SEGMENTATION OF TOUCHING OBJECTS

Donald G. Bailey
Image Analysis Unit
Massey University
Palmerston North

ABSTRACT

Image analysis is frequently used to measure the distribution of size, or some other parameter, of a large number of objects. If the objects are separated, segmentation consists of separating the objects from the background. When the objects are touching, separating objects from the background will result in clusters of objects which then need to be separated into individual objects.

If the size and shape of the objects is known (for example in counting applications), template matching may be used to isolate the individual objects. Problems are encountered in determining the orientation of non-circular objects. For circular objects (of varying sizes) several methods may be used. A Hough transform may be used to transform each boundary pixel into parameter space (x, y, and radius) where clusters corresponding to individual objects may be detected. Arc fitting and segment matching is a refinement of this technique which removes the need for explicitly representing the 3-d parameter space. A further technique is to analyse the concavities and successively split clusters along narrow points.

These techniques are reviewed and compared in the size analysis of fibres.

INTRODUCTION

Image analysis is the use of image processing techniques to obtain quantitative information about an object or group of objects. The two most significant benefits of using image analysis for automatic measurement are the ability to obtain more consistent results (since they are less dependent on operator variability and fatigue) and a reduction in the measurement time (enabling larger sample sizes to be used).

In some applications, all that is required is an average measurement. In such cases, it may not be necessary to explicitly measure each individual object, provided that the number of objects is known (or can be determined). However, if a population distribution is required then each object must be measured individually, regardless of the parameters of each object that are being measured. This requires that each object must be detected individually to be measured. When there is a large number of objects in the field of view, it may be difficult or impractical to prevent adjacent objects from touching. In some circumstances, the preprocessing stages may cause close objects to become joined in the image. In either case, the segmentation process may result in clusters of objects, rather than detecting the objects singly. Therefore, to enable objects to be measured individually, it is necessary to introduce an object separation stage into the algorithm (figure 1). Once separated, the required parameters of each object can readily be measured.

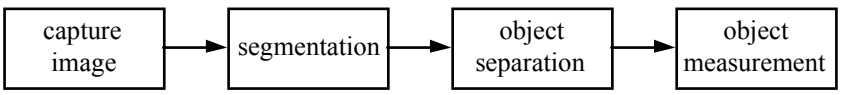

Figure 1: Processing steps required when measuring objects which may be touching.

Depending on the shape of the objects, a variety of techniques is available to separate the objects. These are described in the following sections.

Two assumptions are made that need to be stated explicitly. The first is that any objects that touch do not overlap significantly. Although some of the following techniques will work on overlapping objects, some of the information about the object may be distorted, or even lost (for example optical density in microscope images). Secondly, it is assumed that the image is binary. In other words the clusters of objects have been separated from the background.

TEMPLATE MATCHING

The first approach is to look for a match between the input image and a template representing the object to be detected. It may be implemented in a number of ways - the simplest of which is to scan the template across the image looking for matches. In a sense, this is a form of local filtering, where the template is used as filter kernel. If a linear filter is used, this process is matched filtering, or correlation.

The template may consist of either the boundary of the object, or the full shape. If boundaries are used, edges need to be detected in the input image. Such a method is reasonably insensitive to small gaps in the detected edges provided that most of the boundary is found.

Template matching works well if the objects being separated are the same size, shape and orientation. To detect a range of different shapes or orientations, a set of templates is required. In general, simple template matching techniques are too restrictive in most applications involving natural objects [1] since they do not handle the changes in the size or orientation of the object very well. However the technique may be generalised by using a form of constrained region growing [1] or through the introduction of additional parameters, giving a form of Hough transform.

HOUGH TRANSFORM

The Hough transform turns template matching around so that instead of scanning the template across the image, detected edge points 'vote' for possible objects that they could be a part of. The Hough transform works best if the objects being detected can be represented by a few parameters [2]. For example, a circular object (of arbitrary size) may be represented by three parameters: the horizontal and vertical position of the centre, and the radius.

The object detection (and separation) process consists of assigning each pixel in an image to an individual object or the background. Or, represented in edge terms, associating each edge point with the edge of an individual object. Since each object is represented by a single set of parameters, this means finding a set of parameters that matches a significant number of edge points in the input image. The Hough transform does this by transforming the input edge points into parameter space. Each edge point in the input image is transformed into parameter space by contributing to all possible parameter combinations describing an object that it could be a part of.

This process is perhaps best illustrated by example. If we are detecting circular objects, the parameter space is three-dimensional, with parameters **x**, **y**, and **r** (figure 2a). If we consider a single edge point **P** in the input image, some of the objects that it could be part of are shown in figure 2b. Looking at it a little more systematically, if the point is on the edge of a circle of radius **r**, the centre of the circle must be a distance **r** from **P**. Therefore, the locus of the **x** and **y** parameters for possible circles of radius **r** through **P** fall on a circle of radius **r** centred at **P** (shown as the shaded circle in figure 2c). This is just for a single radius (**r** constant). Extending to three dimensions by making **r** a parameter gives a cone with axis centred on **P** (figure 2d). Therefore, since **P** is on a circle, that circle must have parameters falling

somewhere on the surface of the cone in parameter space. Each edge point in the input image transforms into a cone in parameter space.

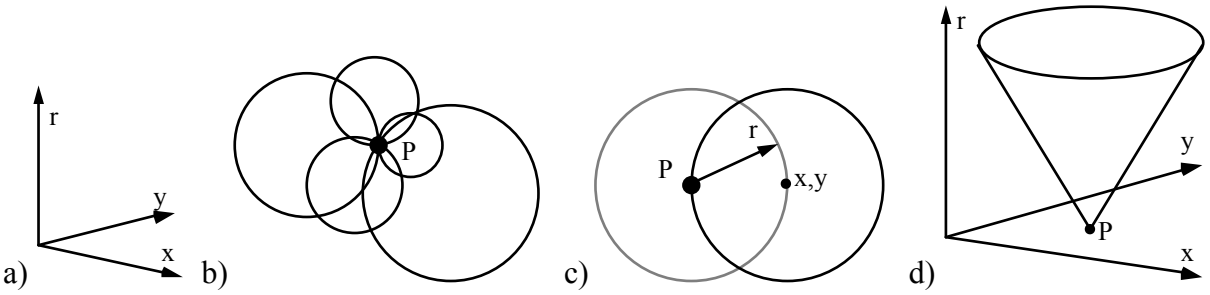

*Figure 2: a) parameter space for detecting circles; b) some of the possible circles through P; c) circles with radius **r** must have centres at radius **r** (shown shaded); d) parameters of all possible circles through **P**.*

The points on the surface of the cone may be considered as parameters of candidate circles through the associated point in the input image. When the contributions from every input edge point are considered, there will be regions in Hough parameter space which fall on multiple cones. Such clusters of points in parameter space represent the likely parameters of individual objects.

To show how the Hough transform may separate touching circles, consider figure 3. When each edge point in the leftmost figure is transformed, each point contributes to a cone shaped region in the Hough parameter space. Since the three dimensional nature of the parameter space makes it difficult to draw on paper, a series of cross sections is shown for different radii. Strong clusters in parameter space (indicated by dark spots in these images) result when the corresponding parameters represent candidate objects common to a large number of edge points. Therefore, these represent the most likely parameters of the individual objects. By detecting such clusters, the parameter values may be used to redraw the detected objects, as shown in the rightmost panel.

*Figure 3: touching circles (left); cross sections of Hough parameter space at radii 4, 6, 8, 10, and 12 pixels - note the sharp peaks at **r** = 6 and 10; detected objects (right).*

If the object is not symmetrical, a further parameter **Ø** may be added for the object orientation. For example, to detect elliptical objects, two further parameters need to be added: one for the orientation **Ø**, and one for the ellipticity **e**. The difficulty with the Hough transform is that searching for clusters in parameter space is not always a trivial problem, especially when there are several parameters. The process of finding all parameters that may be used for each input point is computationally expensive, especially when the parameter space has more than two dimensions. This computational burden may be reduced if the expected parameter values fall within a narrow range.

As can be seen in figure 3, there is quite a lot of background clutter in the parameter space. This may be reduced by using knowledge of the object to simplify the problem. For example, when detecting the circles, we know which is the inside, and which is the outside, so an edge point only needs to contribute to parameters in the direction of the inside. We also know that the direction of the centre of the circle is perpendicular to the edge. So rather than contribute to parameter space in a cone, only a single ray in the direction of the centre is needed. When the contributions from each point are combined, the rays intersect at a point, indicating the required parameters.

*Figure 4: touching circles (left); cross sections of Hough parameter space using the ray method at radii 4, 6, 8, 10, and 12 pixels - note the peaks at **r** = 6 and 10; detected objects (right).*

Figure 4 shows the results of applying the ray method to the touching circles. The Hough parameter space is considerably less cluttered as a result of removing the extraneous information. However, the disadvantage is that the peaks are not as sharp. This results from the limited accuracy in the estimation of the slope at each edge point because of the discrete positions of pixels in a digital image. In this example, the slope was estimated by smoothing over 4 edge points (2 on either side). Ideally, the smoothing should increase as the radius increases to maintain reasonable accuracy.

In figure 4, the Hough parameter space still had three dimensions. Since a lot of the clutter has been removed, it is now possible to compress the parameter space to two dimensions. Rather than drawing the rays in three dimensions, by drawing the rays in two dimensions (ignoring the radius), the peaks will indicate the centre of the circles. In this example, since it is known beforehand that the radius is less than 20 pixels, each edge point contributes to Hough space up to a distance of only 20 pixels. This helps to reduce the clutter in other, unrelated regions of the parameter space. The result of using this method is shown in figure 5. This time, however, there is no information on the radius. Having found the object centres, determining the radii is relatively straight forward.

a) b) c)

Figure 5: a) touching circles; b) the result of performing a Hough transform with only 2 parameters; c) expanded to show the fainter background detail.

The techniques described above can be generalised to arbitrary shaped objects [3]. The objects are represented by a template, which is then parameterised by position (**x** and **y**), and maybe scale (**r**) and orientation (**Ø**). In this way, the Hough transform implements a form of template matching.

CURVE FITTING

The refinements illustrated in figures 4 and 5 may be taken one step further. Rather than transforming each edge point individually, by transforming groups of points as a single unit, it is possible to reduce the clutter in Hough parameter space even further. If the objects are circular, or approximately so, then a circular arc may be fitted to a group of edge points. The parameters of the fitted curve then represent the candidate object. By using only a small number of groups, each containing many edge points, it is possible to do away with the Hough parameter space altogether. If this is done, each group produces a candidate object, and groups producing similar parameter values are combined.

The question then becomes one of deciding how the edge points should be grouped. Adjacent edge points are likely to belong to the same object, however this will not be the case at the boundary of touching objects. If the objects are circular, or elliptical, then edge points with convex regions of the boundary are likely to belong to the same object. Edge points in the concave regions are not used since it is not known which object such points belong to. After

combining groups where necessary, the edge points in the concave regions can be assigned to the closest group. This process is illustrated in figure 6.

a) b) c)

Figure 6: a) touching objects; b) the convex regions on the boundary; c) the resulting segmentation

One problem with this approach is the accuracy of the parameters of the fitted curve. The digitisation noise (caused by having edge points on a rectangular grid) can sometimes affect the parameter values quite significantly. This effect is even more pronounced when there are only a few points in a group. Groups belonging to the same object may end up with quite different parameter values making it difficult to decide whether or not the two groups should be combined.

This last difficulty may be overcome by considering adjacent pairs of groups in the following manner. First, the most complete group is found; that is the group that covers the largest proportion of its associated candidate object. It is then combined with the group on either side on a trial basis. If one of the combinations gives a good fit (the error is below a certain threshold) then the combination is kept, otherwise the group is removed. This process is repeated until there are no more groups left. With a chain of objects, this process removes one object from the end of the chain with each iteration. This procedure has been used with collagen fibrils [4] to separate touching fibrils after segmentation from the background.

CONCAVITY ANALYSIS

When the objects are approximately circular or elliptical, when two (or more) objects touch, there is a distinct concavity in the outline, provided that the objects are of similar size (to within an order of magnitude). In the simplest case, a pair of objects may be separated simply by splitting between the two concavities. If there is a cluster of objects, it is necessary to make several splits. After each split, each of the resulting parts needs to be checked to see if it is still a cluster requiring further splitting.

One method of determining which two concavities to split between is as follows:

1) Obtain the convex hull of the object. The convex hull is the minimum area shape that is completely convex and completely encloses the object [5].
2) Measure the distance of each edge point from the convex hull. The point with the largest distance (deepest concavity) is selected.
3) Measure the distance from the selected point to each edge point in the other concavities. The point with the smallest distance (nearest adjacent concavity) is selected.
4) Step 3 is repeated until the distance between the two points reaches a minimum. This step is required to ensure a good split is found.
5) Split the cluster by breaking along the line between the two selected points.
6) Repeat the process for each of the two smaller clusters produced.

This procedure is illustrated in figure 7, where four circular objects need to be separated. The first splitting leaves two clusters which must each be processed again.

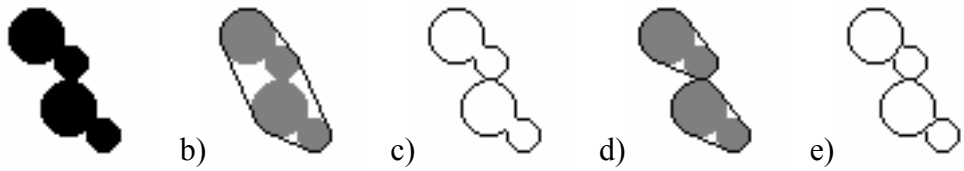

a) b) c) d) e)

Figure 7: a) touching objects; b) convex hull; c) after first split; d) new hulls; e) final result

Separation via concavity analysis does not require the objects to be circular. The only two requirements are that the individual objects are convex, and that they form concave regions when they touch. This method was successfully used to separate wool fibres [6].

DISCUSSION

All of the preceding techniques work reasonably well when the clusters consist of chains of objects. If the objects form rings or clumps, and the internal holes are filled (see figure 7), it may not be possible to separate the objects. This is especially so for the concavity analysis method since the segmentation is not parameter based.

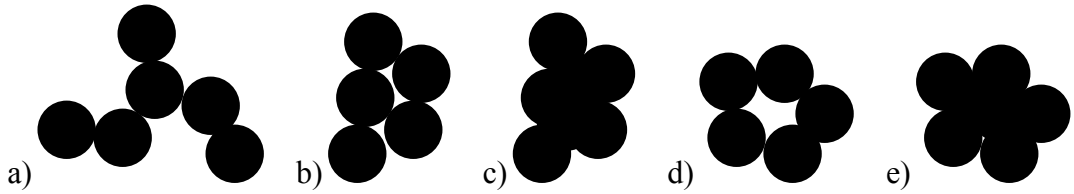

Figure 7: a) a chain of objects; b) a cluster c) with filled holes; d) a ring e) with filled holes.

One factor that must always be taken into consideration is the similarity between the model used and the objects being separated or detected. For example, it is assumed that the objects actually are circular when using the Hough transform to detect circular objects. In practise, however, such a model may only be an approximate representation of the actual object. This is especially the case when dealing with natural objects, such as collagen fibrils [4] and wool fibres [6].

The consequence of this for the Hough transform (and derived parameter based methods) is that the sharp peaks in parameter space become fuzzy and less distinct. This not only makes the peaks harder to detect, but the parameters obtained should be regarded as only approximate. They may be used to separate the touching objects by estimating the line of best split (minimum error between the two objects), however the final measurements should be made of the objects individually.

The best technique for any application depends a lot on the objects being separated. This paper has illustrated some of the techniques available and shown how they work.

REFERENCES

[1] McQueen M.P.C., "A generalization of template matching for recognition of real objects", *Pattern Recognition*, vol 13, pp 139-145 (1981)

[2] Schalkoff R.J., "Digital Image Processing and Computer Vision", John Wiley & Sons, New York (1989)

[3] Ballard D.H., "Generalizing the Hough transform to detect arbitrary shapes", *Pattern Recognition*, vol 13, pp 111-122 (1981)

[4] Bailey D.G., "Recent projects at Massey University's Image Analysis Unit", *7th NZ Image Processing Workshop*, Christchurch (1992)

[5] Arcelli C. and Levialdi S., "Concavity extraction by parallel processing", *IEEE Transactions on Systems, Man and Cybernetics*, Vol 1, pp 394-396 (1971)

[6] Orwin D.F.G. and Bailey D.G., "Measurement of wool cortical cell proportions by image processing", in "Applications of mathematics and physics in the wool industry", ed L.F. Story, WRONZ / The Textile Institute (NZ branch) pp 330-337 (1988)