Red Blood Cell Segmentation from SEM Images

Joost Vromen
Human Media Interaction
University of Twente,
Enschede, The Netherlands
Email: joost.vromen@gmail.com

Brendan McCane
Department of Computer Science
University of Otago.
Dunedin, New Zealand
Email: mccane@cs.otago.ac.nz

***Abstract*—We present a model based contour tracing approach to the problem of automatically segmenting a Scanning Electron Microscope image of red blood cells. We use a second order polynomial model and a simple Bayesian approach to ensure smooth boundaries, and a postprocess ellipse fitting procedure to cull noise contours. Of all contours detected, 95.7% are correct, with a 0.6% false negative rate, and 4.3% false positive rate on 100 sample images involving more than 11000 red blood cells.**

I. Introduction

There is considerable evidence indicating that erythrocyte (red blood cell, RBC) deformability is an important determinant in the filterability of blood and subsequently in the pathology of several diseases including Huntington's disease, Myalgic Encephalomyelitis (ME) and Multiple Sclerosis (MS) [1, 2, 16, 17, 12, 15]. Further, there is significant evidence indicating the shape of erythrocytes is directly related to its deformability and hence its filterability [12, 14]. The difficulty of moving the results of these studies from the laboratory to clinical diagnosis or indeed measuring the effectiveness of treatments is largely due to the large amount of tedious work required to measure the distribution of different erythrocyte cells from a sample. The aim of this project is to automatically estimate the distribution of erythrocyte shapes from scanning electron microscope (SEM) images. This paper presents the result of the first part of this project - namely the segmentation of blood cells from an SEM image.

Figure 1 shows a typical example of the type of image we are dealing with. The blood was obtained and processed following the protocol of [12]. The immediate fixation of the blood upon drawing is extremely important as erythrocytes which are not immediately fixed tend to change morphology [13]. Images were captured at 600 times magnification using a scanning electron microscope. There are several things to notice about the images that make our problem unique. On the positive side, the images are of very good quality. On the negative side, the blood cells are highly overlapping and edge information on the inside of a cell is often just as strong as edge information between two overlapping cells. Further, although most cells are strongly elliptical with smooth boundaries, some cells are not, and it is this latter class that is of special interest in diagnosing and treating disease states. It is the problem of dealing with overlapping cells and variant morphologies that makes our problem interesting and different from previous work in this area.

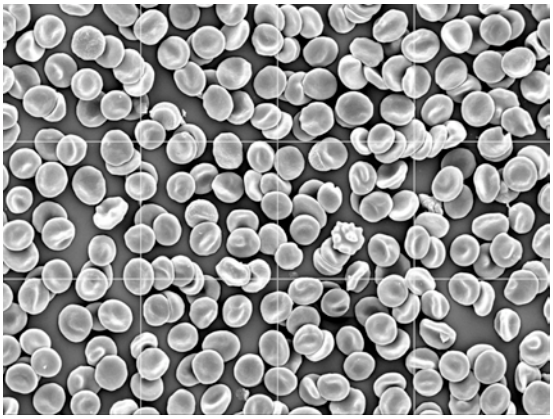

Fig. 1. A typical SEM image of red blood cells. Note that the light gridlines visible on the image were added for the sake of manual counting. These lines are removed by interpolating between adjacent lines prior to running the contour tracer algorithm.

II. Previous Work

Segmentation is one of the major problems in image analysis and considerable research has been performed in trying to solve this problem including methods of segmenting various types of cells or biological images. Aside from very simple techniques such as thresholding or basic edge detection, there are two main approaches to segmentation — region based or edge based.

For region based methods, the prototypical approach is based on morphological operators and watersheds [7, 3]. Malpica et al. [7] describe an algorithm that segments clustered fluorescence microscopy based analytical cytology images using a watershed algorithm and simple morphological operators. Their method works quite well but the images are very different to the quality of our own, and it is assumed that neighbouring cells are adjacent rather than overlapping. Di Ruberto et al. [3] describe a method for analysing and recognising malarial red blood cell light microscopy images using morphological operators and a watershed algorithm. Of most relevance to our problem is the method by which they process clustered or overlapping cells to detect individual cells. A morphological opening operation is performed to enhance boundaries followed by a standard watershed algorithm. The resulting segments are considered to be individual cells. Although producing good results for their application, the RBC

contours produced by their method are quite noisy. This is probably also due to the nature of the images, where RBCs are semi-transparent and therefore boundaries for clustered cells are difficult to locate. Since our images are cleaner in general, and the image gradients are stronger, we have focused more on edge based techniques than region based ones.

For edge based methods, typical approaches in increasing levels of complexity include model fitting [6], deformable templates [5, 1], active contours or level sets [8], and a combination of level sets and watersheds [10, 9].

Jiang and Yang [6] describes a method for finding elliptical cell boundaries using an evolutionary tabu search. The search method is used to improve the localisation of ellipses in an image - points to match are found using Canny's edge detection algorithm. Although their method works quite well, they did not include any images with overlapping or clustered cells (which would violate the elliptical condition). We have also experimented with directly finding ellipses in an RBC image (using the method of [4]), but found that overlapping cells proved too difficult without having performed prior segmentation.

Garrido and de la Blanca [5] describe a method for finding approximately elliptical cell boundaries using a generalised Hough transform combined with a deformable template model. First, edges are detected in the image using Canny's edge detector. Then junction points are removed from the edges and straight line segments extracted. The line segments are used to find circular shapes using a generalised Hough transform, which informs a deformable template for locating more accurate boundaries. Reasonably good results are obtained, although again overlapping cells are minimal in their image set.

McInerney and Terzopoulos [8] describe topology adaptive snakes. There are two stages to this method - the first is the usual snakes method using a simple Euler integration scheme. The second stage consists of reparameterising the T-snake by defining new nodes of the snake as those points where the snake crosses a regular triangular grid of the image. The use of the regular grid allows for consistent rules for changing the topology of the snake. Unfortunately, like all snake algorithms, it requires careful setting of inflationary/deflationary forces so that the contour stops at significant edges in the image but ignores insignificant ones. The method has not been tested on cell images and it's utility on images with many objects is untried.

Park and Keller [10] describe an algorithm that combines watershed segmentation and active contours to segment white blood cells in bone marrow images. Their method works by the user manually specifying multiple initial boundary contours, which are then expanded based on the watershed transformation to form "snake zones" in which the active contour is limited to remain. A global optimum within the snake zone is then found using a dynamic programming approach. The paper reports more accurate results than standard snakes, but because of the requirement for initial placement, it is unsuitable for our application (which includes hundreds of cells).

Nilsson and Heyden [9] describes a method for segmenting dense white blood cell clusters from peripheral blood or bone marrow light microscopy smear images. Their method uses a combination of level-set techniques and the watershed algorithm. The speed function used in their approach is designed specifically for the procedure used to stain the smear images and is therefore not directly usable for our problem. Also, the nature of smear images means that cell clusters do not generally include overlapping cells, but rather adjacent cells are squashed together. Clearly such a preparation method is inappropriate for an application for which the actual shape of the cell is important.

Bronkorsta et al. [1] describe a method for measuring the deformability of red blood cells directly using deformable templates. Optical traps are used to physically deform red blood cells and how quickly they revert back to a natural shape is used as a proxy for the deformability of the cell. The image analysis outlined in [1] is not applicable in our case due to the different imaging environment (capturing a video of deforming cells), but the application is the same. Our approach can be seen as a direct competitor to the method of [1] in that we are trying to solve the same problem (determine the filterability of blood by measuring attributes of erythrocytes), but our methods are very different.

As far as the authors are aware, this is the first attempt to segment scanning electron microscope images of erythrocytes. Many of the aforementioned methods are potentially applicable to our problem, but most do not consider the problem of overlapping cells (as opposed to closely neighbouring/adjacent cells), which is our main problem. Given the nature of the images we are dealing with (good contrast), we have adopted a simpler approach to many of the aforementioned techniques.

III. Method

We are interested in estimating the distribution of different erythrocyte shapes from SEM images. The SEM images were acquired on fine grain film at a magnification of 600 in an ETEC Autoscan electron microscope at 20kV. Further details on sample preparation can be obtained from [12]. The signal-to-noise ratio of the images is approximately, $SNR = 13$, where SNR is calculated as the mean of a foreground cell divided by the standard deviation of a background region. Since it is the distribution that is important, rather than an exact count from a particular slide, we can make some simplifying assumptions. We assume that the distribution of overlapped or obscured cells is identical to the overall distribution. Therefore we can concentrate on detecting and recognising the top-most cells only, and ignore all cells which are obscured. This has the further advantage of making the classification problem easier.

For the sake of simplicity and ease of implementation, we opt for a contour tracing approach. Such an approach allows us to view contour detection as a sequential problem, negotiating perceived problem areas one at a time. A simple contour trace would be a greedy trace, constantly following the highest possible gradient until the traced path intersects itself. However, this approach suffers from problems caused

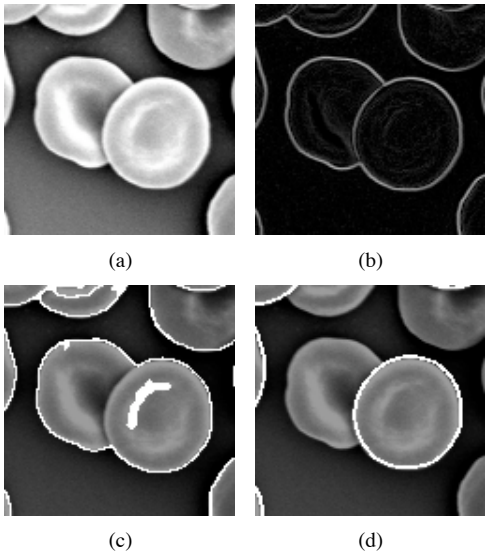

Fig. 2. Contour tracing results on overlapping cells. (a) source image, (b) gradient map, (c) results of greedy contour tracer, (d) results of guided contour tracer

by overlapping cells. Figure 2(a) and 2(b) show a pair of overlapping cells and the resulting gradient magnitude map. Since the cells are situated on a dark background, the change in brightness between a cell and the background will be much higher than the change in brightness between two cells. Consequently, edges separating the cell from its background are much clearer than those separating cell from cell. Junction points between the edges of overlapping cells also lead to gaps of low gradient intensity in edges that show on the gradient map. This results in unsatisfactory contours when using a simple greedy contour tracer, as seen in figure 2(c). For comparison purposes, figure 2(d) shows the same two cells as traced by our method.

These problems stem from the fact that a greedy contour tracer sees the gradients directly adjacent to the most recent point in the path (called the *focus point* for the remainder of this paper) as the only factor in determining the probability of the cell contour continuing in a particular direction. The method presented in this paper avoids this problem by determining that probability using both gradient values close to the focus point and an *a priori* probability distribution based on the shape of the traced path so far. The tracing algorithm can be divided into four steps:

1) Calculate a gradient magnitude map
2) Select starting points for traces
3) Trace the contours
4) Verify and cull the detected contours

A. Gradient map and starting point selection

Because SEM images are monochrome, we need not worry about colour information, only brightness. The gradient map, G, is calculated using a simple finite difference approximation with the two convolution operators $(-1, 1)$ and $(-1, 1)^T$. The images we use have very little noise, so smoothing was not necessary, although in higher noise images, smoothing would be required. Any reasonable gradient estimation could be used. The gradient map is thresholded ($t = 80$) to produce a set of high-gradient pixels. The value of t is not critical but should be low enough to produce a starting point on each cell boundary. We use all pixels with $|G| > t$ as potential starting points, resulting in a single cell possibly being traced multiple times, but maximising the chance of detecting each cell. A post-process is used to group multiply detected cells.

B. Tracing contours

The tracing problem is formulated in a simple Bayesian tracking framework. We wish to locate the set of all unobscured blood cell contours $\{C_i\}$, where each $C_i = \mathbf{x}_{0:N_i} = \{\mathbf{x}_0..\mathbf{x}_{N_i}\}$ consists of a set of image points, $\mathbf{x}_j$, forming a closed contour. Given a starting point $\mathbf{x}_0$, the problem is to iteratively build the contour C_i until a loop is formed. At each stage in the decision process, we wish to find the point $\mathbf{x}_{j+1}$ that maximises

$$p(\mathbf{x}_{j+1}|\mathbf{x}_{0:j}, G) = \frac{p(G|\mathbf{x}_{0:j+1})p(\mathbf{x}_{j+1}|\mathbf{x}_{0:j})}{p(x_{o:j}|G)}.$$

We make several simplifying assumptions: the sequence $\mathbf{x}_{0:j}$ is fixed and therefore can be factored out of the first term in the numerator, and the denominator becomes a constant; the second term is estimated locally rather than globally to allow for deviations in the local shape of the cell. These assumptions leave us with the following:

$$p(\mathbf{x}_{j+1}|\mathbf{x}_{0:j}, G) \sim p(G|\mathbf{x}_{j+1})p(\mathbf{x}_{j+1}|\mathbf{x}_{(j-n_l):j})$$

where n_l is the size of the history window. We recast this problem so that we are seeking the direction in which to make the next step based on the current direction of motion. That is, choose θ to maximise:

$$p(\theta|\mathbf{x}_{0:j}, G) \sim p(G|\theta)p(\theta|\mathbf{x}_{(j-n_l):j}). \quad (1)$$

The first term in Equation 1 is the data dependent term and the second is the model dependent term.

The SEM image of a blood cell is approximately elliptical, at least in a global sense, but there are local variations that are often far from elliptical, and hence an ellipse model is inappropriate to use in the tracing algorithm. Locally the edge of a blood cell may be modeled as a second degree polynomial except in rare pathological cases. Least squares is used to fit a second degree polynomial to the $\mathbf{x}_{(j-n_l):j}$ points in the history window. The tangent to the polynomial, θ_t, defines the preferred step direction or model prior:

$$p(\theta|\mathbf{x}_{j-n_l:j}) \sim c\cos(\theta - \theta_t), \quad (2)$$

where c is an empirically determined constant chosen to prefer a new point that maintains the current direction of curvature: $c = 1.7$ if it lies on the same side of the tangent line as the fitted curve and $c = 1.0$ if it lies on the opposite of the tangent line.

The data dependent term is simply the sum of the gradient magnitude in the direction of θ — in effect a linearisation of the expected model:

$$p(G|\theta) \sim \sum_{m=1}^{l} G(\mathbf{x}_j + m\mathbf{u_i}), \tag{3}$$

where l is the look ahead factor and $l = 5$ in our implementation, $u_i = (\cos\theta, \sin\theta)^T$ is the unit vector in the direction of θ. Figure 3 shows a visualisation of the distribution $p(G|\theta)$ at a junction point in the image.

Combining Equation 2 and Equation 3 gives the posterior distribution of θ (Equation 1). We choose $\hat{\theta}$ as the MAP estimate of Equation 1, from which we choose $\mathbf{x}_{j+1}$:

$$\mathbf{x}_{j+1} = \lfloor \mathbf{x}_j + (\cos\hat{\theta} + 0.5, \sin\hat{\theta} + 0.5)^T \rfloor.$$

The final tracing algorithm is very simple and is given in Algorithm 1. The maximum contour length, L_m, in our experiments is 250 pixels, and was chosen empirically to ensure all valid cells could be found. Strictly speaking, a maximum length is not necessary as other processes would cull any incorrect tracings, but was chosen to terminate a trace that was clearly not a cell. One could embed this process in the machinery of a Kalman filter, or particle filter, using the JetStream [11] method for example, but this has been unnecessary for the practicalities of this problem.

Algorithm 1 The tracing algorithm

```
Input: Gradient G
Input: initial point x_0
Input: maximum length L_m
Output: contour C_i
θ = G(x_0)^⊥
x_1 = ⌊x_0 + (cos θ + 0.5, sin θ + 0.5)^T⌋
C_i = {x_0, x_1}
j = 1
while (C_i not closed and j < L_m)
        θ = arg max p(θ|x_{0:j}, G)
        x_{j+1} = ⌊x_j + (cos θ + 0.5, sin θ + 0.5)^T⌋
        C_i = C_i ∪ x_{j+1}
        j = j + 1
endwhile
```

Because of multiple starting points being generated on the same cell boundary, contours are sometimes detected more than once. Therefore, contours whose centroids are close to a previously found contour are culled. Noise contours are also sometimes detected and are typically highly non-elliptical, due to multiple cells being detected as one, or a single cell being detected incorrectly due to surface markings. These contours are culled by fitting an ellipse [4] and measuring the root mean squared error (rmse) between the detected contour and the ellipse. Contours that are highly non-elliptical are culled. In our case, an ellipse is culled if rmse $> 1 \times 10^5$ using the algebraic distance of [4] as an error metric.

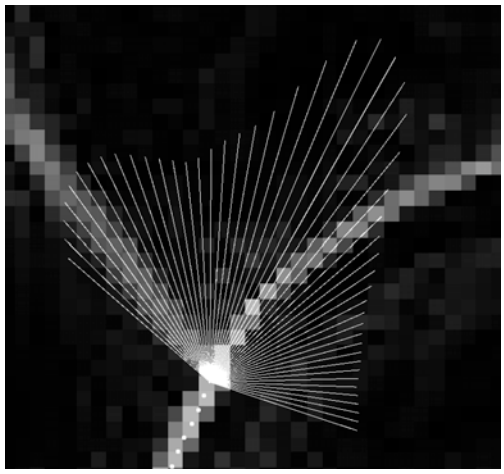

Fig. 3. Visualisation of $P(G|\theta)$ at a junction point. The length of the vector indicates the average of the gradient magnitudes in that direction. Note the local maxima in the directions of the edges.

IV. Results

In evaluating the algorithm, we used a data-set containing 800 SEM images. These images measure 1024 by 768 pixels, gray-scale. For our evaluation, we ran the algorithm on 100 randomly selected images from the data-set. Four representative trace results can be found in figure 4. Figure 4(a) displays the results on an image particularly well suited to this method of segmentation. All cells are situated perpendicular to the electron beam, displaying a flat, ellipse-like surface. Furthermore, the image contains next to no 'shadows' that could distract the trace algorithm. The next (figure 4(b)) is obtained from more or less average quality pictures from the data-set. Still, most of the top cells are found correctly. Figure 4(c) shows the results on an image on the other end of the spectrum: cells are squashed together tightly and are often not lying flat on the surface. This throws off the trace algorithm, and many of the contours found are incorrect.

Results of the trace were superimposed on the original images, and scrutinized by eye. Contours were manually divided into three categories as demonstrated in Figure 5: fully correct; minimally obscured; and incorrect. The 'minimally obscured' category indicates a very slight overlap breaking the contour of this cell. Often, these cells are still correctly detected, as they are visually similar to non-obscured cells. It is reasonable to believe that they can be classified correctly, but because we cannot easily predict how much data the classifier algorithm will need, we categorise them separately. Examples of each category can be seen in figure 5. We also counted the number of contours that were not found, but were nevertheless fully visible. These are false negatives. Results of the evaluation can be found in Table I.

The incorrectly detected contours can be divided into six categories. Figure 6(a), 6(b) and 6(c) are examples of the trace being distracted by gradient data within a cell. This results in contours being smaller than they should be. Figure 6(d) shows cases where the predicted curve was not enough to

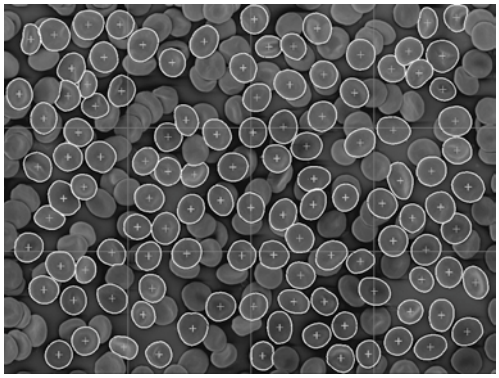

(a) High accuracy

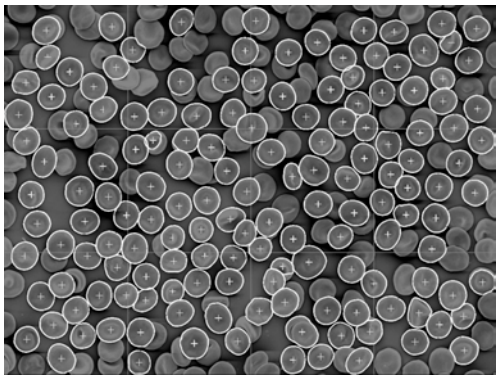

(b) Average accuracy

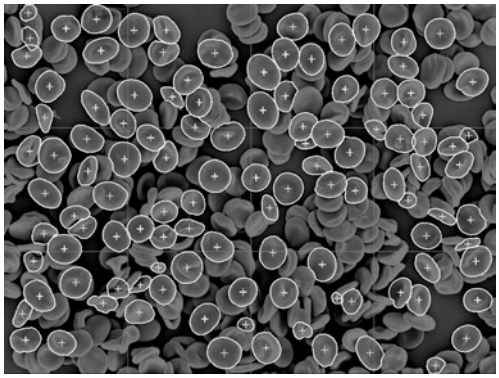

(c) Low accuracy

Fig. 4. Results on three images used in the evaluation

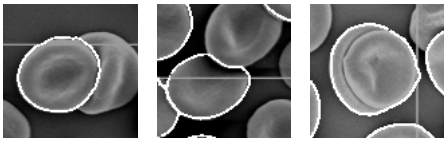

(a) (b) (c)

Fig. 5. Contour categories. (a) Fully correct, (b) minimally obscured, (c) incorrect

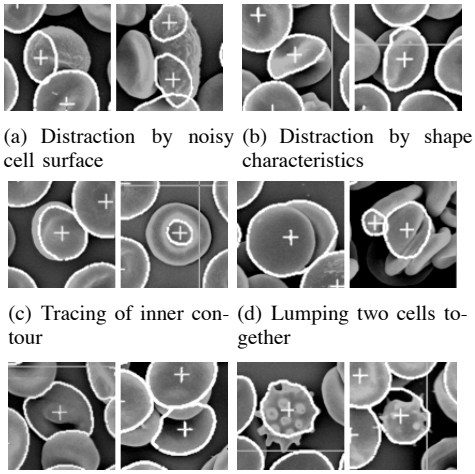

(a) Distraction by noisy cell surface (b) Distraction by shape characteristics

(c) Tracing of inner contour (d) Lumping two cells together

(e) Tracing highly obscured cells (f) Distraction by abrupt changes in curvature

Fig. 6. The six main categories of incorrectly found contours

Category	Total for 100 images	Mean(S.D.) for each image	Percentage
Detected contours	11704	117.7(18.3)	100.0%
Correct	9655	96.6(14.8)	82.5%
Minimally obscured	1548	15.5(6.4)	13.2%
Incorrec	501	5.0(4.1)	4.3%
False negatives	70	0.7(0.8)	0.6%

TABLE I

RESULTS ON 100 RANDOMLY SELECTED IMAGES. COLUMN 2 SHOWS THE TOTAL NUMBER OF BLOOD CELLS DETECTED IN EACH CATEGORY ACROSS ALL IMAGES. COLUMN 3 IS THE MEAN NUMBER AND STANDARD DEVIATION OF DETECTED CONTOURS FROM EACH IMAGE, AND COLUMN 4 IS THE PERCENTAGE OF EACH CATEGORY COMPARED TO THE TOTAL NUMBER FOUND.

guide the algorithm around a junction point. The two cells are still lumped together as one. The curve prediction can also fail on cells with abrupt changes in curvature, as seen in figure 6(f). While many of these cells are traced correctly, some curvature changes can throw off the trace by enough to result in a wrong contour. The few false negatives are mainly caused by one of two things. The first, again, is a contour with abrupt changes in curvature (figure 7(a)). The second is a contour that passes through a low contrast area (figures 7(b) and 7(c)). These areas result in low gradient values, and the trace is prone to be distracted by nearby high-gradient contour lines of other cells.

V. CONCLUSION AND FUTURE WORK

The results of our evaluation are promising. Of all contours detected, 95.7% are either unobscured or minimally obscured cells with only a 0.6% false negative rate, and a 4.3% false positive rate. However, the usability of the minimally obscured contours will have to be tested in further stages of the project. These apparently good results are, for a large part, due to specific properties of our problem domain. The SEM produces low-noise, high-contrast images, and red blood cells shapes are smooth enough to have some benefit from a prediction using

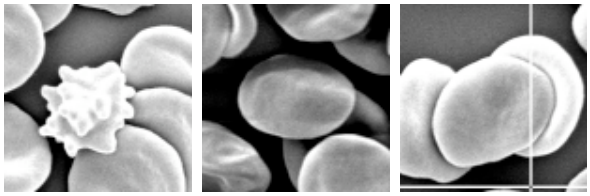

(a) Non-smooth contours (b) Dark areas (c) Bright areas

Fig. 7. Main causes for false negatives

a local curve. The errors that do occur are often caused by noise within a cell. The trace "wants" to curve in towards the cell, and a noisy cell surface will create enough gradient information for it to do so. This kind of error makes up the majority of the incorrectly traced contours.

Choosing the direction vector based on average of the gradient magnitude information over a number of steps in that direction causes a slight smoothing of the contour. Details smaller than the number of averaged steps may be lost. Future work will focus on classification of the resulting contours, as well as improving culling measures to filter out more incorrect contours. In order to find more than just the top-most cells, we envision a multi-pass approach, first marking the top contours, and then repeating the process to find lines that end at previously found cells.

REFERENCES

[1] PJH Bronkorsta, M.J.T. Reinders, E.A. Hendriks, J. Grimbergen, RM Heethaar, and GJ Brakenhoff. On-line detection of red blood cell shape using deformable templates. *Pattern Recognition Letters*, 21(5):413–424, 2000.

[2] S. Chien. Red Cell Deformability and its Relevance to Blood Flow. *Annual Review of Physiology*, 49(1):177–192, 1987.

[3] C. Di Ruberto, A. Dempster, S. Khan, and B. Jarra. Analysis of infected blood cell images using morphological operators. *Image and Vision Computing*, 20(2):133–146, 2002.

[4] Andrew Fitzgibbon, Maurizio Pilu, and Robert B. Fisher. Direct least square fitting of ellipses. *IEEE Transactions on Pattern Analysis and Machine Intelligence*, 21(5), 1999.

[5] A. Garrido and N.P. de la Blanca. Applying deformable templates for cell image segmentation. *Pattern Recognition*, 33(5):821–832, 2000.

[6] T. Jiang and F. Yang. An evolutionary tabu search for cell image segmentation. *Systems, Man and Cybernetics, Part B, IEEE Transactions on*, 32(5):675–678, 2002.

[7] N. Malpica, C.O. de Solorzano, J.J. Vaquero, A. Santos, I. Vallcorba, J.M. Garcia-Sagredo, and F. del Pozo. Applying watershed algorithms to the segmentation of clustered nuclei. *Cytometry*, 28(4):289–297, 1997.

[8] T. McInerney and D. Terzopoulos. T-snakes: Topology adaptive snakes. *Medical Image Analysis*, 4(2):73–91, 2000.

[9] B. Nilsson and A. Heyden. Segmentation of dense leukocyte clusters. *Proceedings of the IEEE Workshop on Mathematical Methods in Biomedical Image Analysis (MMBIA'01)*, 2001.

[10] J. Park and J.M. Keller. Snakes on the watershed. *IEEE Transactions on Pattern Analysis and Machine Intelligence*, 23(10):1201–1205, 2001. ISSN 0162-8828. doi: http://doi.ieeecomputersociety.org/10.1109/34.954609.

[11] P. Perez, A. Blake, and M. Gangnet. Jetstream: Probabilistic contour extraction with particles. *Proc. of International Conference on Computer Vision*, pages 424–531, 2001.

[12] LO Simpson. Blood from healthy animals and humans contains nondiscocytic erythrocytes. *Br J Haematol*, 73 (4):561–4, 1989.

[13] LO Simpson. The effects of saline solutions on red cell shape: a scanning-electron-microscope-based study. *Br J Haematol*, 85(4):832–4, 1993.

[14] LO Simpson. Capillary blood flow; red cell shape: Implications for clinical haematology. *Advances in Physiological Fluid Dynamics. New Delhi, Narosa*, pages 42–47, 1995.

[15] LO Simpson and GP Herbison. The results from red cell shape analyses of blood samples from members of myalgic encephalomyelitis organisations in 4 countries. *J Orthomol Med*, 12:221–226, 1997.

[16] LO Simpson, BI Shand, and RJ Olds. Blood rheology and myalgic encephalomyelitis: a pilot study. *Pathology*, 18(2):190–2, 1986.

[17] LO Simpson, BI Shand, RJ Olds, PW Larking, and MJ Arnott. Red cell and hemorheological changes in multiple sclerosis. *Pathology*, 19(1):51–5, 1987.