or by leaves (Figure 4)



Fig. 4. Detail of a typical image with partial occultation of some of the apples

2.2 Shading

The natural light source (sun) causes shades and non-uniform illumination. In particular, as shown in Figure 5, parts of an apple may appear much darker than the rest.



Fig. 5. Detail of an image showing non-uniform light distribution on the apples

2.3 Saturation

Sunlight can be reflected more intensely from some apples or from the apple surface regions that are perpendicular to the light source and this might cause local saturation. In such regions the red, green and blue components reach their maximum value (or close to it), and all color information is lost (Figure 6). Manual under-exposure of the images eliminates most, but not all, of these cases.



Fig. 6. Detail of an image with saturated regions in which all color information is lost

2.4 Color Inference

The color of an object is influenced by the light reflected from surrounding objects. In the present case, apples that are deeper within the tree will not only appear darker but will also be more similar in color to the leaves surrounding them (Figure 7).



Fig. 7. Detail of an image showing the large color differences that exist between the apples on the outside of the tree and those deeper within the tree

3 The Data

Images were recorded in a Golden Delicious orchard in the Matityahu Research Station located in Northern Israel. The images were taken during two consecutive seasons during the months of June-August, under natural daylight conditions.

A first set of images was taken using the fully automatic mode of operation of the camera (Fujifilm FinePix S8000fd). Eight of these images were selected as "calibration images", which were used to develop the algorithm, and another nine were selected as "validation images". About 70% of the apples present in the calibration images were marked manually to provide "apple" calibration pixels.

A second set of images was taken in the following season with a different camera (Olympus C740UZ) and new photographic parameters in an attempt to overcome some of the problems identified when analyzing the first set of images:

1. The camera shutter was set manually to 0.7 units lower than the automatic camera setting in order to darken all the objects and bring the apple pixels close to the middle of the dynamic range of the sensor.
2. The pictures were taken close to sunset, when there was nearly no direct sunlight and lighting was diffusive.

Nine of these images were used to calibrate the detection algorithm, which was then tested on another eight images.

4 Algorithm Description

The algorithm includes three main steps which are described below. The algorithm was developed in the **ImagingChef** environment (a full description of this imaging development environment can be found at: www.odedcohen.com), which supported the entire development process, from reference object marking to algorithm development and result visualization

4.1 Step 1 - Apple pixels detection by color and texture

In the first step, pixels with high likelihood of belonging to an "apple object" are detected according to their color and texture. To achieve this, a K-nearest-neighbors (KNN) classifier was built using pixels belonging to "apple" and "non-apple" objects that were manually marked in the images belonging to the calibration set. Objects overlapping the apples, such as leaves, branches and the petal area were also marked in order to exclude them from the study process. Figure 8 shows some marked objects on part of a typical image. Although this is not visible in

the image, the properties defined for each object include (in addition to the object type) the relative depth of the object, so that the analysis takes into account partial occultations.



Fig. 8. Detail of an image with objects marked using the ImagingChef software. In addition to the apples (indicated by circles), regions that should be ignored at the calibration stage (such as apples that are barely visible or parts of apples that are occulted by leaves and branches) are marked by rectangles or ellipses.

Each pixel in the calibration set was characterized by its Red-Green-Blue (RGB) components and its smoothness, which was estimated as the variance of the RGB components on a small area (6-by-6 pixels).

An example of the classification result is shown in Figure 9, in which the lighter pixels have a higher probability of belonging to an apple object.




Fig. 9. (Left) Original image and (Right) image showing the classification results, in which the lighter pixels have a higher probability to belong to an apple. The yellow and pink arrows point to typical incorrect classifications as explained in the text

Two typical problems are visible in Figure 9:

1. Saturated (bright) areas get low probability, as their color is practically white (e.g. area indicated by yellow arrow);
2. Pixels that belong to the bottom side of leaves are incorrectly recognized as "most probably apple" since this side of the leaves is lighter and its color is similar to that of apples (e.g. area indicated by pink arrow). In order to minimize this problem, the camera should be located high enough and pointing downward so that only the top side of the leaves is visible.

4.2 Step 2 - Apple surface detection by seed area detection and growth

In this step connected sets (blobs) of apple pixels are detected and extended to cover the area of the apple to which

they belong. Each “seed area” consists of a connected set of pixels that have a high probability of belonging to an apple object. Ideally, each apple should result in one seed area that should cover most, if not all, of the apple. In practice, parts of the apple surface might be considerably darker or lighter, because of the amount of light that reaches it and of the way this light is reflected. Such areas are misclassified as "non-apple", resulting in a smaller seed area. Also, some apples appear as split objects due to partial occultation by branches or leaves, which results in several seed areas for the same apple (Figure 10).



Fig. 10. Detail of an image that shows that a single apple may contain more than one seed region (purple contour) as a result of partial occultation

Figure 11 shows typical results of the seed areas detection procedure. It can be seen that both saturated and darker pixels of the apple surfaces are not included in the seed areas.



Fig. 11. (Left) Typical image. (Right) Detected "seed areas" (in orange)

After detecting the seed areas, these areas are expanded to contain similar neighboring pixels. This compensates for the misclassification of the very bright and very dark pixels described above. The expansion is performed by extending each blob to include neighboring pixels with low variance. Figure 12a shows a variance map of Figure 11a in which the darker pixels indicate a higher variance while smooth surfaces appear as lighter regions. Figure 12b shows the results of the seed area expansion: both saturated and dark areas around the seed areas are now included in the "apple" area.



Fig. 12. (Left) Variance map of Picture 11a. (Right) Seed areas obtained by expanding the areas shown in Picture 11b using the variance map shown in Picture 12a

4.3 Step 3 - Contour segmentation and apple detection

A seed area may correspond to any of the following situations (Figure 13):

- a. A fully visible apple.
- b. Part of an apple. Other parts of the same apple might be associated with other seed areas.
- c. Several apples merged into one blob.
- d. Apple and a leaf merged into one blob.
- e. Misclassified non-apple pixels, most usually leaf under intense light.




Fig. 13. Typical seed areas

Transforming the blobs into apples requires shape analysis. For this, the contour of each blob is segmented into the following components (Figure 14):

- arcs,
- linear segments,
- anamorphic segments.

In the next step, the arcs are grouped into circles and at the end of this stage each circle indicates one apple.



Fig. 14. Contours obtained after segmenting the contour of each seed region into arcs, linear segments and anamorphic segments

5 Results

5.1 First set of images – images recorded in automatic exposure mode

The results are summarized in Table 1, which shows the number of correctly detected apples and the number of false positive detection (incorrect detection of apple where there is none). These results are compared to a number of apples visible to a human observer inspecting the images for about 20-30 seconds. The overall rate of correct detection is above 85% for both the calibration and validation images while the false positive rate is 12% and 19%, respectively.

Table 1: Results of the analysis of the images recorded using automatic light exposure

	Calibration images		
	Number of apples visible	Number of apples detected	Number of false detections
Image 1	31	26	2
Image 2	29	25	2
Image 3	45	34	1
Image 4	37	33	8
Image 5	22	17	4
Image 6	31	26	3
Image 7	19	19	2
Image 8	49	44	9
Total	263	224 (85%)	31 (12%)
	Validation images		
Image 1	44	44	6
Image 2	26	23	4
Image 3	40	38	12
Image 4	19	17	9
Image 5	22	21	5
Image 6	60	53	4
Image 7	47	37	10
Image 8	38	33	0
Image 9	44	34	11
Total	296	256 (86%)	55 (19%)

A detailed analysis of the results led to the following conclusions:

- As apples are much lighter than leaves, automatic setting of the exposure by the camera resulted in apple pixel values too close to the end of the dynamic range of the sensor, which often caused some of the pixels to be saturated.
- The light was not diffusive enough. Even though the pictures were taken early in the morning, the light distribution was still not uniform enough to enable better detection.

5.2 Second set of images – images recorded with manually-lowered light exposure

Adjusting the photography parameters as detailed in Section 3 improved the pictures quality for subsequent analysis. There were no cases of color saturation and the apples appeared as more uniform. As a result, the overall rate of correct detection increased to 90% and the false positive rate decreased to less than 10% (Table 2).

Table 2: Results of the analysis of the images recorded under diffuse light and after setting manually the camera shutter 0.7 units lower than prescribed by the automatic setting

	Calibration images		
	Number of apples visible	Number of apples detected	Number of false detections
Image 1	66	61	4
Image 2	68	61	7
Image 3	29	28	7
Image 4	28	25	5
Image 5	37	33	2
Image 6	13	13	3
Image 7	36	33	3
Image 8	58	53	7
Image 9	29	28	1
Total	364	335 (92%)	39 (11%)
	Validation images		
Image 1	57	47	12
Image 2	48	39	2
Image 3	60	52	4
Image 4	17	14	2
Image 5	31	29	2
Image 6	24	17	1
Image 7	53	48	1
Image 8	62	57	7
Total	295	256 (87%)	19 (6%)

6 Conclusions

Apple detection in images taken under natural daylight conditions has two main inherent difficulties: (1) the natural light is not diffusive enough and might cause shades and saturation, and (2) the apples have different shapes and colors, overlap other objects, and are rarely fully visible. Nonetheless, the present algorithm is capable of detecting correctly more than 85% of the apples present in an image obtained using the automatic settings of a standard color camera. The performance can be improved by taking care of recording the images under diffuse

light conditions and manually lowering the exposure in order to avoid light saturation. In such cases close to 90% of the apples are correctly detected. The proposed algorithm should be further validated using a larger dataset and its extension to other apple varieties should be considered.

References

1. Annamalai, P., Lee, W.S. Citrus yield mapping system using machine vision, ASAE paper 03-1002, 2003.
2. Bulanon, D. M., Kataoka, T., Zhang, S. Ota, Y. Hiroma, T. Optimal thresholding for the automatic recognition of apple fruits, ASAE paper 01-3133 (2001).
3. Jimenez, A.R., Ceres, R., Pons, J.L. A survey of computer vision methods for locating fruit on trees, Transactions of the ASAE 43: 1991-1920 (2000).
4. Kim, Y., Reid, J. Apple Yield Mapping Using a Multispectral Imaging Sensor. Proceedings of the AgEng'04 Conference (2004).
5. Pla, F., Juste, F., Ferri, F., Vicens, M. Colour segmentation based on a light reflection model to locate citrus fruits for robotic harvesting. Computers and Electronics in Agriculture 9: 53-70 (1993).
6. Safren, O., Alchanatis, V., Ostrovsky, V., Levi, O. Detection of green apples in hyperspectral images of apple-tree foliage using machine vision. Transactions of the ASABE 50: 2303-2313 (2007).
7. Stajnko, D., Lakota, M. Hocevar, M. Estimation of number and diameter of apple fruits in an orchard during the growing season by thermal imaging. Computers and Electronics in Agriculture 42: 31-42 (2004).