Automatic Can Separation

Kjersti Aas Line Eikvil Otto Milvang

Norwegian Computing Center, P.O. Box 114 Blindern, N-0314 Oslo, Norway
e-mail: Kjersti.Aas@nr.no Tel: (+47) 22 85 25 00 Fax: (+47) 22 69 76 60

Abstract

A system for analysing can lids is presented. Images of the can lids are captured by a CCD camera. The images are then analysed to determined whether the cans are with or without imprinted text on the lid. The system is also capable of detecting whether the can is viewed from the top or the bottom. The performance of the system is promising, yet there are possibilities for improvements.

1 Introduction

Tomra Systems AS produces machines for receiving return cans for the American market. Here, a deposit is payed on the cans in some states, while there is no deposit in other states. Obviously, money should only be reimbursed for those cans for which a deposit has been payed. To avoid reimbursing money for a non-deposit can from one state in another state, it is therefore necessary to be able to distinguish between cans purchased in states with and without a deposit arrangement.

Today, the can machines are equipped with a bar code reader, which reads the deposit information about the product contained in a bar code printed on the cans. An alternative solution, which may be cheaper, is to imprint this information on the can lids. The issue of the work described in this paper, was to design a machine vision system for interpretation of such information.

The lids of the cans are made of aluminium, resulting in a very shiny surface. This creates reflexes and low contrast making the problem of recognition difficult. In [1] stamped characters are recognized by analysing range images. In the can machine this is, for economical reasons, not an option, and a CCD camera had to be used for the image acquisition. Hence, overcoming the problems of low contrast combined with sharp reflexes was a challenge.

To obtain separation of cans with and without refund, several problems must be solved.

- The system should distinguish between the lid and the bottom of a can. This is necessary as cans may be inserted the wrong way into the can machine.
- The system should separate cans with information on the lid from those without, as cans without refund may often lack information on the lid.
- The final step would be to interprete the information on the lid to decide on the amount to be refunded, using methods from symbol recognition.

In this study[1] we have concentrated on solving the first two problems, and on finding robust methods which will work in a final system. Also, being able to solve these problems are absolutely crucial for the complete system to work. For interpretation of the information on the lids, a standardization of the information is necessary. Currently this information may appear in any form, and the recognition problem was therefore not handled at this stage.

2 Data

The available space for a camera within the can machine is limited, and it is therefore impossible to obtain images viewing the cans from directly above. The camera must be placed at an angle, and the resulting view of the can is illustrated in Figure 1. A mirror is used to reflect the image of the can back to the position where there is room for the camera.

The image acquisition is performed by a CCD camera. The images used in this study has a resolution of

[1]The project has been supported by the Norwegian Research Council

512×512 pixels, and for the ellipse describing the top of the can, the minor and major axes are typically 110 and 190 pixels, respectively. The cans are illuminated by red light emitting diodes. A few examples of the can images in our dataset are shown in Figure 2.

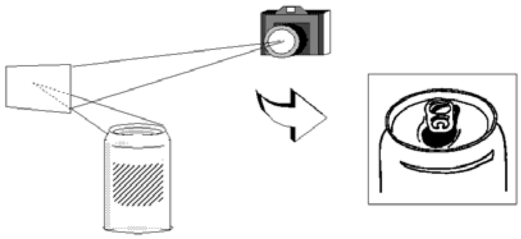

Figure 1: *View of the can.*

3 Methods

Prior to any analysis, the can must be located in the image and segmented from the background. The lid or the bottom of the can is described by an ellipse, and it is only the information within this ellipse which is needed for the analysis. Hence, the process starts by determining this ellipse, and this is described in Section 3.1.

When the ellipse is estimated, the next step is to determine whether the can is viewed from the top or the bottom. This is decided by searching the estimated ellipse for the hole which only will be present in the lids. Section 3.2 describes the hole detection.

The lids are then analysed to identify the regions of the lid which may contain information. This process, which is described in Secton 3.3, results in two regions, one on each side of the hole in the lid. One of these regions may be covered by the pull tab of the lid. If this is the case, symbols printed in this area will not be visible. Hence, the state-specific information should always be printed on both halves of the lid. To extract the region which is not covered by the pull tab and to decide whether this region contains any information, we use an edge detector. This process is described in Section 3.4.

3.1 Ellipse estimation

For the estimation of the ellipse defining the can top, two-level thresholding [2] is used to segment the can from the background. From the thresholded image shown in Figure 3, a part of the contour of the top of the can is extracted. Based on this contour segment, the ellipse describing the lid (or bottom, if the can has been inserted the wrong way into the can machine) is estimated.

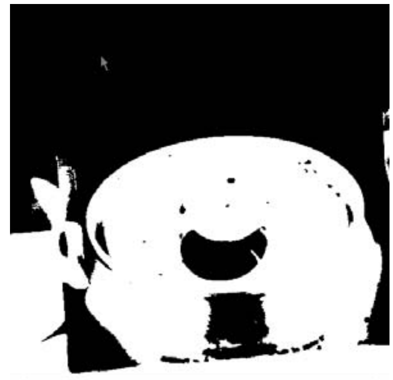

Figure 3: *Principles for hole detection.*

The problem of fitting conic sections to scattered data has arisen in several applications described in the literature. We have used a method described by Bookstein in [3], where the sum-of-squares of the quadratic form

$$Q(x,y) = A\,x^2 + B\,x\,y + C\,y^2 + D\,x + E\,y + F = 0 \quad (1)$$

is minimized. The ellipse fitted to the contour extracted from the can in Figure 3 is shown in Figure 4. Having estimated the ellipse, a sub-image is determined as the smallest possible rectangle enclosing the ellipse, and all further processing takes place on this sub-image.

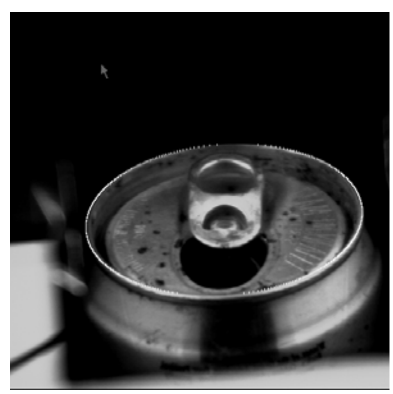

Figure 4: *Estimated ellipse marked with white.*

3.2 Hole detection

For the hole detection, the sub-image is first thresholded by a single threshold chosen such that 20% of the pixels will be below the threshold (black) and 80% of the pixels will be above (white).

Figure 2: *Examples of the can images in our data set. From left to right: Can lid with text, can lid without text, can bottom.*

From this thresholded image, all the contiguous regions of pixels with the same pixel value are extracted. The hole will be represented as a black region, and we therefore investigate the regions containing black pixels. We assume that a hole covers between 4 and 14 percent of the area of the ellipse describing the can lid. For a black region to be a hole its size must lie between these limits. We also put some requirements on the shape of the hole.

For the regions satisfying these criteria, the centroid is computed. The centroid has to be within a certain distance (specified as an ellipse) from the centre of the can. If more than one region satisfies these requirements, the one with the greatest area is chosen. If no region satisfies these requirements, we assume that the can is viewed from the bottom.

Figure 5 shows the thresholded image of a can from which the hole is found. The centre of the can lid and the centroid of the identified hole are marked with '+'.

Figure 5: *Principles for hole detection.*

3.3 Determination of regions of interest

If no hole has been detected in the previous step, the can is classified as a 'bottom'. Otherwise, we assume that the can is viewed from the top and we proceed to determine the regions which may contain text.

The text on a lid will always be printed within the two sectors on each side of the hole (Figure 6a). For the further analysis of the lid, we are therefore interested in looking at these two sectors, only. To determine the two sectors, we define an inner ellipse enclosing the hole, an outer ellipse enclosing the flat area of the lid, and two straight lines intersecting at the centre of the can. The lines divide the lid into four sections of which neither the one covering the hole nor the opposite will contain any text. The principles are illustrated in Figure 6b.

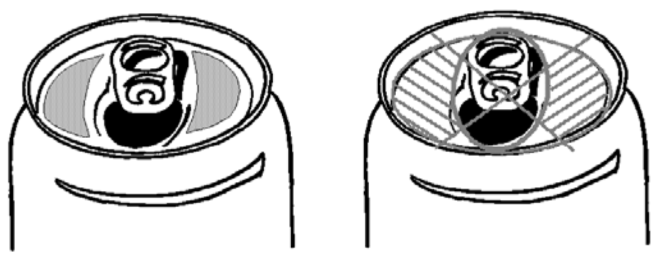

Figure 6: ***a:*** *Text on the can lid is only printed within the areas marked with shading.* ***b:*** *We delimit the interesting sectors by an inner and an outer ellipse and two sectors the size of the hole.*

First, the angle from the centre of the can to the centroid of the hole is computed. This angle determines the rotation of the inner ellipse. The dimensions of the inner ellipse are chosen such that the ellipse approximately covers the hole. The outer ellipse will have the same orientation as the initially estimated ellipse.

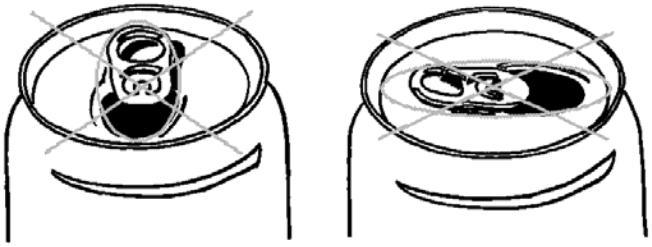

Figure 7: ***a:*** *Direction of hole giving the widest possible sector.* ***b:*** *Direction of hole giving the smallest possible sector.*

The size of the inner ellipse and the width of the sector covering the hole (and its opposite) are defined to vary with the location of the hole. When the can is rotated as in Figure 7a, the hole is at its widest, and when the can is rotated 90 degrees from this position like shown in Figure 7b, the hole is at its most narrow. This means that we define the widest sector when the direction from the centre to the hole is parallell to the y-axis and then the width continuously decreases until this direction is parallell to the x-axis.

The masks defining the resulting sectors for the can from Figure 4 are shown in Figure 8. In Figure 9 the contents of the sectors are shown. In the further processing only the contents of the identified sectors are analysed.

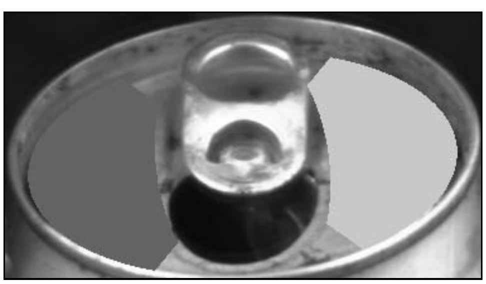

Figure 8: *Masks defining search sectors.*

3.4 Edge detection

The pull tab of the lid may cover one of the regions determined in the previous section. As information printed in this region may not be visible, the next step is to identify the region which is not covered by the pull tab. This is obtained through edge detection. We use the Sobel operator [4], which has the advantage of providing both a differencing and a smoothing effect. The smoothing is useful as derivatives generally enhance noise.

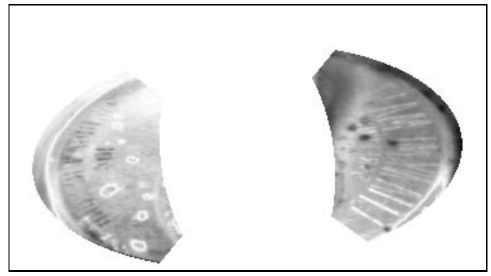

Figure 9: *The contents of the search sectors.*

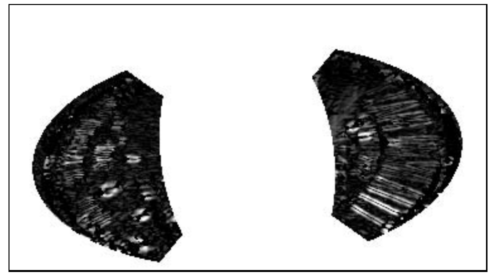

Figure 10: *The gradient information of the search sectors.*

In addition to the selection of an appropriate region, the purpose of the edge detection is to determine whether this region contains deposite symbols or not. The effect of shade and reflected light may result in sharp edges caused by the outer edge of the lid. Hence, before making any decisions based on the gradient magnitude image, these edges must be separated from edges caused by text. Since the non-text edges will follow the elliptic shape of the lid, an edge pixel should be removed if the vector from the centre of the lid to this pixel is perpendicular to the direction of the pixel. The result of using the gradient direction information to remove edges from the gradient magnitude image is shown in Figure 10.

The uncovered region is identified by computing the mean edge value for both sectors and finding the sector with the lowest mean value. Cans with a mean value below a predefined limit are assumed to have no text imprinted.

4 Experimental Results

In the testing of the system, the camera within the can machine was connected to a PC on which the software was implemented. During the test situation, the images were written to file for later analysis of the results.

Our test set consisted of 40 different images of cans. In 35 of the images, the cans were viewed from the top, while the remaining 5 images were of can bottoms. 27 of the can tops had text imprinted. The cans were rotated at different angles, with the pull tab located at different positions. On this set of images we tested the method for separation between can tops and can bottoms and the method for separation between cans with and without text.

The 5 images of can bottoms were all correctly classified. One of the can tops was classified to 'bottom'. The outline of this can was missing due to wrong adjustment of the mirror and camera in the setup. Hence, the extracted contour used for the ellipse estimation was incorrect leading to an erroneous ellipse and subsequent misclassification.

The 8 cans without text were all correctly classified, while five of the cans with text were misclassified. One error resulted from wrong adjustment of the mirror and camera in the setup, while the remaining four errors were due to low contrast in the images, making the text almost invisible.

The estimation of the search sectors was very accurate for all images except for those where the outline of the can was missing.

5 Summary and Discussion

The aim of this study was to arrive at robust methods which would distinguish the bottom from the lid of can, separate cans with imprinted information on the lid from those without and locate the region of the lid where imprinted information was visible (not covered by the pull tab). Recognition of the symbols imprinted on the lids were not an issue at this stage, as this information is not yet standardized. Testing of the system was performed by connecting the CCD camera within the can machine to a PC on which the methods for analysis were implemented.

The analysis starts by determining the ellipse describing the top (or the bottom) of the can. Based on this ellipse a sub-image is selected for further processing. A search for the hole in the can lid determines whether the can is viewed from the top or the bottom. In the former case, the location of the hole is used to define the regions which should be searched for text or graphics. To determine whether the lid has text imprinted, the gradient information in these regions is used. The methods used for the analysis are all fairly fast. This was necessary for them to be suitable for a real time application.

Experimental results show that the method distinguishes very well between can tops and can bottoms. The separation of cans with and without text is slightly harder, but the preliminary results are good. Out of 35 images of can lids, 28 are correctly classified. Two of these errors occured because the adjustment of the mirror and camera was inaccurate, resulting in an image where a bit of the outer contour of the can lid was missing. The current method used in the analysis is highly dependent on that the entire upper contour of the can lid is visible in the image resulting from the CCD camera. It is therefore important to be careful with the camera setup.

The remaining four misclassifications were due to low contrast in the images. As the pull tab of the can may cover one of the sectors analysed for text, it is important that the text is visible in both sectors. To avoid disturbing reflexes from the bright aluminium can, diffuse lighting was used. However, it might seem that not only the disturbing reflexes were removed, but also too much of the light reflected from the imprinted text. Still, the light setting has not been optimized for the current method and there should be room for improvements.

Concluding this study, the preliminary results show that the method enables separation between can lids and can bottoms, distinguishes between cans with and without text and is able to extract the region where the imprinted information is visible.

References

[1] T. Houriuchi, K. Toraichi, H. Yamada, K. Yamamoto, *Stamped Character Recognition Method using Range Images*, Proceedings of the 2nd ICDAR, October, 1993.

[2] S. S. Reddi, S. F. Rudin and H. R. Keshavan, *An Optimal Multiple Threshold Scheme for Image Segmentation*, IEEE Transactions on Systems, Man and Cybernetics, Vol SMC-14, pp 661–665, 1984.

[3] F. L. Bookstein, *Fitting Conic Sections to Scattered Data*, Computer Graphics and Image Processing Vol. 9, pp 56–71, 1979.

[4] R. C. Gonzalez, R. E. Woods, *Digital Image Processing*, Addison-Wesley, New York, 1992.