Object Segmentation Within Microscope Images of Palynofacies

J.J. Charles[1], L.I. Kuncheva[1], B. Wells[2], I.S. Lim[1]

Abstract

Many industries require identifying objects contained within microscope images. Our task is to locate and count the pieces of inertinite and vitrinite in thin section images of rocks. The classical watershed algorithm oversegments the objects because of their irregular shapes. In this paper we propose a method for locating multiple objects in a black and white image while accounting for possible overlapping or touching. The method, called Centre Supported Segmentation (CSS), eliminates oversegmentation and is robust against differences in size and shape of the objects.

Key words: Image processing; object segmentation; watershed; palynomorphs

1 Introduction

Accurate recognition of biological remains found in palaeosediments is the basis of palynology and micropaleontology and this underpins the interpretation of palaeoenvironment, chronostratigraphy and much more. The fundamental importance of accurate recognition of fossil material under the microscope has spured considerable effort into automating the task. In the last 25 years some significant developements have been made in recognising specific types of fossil material under ideal conditions (England et al., 1979; France et al., 2000; Weller et al., 2005).

Nevertheless, obtaining statistically significant, unbiased and reproducible results from automated analysis of microscope images is still regarded as a challenge. The task of identification is expected to require the skills of highly qualified and experienced human eyes and brains for some time. Even in the unrealistic case of individual, whole, undeformed specimens, difficulties arise

[1] School of Informatics, University of Wales, Bangor, LL57 1UT, United Kingdom
[2] Conwy Valley Systems, United Kingdom

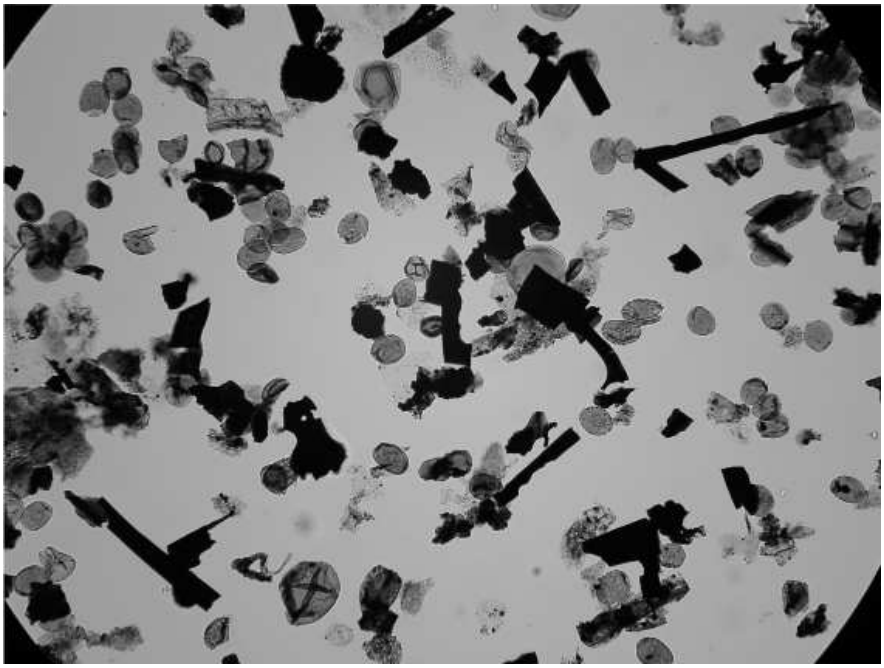

Fig. 1. Microscope image of palynofacies. The dark objects are referred to as "kerogen"

from the diversity of the species to be recognised, the variability in the image acquisition techniques as well as the subjectivity of the visual analysis (Bollmann et al., 2004). The actual geological task is made considerably more difficult given that the material being analysed has arisen from biological remains. These remains are subjected to initial distress at time of deposition and subsequently altered and deformed by burial stresses and tectonic deformation. Furthermore the remains are then retrieved from their current position deep in the Earth by techniques which were not designed primarily for optimum sample preservation. Finally, the processed material is arranged haphazardly on a slide, with both material of interest and other materials overlapping and partially hiding each other.

Starting from the original geological task of a mixture of objects arranged randomly, we seek to create a collection of individual specimens categorised as objects of interest and objects not of interest. In addition to the task facing the geologist, the image analysis task has to address the reality that each slide is seen under different conditions e.g. variations in light intensity, colour balance, background, etc. These variations need to be accounted for in order to create conditions which are sufficiently similar or standard for the algorithms to be useable and the results to be statistically valid. Only then can the large body of work on recognition of individual specimens be made commercially useful.

Figure 1 shows a typical image of an assemblage of objects retrieved by sieve analysis from a sediment. The material consists of light or transparent ob-

jects (palynomorphs and amorphous organic matter) and opaque humic kerogen which can be subdivided into inertinite and vitrinite. The dark objects have irregular shapes and different sizes. Sometimes small dark objects appear within light objects. The light objects, on the other hand, vary in texture, intensity and transparency. Our goal in this study is to devise an algorithm for automatic identification of dark objects in the image by placing a "centre" point in each object. Automated image analysis of this type is a step towards a completely automatic classification system for palynological images (Weller et al., 2005). Up to now, human intervention has been assumed in detecting the objects in the image. At the next stage, automatic classification of the cropped objects can be attempted based on expert knowledge, extraction of salient features and various machine learning and pattern recognition methods (Weller et al., 2005).

Counting objects in an image is straightforward for disconnected objects. However, counting connected or overlapping objects can prove difficult. The standard approach to this task consists of two steps. First, background removal (segmentation into background and foreground) is carried out so that the objects of interest appear as a black foreground. Second, the foreground is further segmented by the watershed transform (Vincent and Soille, 1991) to identify separate objects. The black and white image is divided into regions, each region containing one object. We show that watershed transform oversegments the images and hence we propose a new method which finds adequate object centres.

The rest of the paper is organised as follows. Section 2 explains the type of objects in the image. Section 3 explains the background removal step. In Section 4 we demonstrate the drawbacks of the watershed method for segmenting the foreground objects. Section 5 introduces our algorithm for finding object centres. The results of an experiment are presented in Section 7. Section 8 offers conclusions and outlines future research directions.

2 Inertinite and vitrinite

Inertinite and vitrinite are the visible remains of plant materials that are common components (macerals) of coal. They are formed from fossilised woody tissues that have been preserved under different redox conditions. Vitrinite is formed from humic acid gels preserved under everwet reducing conditions, inertinite is formed where slight drying (and hence oxidation) has occurred. Vitrinite is less carbon rich than consanguineous inertinite, so that under transmitted light microscopy vitrinite is often seen as brown flakes (usually homogenous but sometimes showing residual cell structure) and inertinite black. Under reflected light, inertinite has a higher reflectance than vitrinite, and of-

ten shows shapes based on thick (lignitic) plant cell walls (Stach et al., 1982).

3 Background removal

The common procedure for background removal is to apply an intensity threshold on the greyscale image (Gonzalez et al., 2004; Bollmann et al., 2004). Global thresholding may bring as foreground some of the less illuminated parts of the image together with all their content. To avoid this, background illumination should be equalised. An empty slide can be stored and used for background correction (Weller et al., 2005). The problem with the empty-slide method is that it is specific for the current settings of the microscope and will not produce accurate results in images of a different origin. Further methods include blurring and normalisation (SYSTAT Software Inc, 2002; Zawada, 2003) as well as fitting a 2-dimensional quadratic function [3].

Here we use a recently proposed approach for illumination correction in microscope images (Charles et al., 2006a). Each pixel in the image has to be labelled as foreground or background. The background is typically brighter in the middle and darker towards the edges. If a constant threshold on the grey level intensity is used throughout the image some objects near the edges of the image will be lost within a "rim" labelled as foreground. As the light intensity fades with the square of the distance from the source, a quadratic function can be used to model the illumination. However, the single function model can be too coarse and inaccurate, especially when the foreground occupies a substantial part of the image. Instead we construct an estimate of the background by fitting parabolas to the background intensities in every horizontal row of pixels. Thus the fitted parabolas form a new image of the background that we call the 'horizontal estimate'. Similarly we can form a 'vertical estimate' using columns of pixels. Let A be the original grey scale image and $A(p)$ be the grey level intensity of pixel p. For any pixel p, the background intensity is calculated as $A_m(p) = (A_h(p) + A_v(p))/2$, where tags h and v denote the horizontal and the vertical estimates, respectively. To eliminate the background effect, a normalised image, A_n, is calculated as $A_n(p) = A(p)/A_m(p)$ and is subsequently rescaled between 0 and 255.

The histogram of the grey level intensities of A_n has two distinct peaks, the left peak corresponding to the inertinite and vitrinte material and the right peak representing everything else. The threshold for foreground/background segmentation, T, is chosen as the intensity corresponding to the minimum between the peaks. To estimate T, the histogram is smoothed once with a

[3] The "Montage" project of the US National Virtual Observatory, `http://montage.ipac.caltech.edu/baseline.html`

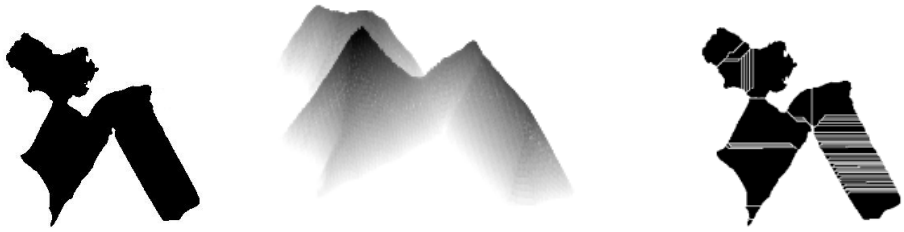

(a) cropped segment (b) distance surface (c) watershed segmentation

Fig. 2. (a) Cropped subimage from Figure 1, (b) illustrates its distance function as a surface and (c) shows the result of applying the watershed algorithm to the negative of the distance function in (b), (56 regions identified)

filter of length 3, the curve segment between the two peaks is identified and a parabola is fitted. T is the minimum of the fitted curve calculated from the coefficients.

Thresholding will result in a black and white image (B/W), black regions representing the vitrinite/inertinite pieces and white regions representing everything else.

4 Watershed segmentation

First, a distance transform is applied (Borgefors, 1986) to the segmented image. The distance function D for a pixel, p, denoted $D(p)$, gives the Euclidean distance to the nearest white pixel. This result can then be thought of as a surface with $D(p)$ being the height for pixel p. The watershed transform is applied to $-D(p)$ and fills the troughs with "water" and finds the watershed ridge lines. These lines partition the image into regions so that one object is contained in each region. This process is very effective for segmenting touching objects with circular shapes. However, rectangular objects cause difficulties and oversegmentation occurs due to the increased number of regional minima in the distance transform.

The inadequacy of the watershed method for detecting dark objects in palynomorph images is illustrated in Figure 2. The number of regions found in the image is 56 while we are looking for just three objects.

5 Centre Supported Segmentation (CSS)

We propose an alternative segmentation method based on identifying a centre point for each object. The Centre Supported Segmentation (CSS) is applied

on the black and white image where the black foreground are the objects to be segmented. The result from CSS is a set C of centres of objects. The centre of an object is needed for several reasons:- (1) counting the number of objects, (2) viewing an object by moving the scanning camera to the centre and (3) cropping the object for further analysis and classification.

Definition. Any pixel p with the largest distance $D(p)$ within the object is called a **centre** of this object.

For example, the centre of a filled-in circle will be its geometrical centre. However, a doughnut-shaped object will have infinitely many centres, none of which will be the geometrical centre of the figure.

5.1 CSS algorithm

A pixel p with coordinates (x, y) has an *8-neighbourhood* consisting of the set $N_8(p) = \{(x+1, y), (x-1, y), (x, y+1), (x, y-1), (x+1, y-1), (x+1, y+1), (x-1, y-1), (x-1, y+1)\}$. Two pixels p and q are *8-connected* if there exists a path of pixels between p and q where each pixel in the path is of the same intensity and in the 8-neighbourhood of the next pixel. A set of pixels that are all connected to one another is called a *connected component*.

CSS is applied to the distance function of the B/W image. It first identifies the centres of all possible objects and then filters out the centres which are likely to be noise. Two lists are created: the list C of centres and the list V of their *merging heights*. List C is constructed in the following way. Initially C is empty. Suppose that $m = D(q_1) = \max_p D(p)$ is the unique maximum of D across the whole image. By thresholding D at m all pixels p where $D(p) >= m$ are set to black the rest are set to white, D will define a black and white image B, containing one black point at q_1. Pixel q_1 is taken to be the first centre in C. The next maximum height, $m' = \max_{p \neq q_1} D(p)$ is identified. The new black and white image B' resulting from the thresholding with m' will contain more black points, including points around pixel q_1. Since the object with centre q_1 is already accounted for, we remove this connected component from B'. Let Z be the set of pixels of a connected component in B'. The centre of the object represented by Z is q_2, where $D(q_2) = \max_p D(p), \forall p \in Z$. The new centres are appended to the current C in descending order of their heights. A subsequent threshold $m'' < m'$ is applied to produce image B'' from D, and all objects with centres in C are removed from B'' in the order they are stored in C. The remaining connected components are used to find new object centres, and so on.

If we apply this process to the image in Figure 2 (a), 56 centres will be found,

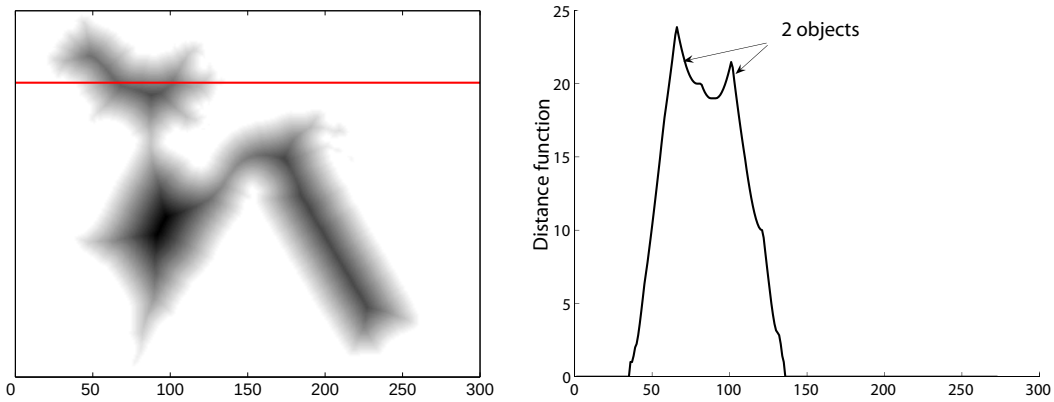

(a) (b)

Fig. 3. (a) displays the distance image of Figure 2, a projection line is plotted through the top of the object. (b) shows the distance function along the projection line indicating how a jagged peak could cause oversegmentation in CSS

each one located in its own separate region defined by the watershed algorithm (subplot (c)). As with the watershed method, small shape irregularities on the periphery of the object will result in a jagged peak of the distance function with many local maxima of similar heights. Each little spike will generate a centre. Figure 3 shows an example. Subplot (a) displays the distance image of the example in Figure 2 (a), and plots a projection line through the top object (a piece of vitrinite). Subplot (b) gives the distance function across the projection line. In this case two distinct connected components will be found, generating two object centres instead of one.

Oversegmentation can be avoided by attaching to each centre a *merging height*. The merging height of centre q_i, denoted v_i, $(v_i \leq D(q_i))$, is the lowest height at which q_i defines a connected component disjoined from any connected components of q_j such that $D(q_j) > D(q_i)$. For any value lower than v_i, q_i and another centre at a larger D will share a connected component. Figuratively speaking, the object of smaller size (smaller peak $D(q_i)$) is eclipsed by an object of a bigger size $(D(q_j) > D(q_i))$.

The merging heights of centres of large objects will be low even if they overlap with smaller objects. The merging height of large objects will be updated until they are joined to a larger object. On the other hand, centres of smaller objects corresponding to noise at the peak will have high merging heights. For example in Figure 3 (b) two object centres will be found, the centre on the left will have a merging height of 0 but the centre on the right will have a merging height of 19 which corresponds to the height of the trough between the two peaks. The centres with large v_i will be candidates for elimination.

The algorithm for identifying the centres C and their merging heights V is

CENTRE SUPPORTED SEGMENTATION (CSS) ALGORITHM: C AND V

(1) Given is B/W image B. Initialise $C = \emptyset$, $V = \emptyset$

(2) Find the distance image D for B and sort all the distinct distance values in descending order: $m_1 > m_2 > \ldots > m_T$.

(3) For $i = 1 : T$

  (a) Find B/W image B_i by thresholding D at m_i.
  (b) For $j = 1 : |C|$ (each centre in C)
    (i) find the connected component Z_j
    (ii) if Z_j has no intersection with any of $Z_1, \ldots, Z_{j-1}$, then set $v_j = m_i$
  (c) Remove all connected components $Z_1, \ldots, Z_{|C|}$ from B_i
  (d) Find the remaining connected components in B_i, e.g., $Z_{i,1}, \ldots, Z_{i,k}$. For each component, find $q_{i,t} = \arg\max_{p \in Z_{i,t}} D(p)$ as the centre of this component, $t = 1, \ldots, k$.*
  (e) Augment C and V

$$C \leftarrow C \cup \{q_{i,1}, \ldots, q_{i,k}\}, \quad V \leftarrow V \cup \underbrace{\{m_i, \ldots, m_i\}}_{k}$$

(4) Return C and V

* Note: As each possible distance is checked, the connected components at distance m_i, after removing all connected components at (3.c) will all be exactly at height m_i. These may be single points or clusters of points at the same height. Thus any point from $Z_{i,t}$ can serve as the centre of the component.

Fig. 4. The Centre Supported Segmentation (CSS) algorithm. Stage 1: obtaining centres C and merging heighths V.

detailed in Figure 4.

In Stage 2 of the CSS algorithm redundant centres are eliminated. A threshold s can be applied to account for the minimum allowable size of an object. All objects with centres q, such that $D(q) < s$, are discarded. If the algorithm is run with $s = 0$, it will find all the specs in the image as objects of interest. The value of s can be estimated subjectively or can be learned from a sample of training images where the objects of interest have been pre-labelled by hand. Weller et al. (2005) propose an empirical threshold of 14 μm.

5.2 Overlap

The cropped image in Figure 2 (a) looks like three touching objects, however it may also be a genuine set of 56 tightly packed objects. We introduce a

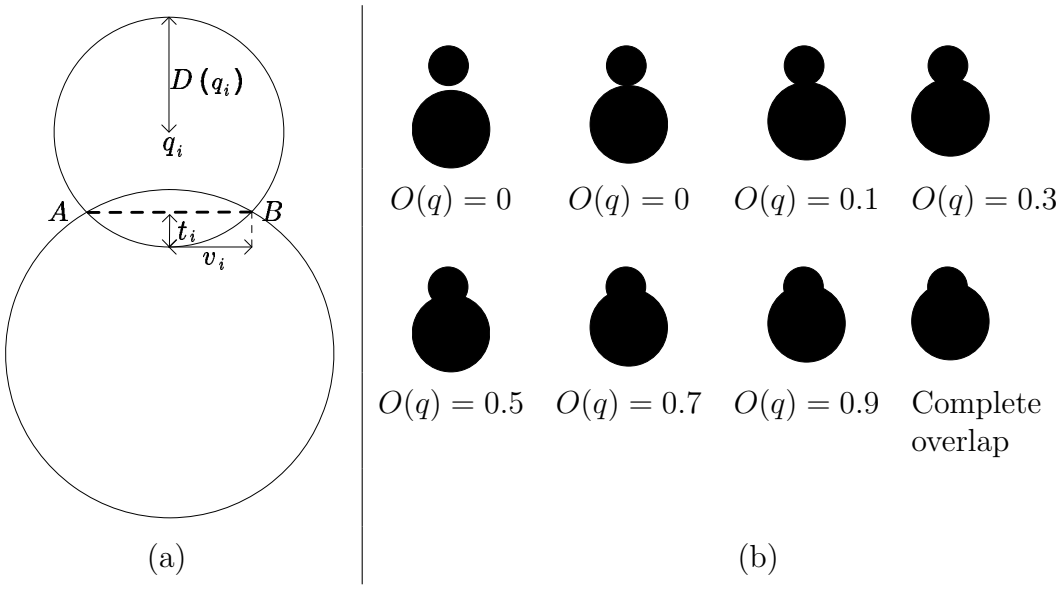

Fig. 5. (a) defines the concept of overlap, (b) illustrates the value of overlap $O(q)$

parameter d to determine which centres need to be removed. The degree of overlap is defined using two intersecting circles as demonstrated in Figure 5 (a), and is measured with respect to the smaller circle. If the two circles are of the same size, each of the two can be chosen. The overlap value is defined by the two intersection points A and B. Denote by t_i the minimum distance from the mid-point of the segment AB to the edge of the smaller circle. The overlap is defined as the ratio of t_i to the radius of the circle, $D(q_i)$. The length t_i is found using v_i, the merging height of the centre q_i. The merging height is half the length of AB and is shown as v_i in the diagram. Since $t_i = D(q_i) - \sqrt{D(q_i)^2 - v_i^2}$, the degree of overlap for a centre q_i is

$$O(q_i) = 1 - \sqrt{1 - (v_i/D(q_i))^2}.$$

By definition $D(q_i) > v_i$ and so $O(q_i) \in [0, 1)$. A centre q_i with overlap $O(q_i) = 0$ means that the object is isolated. As the overlap approaches 1, the object is increasingly covered by a larger item. The overlap of two circles is demonstrated in Figure 5 (b). The small circle is increasingly covered by the larger circle. The overlap value $O(q)$ is also shown. As soon as the two circles merge so that $q_i = t_i$, the CSS algorithm will continue to recognise one object, in this case we have complete overlap.

5.3 Filtering centres

To remove centers from C we set a limit d on the amount of overlap such that if $O(q_i) \leq d$ then centre q_i is kept in C and discared otherwise. If we wish to

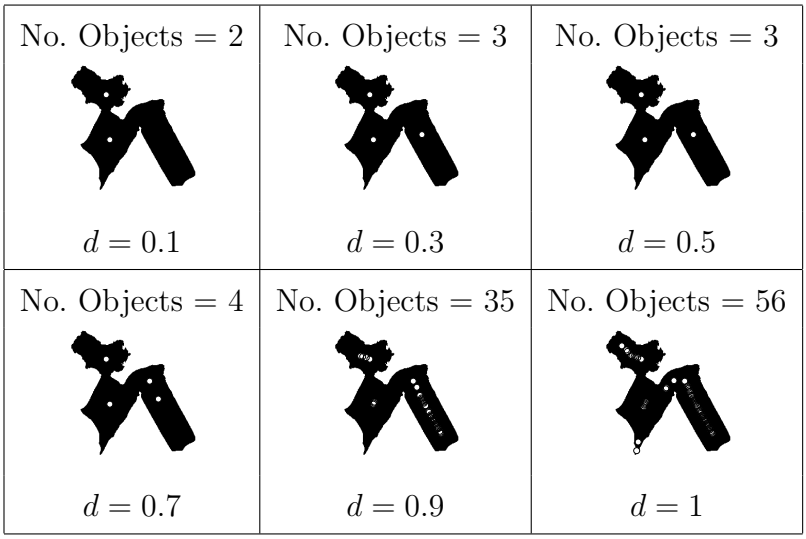

Fig. 6. Effect of the allowable overlap d on the image segmentation. Best segmentation (3 centres) is obtained for d between 0.3 and 0.5

segment an image to its maximum detail then the threshold is set to $d = 1$. This would return the same number of segments as the watershed algorithm. By adjusting d oversegmentation can be prevented. The "noisy" centres occur due to small deformations in the shape and this will correspond to relatively large merging heights yielding large overlap values. Hence setting a threshold d not only specifies the connectivity of objects but also eliminates the "noisy" centres. The effect of setting a maximum allowable overlap can be seen in Figure 6. At $d = 0.5$ we have a "correct" (intuitive) segmentation into three objects and at the limiting case of $d = 1$ we obtain the 56 segments that are produced by the watershed algorithm.

6 Extracting objects

After filtering the centres the corresponding individual objects can be extracted. The watershed algorithm is applied to the negative of the distance image $D(p)_{neg} = -D(p)$ modified so that only regional minima occur at each centre. This technique is known as *image imposition* (Soille, 2003). In this way the centres are used as markers and the watershed algorithm will segment according to the position of the centres. The watershed ridge lines will form boundaries for each object. To extract the objects, the watershed boundaries are overlaid in white on top of the black and white image. Thus the new B/W image will consist of black connected components corresponding to objects, which can be easily extracted for further analysis. [4]

[4] Matlab function *bwlabel* can be used to label all connected components

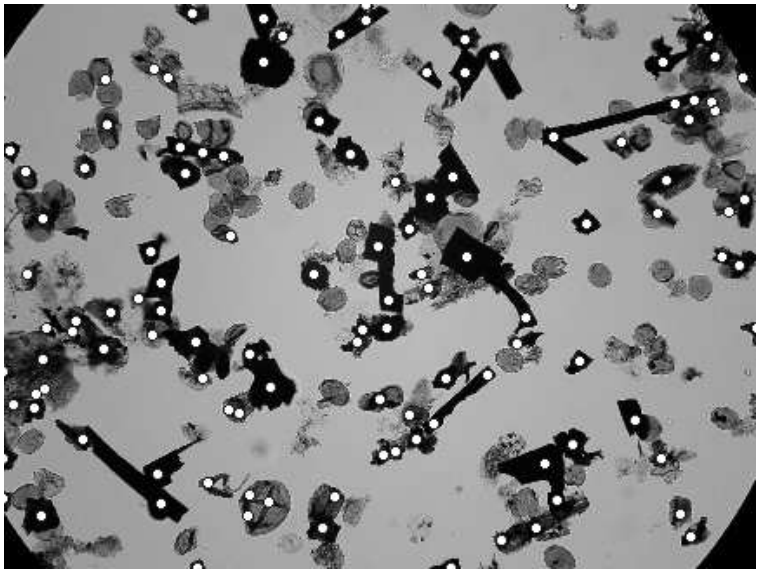

Fig. 7. Centres found by CSS

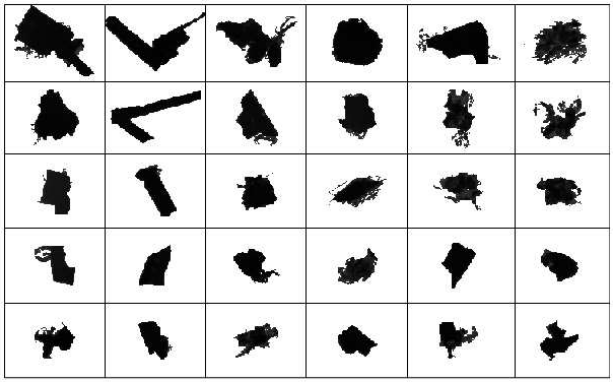

Fig. 8. Objects extracted using image imposition with centres acting as markers

7 Experimental results

7.1 An illustration

Figure 7 displays the centres of dark objects within the image in Figure 1. The original image of size 1704×2272 pixels was sampled every 3 pixels. By sampling the image we decrease the processing time for CSS, the sampling number is not sensitive to the output of CSS and is sufficient provided the sampled image can still be segmented into individual objects by human eye. Once the coordinates of centres have been found for the sampled image they are mapped back to the original image so that extraction of objects can proceed. The chosen parameter values were $s = 4$ and $d = 0.5$. The objects were then extracted, the first 30 largest objects are shown in Figure 8.

Table 1
Segmentation results for the watershed method and Centre Supported Segmentation (CSS) on 6 microscope images of palynofacies

Image	Watershed		CSS	
No	$S(C, C^*)$	time [s]	$S(C, C^*)$	time [s]
1	0.43	7.37	0.14	48.87
2	0.58	7.32	0.22	107.07
3	0.60	7.59	0.29	94.21
4	0.55	7.31	0.13	73.05
5	0.53	7.39	0.22	87.61
6	0.50	7.30	0.11	31.90

Note: Small values of $S(C, C^*)$ indicate better match between the obtained (C) the ideal (C^*) centres

7.2 A comparison with the watershed method

To evaluate the quality of the obtained set of centres C against a known set of centres C^* we use a recently proposed measure $S(C, C^*)$ (Charles et al., 2006b). The three components of the measure evaluate the under/over segmentation of the objects, the proportion of centres placed in the background rather than in objects, and the distance between the guessed and the true centres. The measure varies between 0 and 1 where 0 signifies perfect match while 1 means that set C is useless. Table 1 shows the results on several microscope images of palynomorphs. "Ideal centres" C^* were manually placed on each image after the background was removed. The values of S and the processing times indicate that CSS is slower but consistently more accurate than the watershed segmentation method.

8 Conclusion

We propose an algorithm for automatic finding centres of objects in an image and extracting the corresponding objects. The algorithm, called Centre Supported Segmentation (CSS), remedies the oversegmentation problem of the watershed method traditionally used for such segmentation. The experiments show that CSS is more accurate although slower than the watershed method.

The method for finding centres is a step within building an automatic system for identifying palynomorphs and humic kerogen in images of rock samples captured through a microscope. Once the objects are extracted, classification

methods can be applied to determine their types. In this study we demonstrated how CSS operates for extracting dark objects, which could be further classed into inertinite and vitrinite. After the dark objects are removed, the remaining image will contain palynomorphs and amorphous objects. These objects can be extracted in turn using again CSS. As CSS is not meant to work on-line, the processing time is not crucial. Nevertheless future effort will be also directed towards speeding up CSS.

Application of CSS may be sought in various domains, e.g. segmenting cell nuclei and setting the intial postion of active contours (Clocksin, 2003), seperating pollen grains for automated analysis (France et al., 2000) and marker-controlled segmentation (Gonzalez et al., 2004).

References

Bollmann, J., Quinn, P. S., Vela, M., et al, B. B., 2004. Automated particle analysis: Calcareous microfossils. Image Analysis, Sediments and Paleoenvironments 7, 229–252.

Borgefors, G., June 1986. Distance transformations in digital images. Computer Vision Graphics and Image Processing 34 (3), 344–371.

Charles, J. J., Kuncheva, L. I., Wells, B., Lim, I. S., 2006a. Background removal in microscopy images, (submitted), `http://www.informatics.bangor.ac.uk/~kuncheva/papers/ckwlsubmitted06.pdf`.

Charles, J. J., Kuncheva, L. I., Wells, B., Lim, I. S., 2006b. An evaluation measure of image segmentation based on object centres. In: Proc. International Conference on Image Analysis and Recognition ICIAR. Portugal.

Clocksin, W., 2003. Automatic segmentation of overlapping nuclei with high background variation using robust estimation and flexible contour models. iciap 00, 682.

England, B., Mikka, R., Bagnall, E., August 1979. Petrographic characterization of coal using automatic image analysis. Journal of microscopy-Oxford 116, 329–336.

France, I., Duller, A., Duller, G., Lamb, H., February 2000. A new approach to automated pollen analysis. Quaternary Science Reviews 19 (6), 537–546.

Gonzalez, R. C., Woods, R. E., Eddins, S. L., 2004. Digital Image Processing Using Matlab. Pearson Education, Inc.

Soille, P., 2003. Morphological Image Analsis: Principles and Applications. Springer-Verlag, NY.

Stach, E., Mackowsky, M. T., Teichmuller, M., Taylor, G. H., Chandra, D., Teichmuller, R., 1982. Stach's Textbook of Coal Petrology. Lubrecht and Cramer Ltd.

SYSTAT Software Inc, 2002. Sigmascan`http://www.systat.com/products/SigmaScan/`.

Vincent, L., Soille, P., June 1991. Watersheds in digital spaces: an efficient algorithm based on immersion simulations. IEEE Transactions on Pattern Analysis and Machine Intelligence 13 (6), 583–598.

Weller, A. F., Corcoran, J., Harris, A. J., Ware, J. A., 2005. The semi-automated classification of sedimentary organic matter in palynological preparations. Computers & Geosciences 31 (10), 1213–1223.

Zawada, D. G., October 2003. Image processing of underwater multispectral imagery. IEEE Journal of Oceanic Engineering 28 (4), 583–594.