Segmentation of Wear Particle Images from Used Oil Filtergrams

Elliot Duff
CSIRO Manufacturing Science and Technology
PO Box 883 Kenmore QLD 4069 Australia
Elliot.Duff@csiro.au

Abstract

In this paper, a robust technique for the image segmentation of wear particle from used oil filtergrams is presented. The different wear modes that occur within an engine can be identified by the shape, colour and texture of particles that are found in the oil. To automate the identification of these particles it is necessary to develop image processing software that is robust enough to extract the shape and surface features of each particle, accurately and independently of the quality and content of the filtergram image. To do this, a top-down approach has been chosen, which establishes the nature of any distortion or attack upon the particle and then attempts to recover the original image. This is achieved with a technique of over-segmentation and graph-based reconstruction.

Keywords: segmentation, watershed, split and merge, region adjacency graph RAG).

1 Introduction

The microscopic inspection of wear debris on a filtergram slide[1] can provide extensive and timely information on the condition of mechanical components lubricated by oil. In the past, this technique has been limited because the inspection is performed by a human operator, which can be both expensive and time-consuming. Automation can change this and transform the technique into an analytical tool, with applications over a wide range of equipment. In analysing the filtergram, the different wear modes can be identified by the shape, colour and texture of each particle [3]. The segmentation of these particles from a digital image is the first step towards the automation of such analysis. This process can be complicated by the presence of contamination, sludge and other unidentifiable debris. Furthermore, what is seen under the microscope is the result of physical, thermal and chemical attack whilst the particle is in the lubricant, and distortions generated by the vision system used to acquire the image. To automate the process of identification the software needs to account for all of these factors. To do this, it can either:

- develop different classification rules for each type of attack/distortion, or
- establish the nature of attack/distortion and try to recover the original image.

The first solution is a "brute force" approach because it requires a large number of different rules. It will first require: a rule for a particle that is "fresh" (ie. has not been attacked in the lubricant or distorted by defects in the vision system); another rule for a particle that has been crushed (eg. in the bearing); and then, another rule for a particle that is out of focus, and so on. This solution, although simple, has the disadvantage that it cannot handle situations for which it has not been trained. The second solution is a more intelligent approach. The software will need to be made aware of the different modes of attack and distortion, and the effect that these modes can have upon the particle. This solution, although complex, has more potential because it closely mimics the human decision making process. It can be achieved in the following steps:

Acquisition, acquire digital image of particle;

Over-segmentation, split the image up into a number of small parts;

Reconstruction, merge parts based upon their features;

Analysis, perform shape and surface analysis on the rebuilt particle image;

Classification, use knowledge gained in reconstruction and analysis.

The first four steps are described in the Sections 2 to 5 respectively. Classification is not within the scope of this report. Section 6 describes the results of the software, followed by Section 7 which lists a number of conclusions.

[1] Used oil is passed through a very fine paper filter. This filter is washed with a solvent removing the oil but leaving small particles on the paper. The paper is made transparent with a clarifying agent.

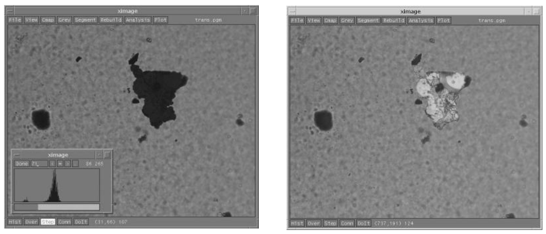

Figure 1. Image of particle taken with transmission (a) and reflected (b) light source.

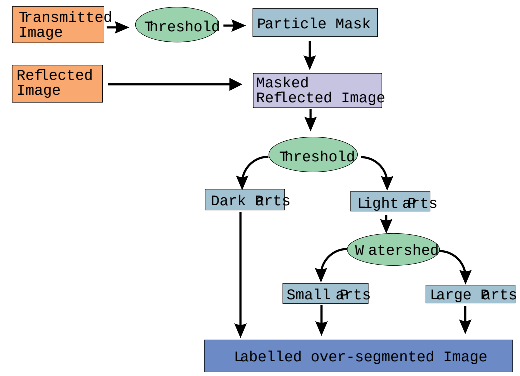

Figure 2. Flow diagram to over-segment (split) image.

2 Acquisition

In practice, the quality of filtergram images can be significantly affected by the lighting conditions. For a human operator, both transmitted and reflected light are required to correctly identify each particle on the filtergram: transmitted light is used to identify transparent sludge and visual aberrations on the slide; and reflected light is used to examine the colour and surface texture of each particle. For the human eye, the transmitted and reflected light can be distinguished by using green and red light respectively (bichromatic microscope [5]). White reflected light is subsequently used to establish the colour of the particle. Unfortunately, when this technique is applied to the computer, the segmentation of the image has proved to be very difficult. Problems occur as a result of changes in texture, shadow and/or a lack of focus at the edge of the particle. For the human operator, these effects do not normally present a problem because the brain is able to interpolate the shape of the particle from a number of visual clues. Rather than trying to interpret these clues, another solution presents itself by recognizing the fact that a computer is able to record the transmitted and reflected image separately (Figure 1). In this situation, the computer can use the transmitted image (Figure 1a) to generate a mask that defines the outer edge (outline) of the particle. This mask can be applied to the reflected image (Figure 1b) to measure the shape, colour and texture of the particle. Although this technique has the advantage that the segmentation of the transmitted image is relatively straight forward, in practice it is not very robust — less than 40% of the test images were correctly segmented. This failure lead to the development of the following split and merge technique.

3 Over-segmentation

Over segmentation involves splitting the image into a number of smaller components (parts). This process can be based upon colour, texture, shape or size of any region in the image. Currently, over-segmentation is achieved by segmenting the particle from the background, then segmenting the dark parts of the particle from the light, and finally segmenting the small parts of the particle from the large. This process is shown in the flow diagram in Figure 2.

3.1 Segmentation of transmitted image

The first step is to segment the transmitted image (Figure 1a) ie. to isolate the wear particle from the background. To do this, a grey level threshold needs to be established. It is inappropriate to use a fixed threshold because the lighting conditions and contrast are affected by the different levels of contamination in the oil. A popular automatic threshold technique is to find the minimum between two peaks in the brightness histogram. In a normal image, two peaks are expected; one peak for the dark particle, and another for the background light source. An appropriate threshold is the minimum between these two peaks. In theory, finding this minimum is simple, but in practice, it is complicated by the presence of noise in the image. Such noise can introduce extra peaks and troughs to the histogram. Although it is possible to find the minimum by smoothing the histogram, it is difficult to decide how much smoothing to use. If the data is very noisy, then a great deal of smoothing is required, but if one of the peaks is very small, it can be smoothed away altogether. The best solution is to start out with a small amount of smoothing, calculate the derivative of the histogram (the min and max) and then count the zero crossings. Smoothing is progressively increased whilst there are more than three zero crossings, until there are exactly three zero crossings; two maxima and one minima.[2]

[2]Sometimes this criteria will fail, eg. if there are three distinct populations, or if the contrast is very low and the two populations merge. In both cases, an alternative thresholding routine is used; the Kmeans technique. It assumes that two populations exist in the grey level histogram and then

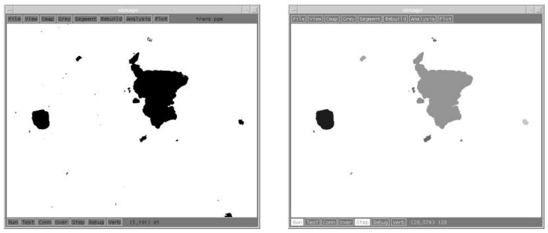

Figure 3. Automatic threshold (a) and labelling (b) of transmission image.

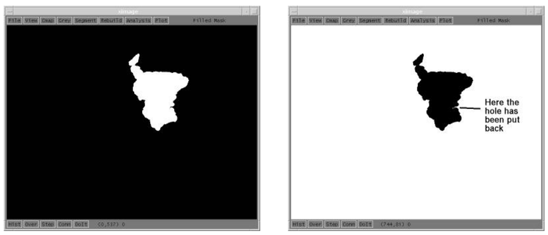

Figure 4. Background filling (a) and hole removal (b).

The effect of automatic thresholding upon the transmitted image can be seen in Figure 3a.

The next step is to isolate connected regions (ie. to label each particle in the image). This is shown in Figure 3b, where each particle has been assigned a different grey level. Since only the largest particle is of interest in this analysis, the area of each particle is measured and all but the largest particle is removed. The next step is to find if there are any holes in the particle. Holes in this mask may either be real or the result of internal reflections in the microscope optics. To find these holes, the previous image is inverted and the background filled to generate a mask that represents the area *outside the particle* (Figure 4a). If this mask is then subtract from previous image, the resulting image represents holes that are *within the particle*. These holes can be labelled and sorted. In this case, there is only one hole. The shape and colour of this hole is calculated. If it appears that it not physically real (i.e. noise or back scattered light) then the hole is removed from the mask of *the particle*. In this case, the hole is real (ie. it consists of light from the transmitted light source) and it is left in the particle mask (Figure 4b). This information is very important for classification because it infers that the particle has been crushed.

---

finds the centre of the two populations by iteratively moving the centres and then calculating the differences between all of the points and the centre. The threshold is the midpoint of the two centres. This technique will always find a midpoint between two populations, even if two populations do not exist.

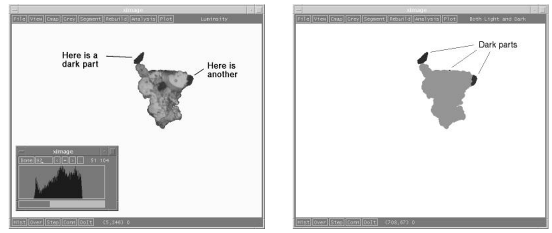

Figure 5. Masked luminosity (a) and labelled dark and light parts (b).

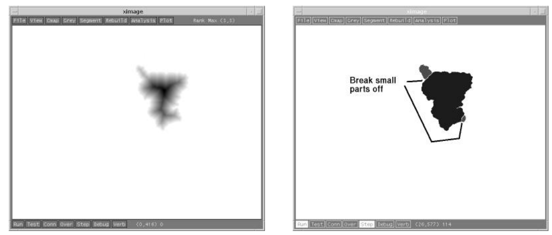

Figure 6. Distance transform (a) and watershed (b).

3.2 Segmentation of reflected image

The next step is to segment the reflected image (Figure 1b) ie. to segment the dark parts from the light parts. Since the background image is of no interest in this analysis, the luminosity image can be masked with the binary image generated in the previous section (Figure 4b). The results of this operation is shown in Figure 5a. The histogram in the bottom left hand corner of this image demonstrates that the reflected light from the particle has at least two populations. Therefore, the image can be segmented using the minima between the two peaks, to isolate the light parts of the particle from the dark. The light parts are filled, sorted and labelled. If the image is inverted, it is possible to segment and label the dark parts. If the dark and light parts are added together, a labelled image of both dark and light parts is generated (Figure 5b). Unlike the segmentation of the transmitted image, segmentation of the reflected image will only proceed if there are two clearly defined peaks in the brightness histogram.

3.3 Watershed segmentation

The final step is to perform watershed segmentation (Figure 5b) ie. to segment the small parts of the particle from the large. The watershed segmentation is based upon a distance transform of the light parts of the particle determined in the previous segmentation (Figure 5b). In a distance transformation, the value of each pixel (grey scale) is calculated

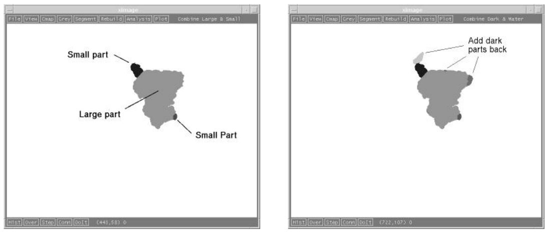

Figure 7. Labelled image of small and large parts (a) and all parts together (b).

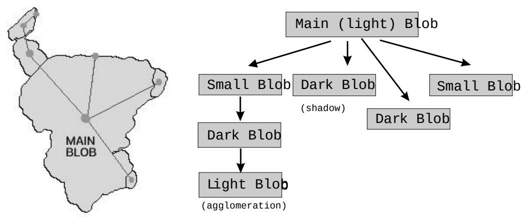

Figure 8. Connectivity of parts (a) and region adjacency graph (b).

from the distance to the nearest edge (Figure 6a). To understand watershed segmentation it is useful to visualize the maxima in the distance transform as valleys on a plain. As it "rains" each valley fills with a different "colour" of water. When different "colours" of water meet, along the watershed between each valley, a line is drawn and the image is segmented. To reduce the complexity of the problem, rather than seeding each valley with a different colour, a threshold is established (based upon the distance histogram) and each valley is labelled as either deep (large part) or shallow (small part). As it "rains", small valleys are segmented from the large along the watershed line that divides them. The effect of this technique can be seen in Figure 6b. Since the large and small parts are now isolated, it is a simple matter to separate, sort and label the small and large parts. The small and large parts are recombined and relabelled in Figure 7a. Since the watershed routine was only performed on the light parts of the reflected segmentation, it is possible to combine, the small, the large, and now the dark parts of the particle into one labelled image (Figure 7b).

4 Reconstruction

Now that the image has been split into a number of parts reconstruction can proceed. Reconstruction of the particle is based upon the features of each part. The decision to include (or reject) a part can be based upon the shape, colour and/or texture of each part. It can also be related to the spatial relationship between each part, and a database

Figure 9. Absorb shadows in labelled image (a) and average luminosity of parts (b).

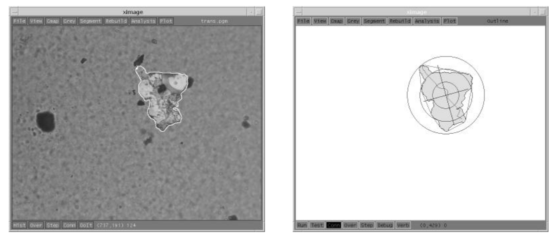

Figure 10. Outline of final mask on reflected image (a) and shape analysis (b).

of sensible arrangements, ie. arrangements that are physically meaningful. This reconstruction is aided by establishing the connectivity of each part to one another. Such a relationship is shown in Figure 8a where a line is drawn from the centre of each part that shares a common border with its neighbour. Here each part of the segmented image corresponds to a vertex, whilst edges contain adjacency information between different parts. This type of relationship is commonly referred to as a region adjacency graph (RAG). It has been used widely in image analysis [2] and has considerable application in the medical field [6]. More recently, it is proving useful for video encoding [4]. Work related to this application can be found in [1] which uses RAG to reconstruction images based upon texture, and [7] which uses RAG to address the problems associated with lighting and focus in microscopy to extract line networks from noisy low-contrast images.

Once the RAG has been created (Figure 8b) the first task is to remove shadows along the edge of the particle. A shadow is defined as a part that lies between a light part and the background. It is typically thin with a high contact length. Given these conditions, if a part is identified as a shadow, the part is absorbed into its neighbour — without modification to the parent properties. This is demonstrated in Figure 9a where a small shadow along the top right hand side of the main part has been absorbed. The next task is move down the graph, starting at the largest part, absorbing parts that appear to belong and removing those that do not. If they belong, the part is absorbed into its neighbour by adding its properties and inheriting its neighbours. If not, the part is removed. When a part is removed from the RAG, the graph is broken and the remaining parts on that particular branch are ignored. This is called pruning and is a desirable behaviour when there is an agglomeration of particles. Currently, the absorption and rejection criteria are based upon the average colour and luminosity of each part. The average luminosity of each part is shown in Figure 9b. The effect of pruning and absorption upon the graph can be seen in Figure 10a where an outline of the "final mask" is overlayed upon the reflected image. In this image, the dark part on the top left hand corner and right side have been removed. These two parts are clearly an agglomeration and the software is correct in removing them.

5 Analysis

Now that the particle has been rebuilt it is possible to perform shape analysis upon the final mask (particle outline). This is shown in Figure 10b where: the cross represents the minimum and maximum moment; the circles defines the minimum and maximum radii; the line down the centre defines the skeleton; and the taut string around the particle is calculated from Feret analysis. The degree of confidence in the shape of the particle can be used to weigh the relative importance of different features. For example, if the software experiences difficulty when thresholding the reflected image, then it would not be appropriate to rely on any feature derived from the profile of the particle to make a classification. A good example of this is highlighted with the fractal dimension of shape. This is one of key identifiers for classification of different wear modes, however it is very sensitive to the spatial accuracy of the profile. On the other hand, the fibre length, which is also a key identifier is relatively tolerant to errors in the profile. Confidence in a feature can be gained by looking at the distribution of the histogram at each level of segmentation, or by examining the individual statistics of each part. During reconstruction, parts are absorbed or rejected based upon the probability that it is a part of the "target particle". Even if the part is absorbed, the probability can be used as a level of confidence that the part is "real".

For textural and colour features, it is not necessary to recover the outline of the whole particle, but only a part that is believed to be representative of the actual particle. To guarantee that the part is valid and contains useful textural and colour information, it is important that the part be in focus and adequately lit. Confidence in these conditions can be made by making comparisons between different parts. If surface analysis is performed upon each part of the particle then the following comparisons can be made: firstly, if the parts have different textural scale, then the part with the lower resolution is probably out of focus; secondly, if one part is dark and the other light, then the dark part is probably in shadow; and finally, if two parts have different textural orientations, then it is likely that the particle has either been folded, or there is an agglomeration of two or more particles.

6 Results

Over 90 pairs of transmission and reflected images were acquired with a high resolution digital camera (Kontron ProgRes 3008) fitted to a Nikon bichromatic microscope. In more that 50% of cases, over-segmentation and reconstruction had an effect upon the final shape of the particle outline, when compared to simple transmission segmentation (see Section 2). If segmentation is deemed to be successful when in agreement with the human operator[3] then:

- 72% were correctly segmented
- 20% were incorrectly segmented.

Of the remaining images, five image pairs could not be identified nor segmented by the human operator and three image pairs were influenced by hardware failure, ie. where the failure has occurred during the acquisition of the image. One of the failures was due to a misalignment of the transmitted and reflected image, whilst the other two were caused by a problem with the transmitted light source and/or camera gain.

6.1 Examples

An example of successful segmentation is shown in Figure 11, where most of the oil-sludge has been removed from the particle except for a small ring on the bottom right hand corner.Another example of successful segmentation is shown in Figure 12, where all of the agglomeration on the left hand side of the particle has been removed.

The segmentation software was not always able to deal correctly with agglomeration. An example of this situation is shown in Figure 13a, where the software has incorrectly assumed that the hole in the centre of the particle is real, and not just an agglomeration of different wear particles. Unfortunately, shape analysis of this image would not be able to identify the curled nature of one of the particles (this is an example of cutting wear which is very serious mode of failure). Another example is shown in Figure 13b, where the software incorrectly identifies the shadows on the left

[3]The software draws a line on the image where it will segment the particle. The segmentation is successful if a human operator agrees with the position of this line. Of course there will be cases where the position will differ slightly to the position chosen by the human operator, but if the error does not greatly influence the shape of the particle (i.e. small bumps) then the software will still be deemed to be successful.

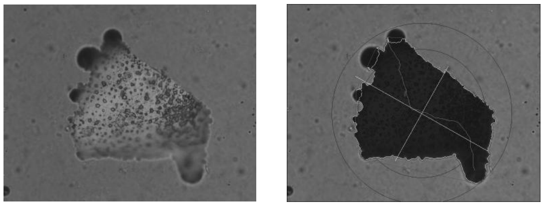

Figure 11. Segmentation of oil sludge.

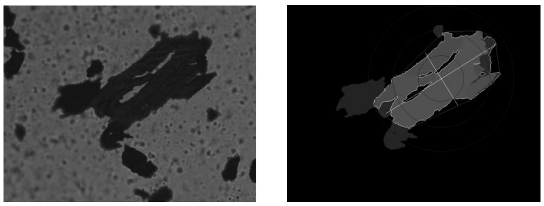

Figure 12. Segmentation of agglomeration.

and right hand sides of the particle as a dark agglomeration. Currently, the software identifies a shadow as a thin part that lies between the particle and the background. In this case, the parts are rejected because they are too thick. To prevent this from happening the software could relax the "shadow thickness" rule, but this would have a side effect of incorrectly identifying an agglomeration as a shadow in a different situation. This is a typical dilemma faced by rule-based systems and requires further research.

7 Summary

In this paper, we have discussed the problems associated with the segmentation of wear particles on used oil filtergrams. As discussed in Section 2, simple bichromatic segmentation was unable to segment more than 40% of the test sample. Most of the problems were due to the fact that:

- the particles are often surrounded by contamination, sludge and other unidentifiable debris;
- the particle can be distorted by various forms of attack whilst in the lubricant;
- the image of the particle is distorted by the vision system used to acquire the image.

To overcome these problems a robust segmentation algorithm has been developed based upon over-segmentation and graph-based reconstruction. This has lifted the success rate of segmentation to greater than 70%. This success will enable work to proceed on the difficult task of classification, which will be made easier, by using some of the knowledge gained in the reconstruction process. Another advantage of this technique is that it provides information regarding

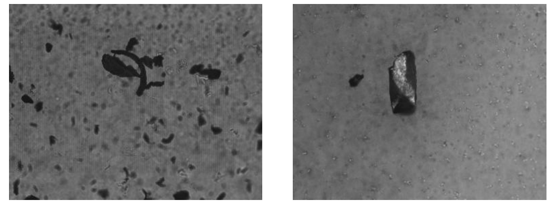

Figure 13. Examples of failure, agglomeration (a) and shadows (b).

the level of confidence that the particle has been correctly segmented. This is a valuable piece of information for any automated system so that it can reject the 10% of images which normally cannot be identified by the human operator.

8 Acknowledgements

I would like to thank ACARP for their funding of this project and Mr. Peter Ball of Machine Reliability Services for both his advice and the use of his equipment.

References

[1] Patrick Chen and Theodosios Pavlidis. Segmentation by texture using a co-occurence matrix and a split-and-merge algorithm. *Computer graphics and Image Processing*, 1979.

[2] Theodosios Pavlidis. *Structural Pattern Recognition*. Springer-Verlag, 1977.

[3] B J Roylance and S Raadnui. The morphological attributes of wear particles - their role in identifying wear mechanisms. *Wear*, 175:115–121, 1995.

[4] P Salembier and L Garrido. *IEEE transactions on Image processing*, 9(4):561–576, 2000.

[5] Douglas Scott, William W Seifert, and Vernon C Westcott. The particles of wear. *Wear*, page 88, 1975.

[6] Milan Sonkar, Gobal Sindaramoorthy, and Eric Hoffman. Knowledge-based segmenation of intrathoracic airways from multidimensional high resolution ct images. *Medical Imaging 1994: Physiology and Function from Multidimensional Images*, pages 73–85, 1994. SPIE.

[7] Timo Viero and Dominique Jeulin. Morthological extraction of line networks from noisy low-contrast images. *Journal of Visual Communication and Image Representation*, 6(4), 1995.