Med Biol Eng Comput (2006) 44: 427–436
DOI 10.1007/s11517-006-0044-2

ORIGINAL ARTICLE

Nicholas E. Ross · Charles J. Pritchard
David M. Rubin · Adriano G. Dusé

Automated image processing method for the diagnosis and classification of malaria on thin blood smears

Received: 29 July 2005 / Accepted: 12 March 2006 / Published online: 8 April 2006

Abstract Malaria is a serious global health problem, and rapid, accurate diagnosis is required to control the disease. An image processing algorithm to automate the diagnosis of malaria on thin blood smears is developed. The image classification system is designed to positively identify malaria parasites present in thin blood smears, and differentiate the species of malaria. Images are acquired using a charge-coupled device camera connected to a light microscope. Morphological and novel threshold selection techniques are used to identify erythrocytes (red blood cells) and possible parasites present on microscopic slides. Image features based on colour, texture and the geometry of the cells and parasites are generated, as well as features that make use of a priori knowledge of the classification problem and mimic features used by human technicians. A two-stage tree classifier using backpropogation feedforward neural networks distinguishes between true and false positives, and then diagnoses the species (*Plasmodium falciparum*, *P. vivax*, *P. ovale* or *P. malariae*) of the infection. Malaria samples obtained from the Department of Clinical Microbiology and Infectious Diseases at the University of the Witwatersrand Medical School are used for training and testing of the system. Infected erythrocytes are positively identified with a sensitivity of 85% and a positive predictive value (PPV) of 81%, which makes the method highly sensitive at diagnosing a complete sample provided many views are analysed. Species were correctly determined for 11 out of 15 samples.

Nicholas E. Ross conducted research while at the University of the Witwatersrand School of Electrical and Information Engineering

N. E. Ross (✉)
Department of Engineering, Cambridge University,
Trumpington Street, Cambridge CB2 1PZ, UK
E-mail: ner25@cam.ac.uk

C. J. Pritchard · D. M. Rubin
School of Electrical and Information Engineering,
University of the Witwatersrand, Private Bag 3,
Johannesburg, Wits 2050, South Africa

A. G. Dusé
Department of Clinical Microbiology and Infectious Diseases,
University of the Witwatersrand Medical School,
7 York Road, Parktown 2193, South Africa

1 Introduction

Malaria is a serious global health problem, causing between 1.5 and 2.7 million deaths every year [8]. The worldwide annual economic burden of malaria, calculated to include spending on prevention and treatment as well as loss of productivity due to illness, was estimated at US$800 million in 1995 [5]. Rapid and accurate diagnosis which facilitates prompt treatment is an essential requirement to control malaria [8].

A number of new methods have been developed in recent years for the diagnosis of malaria. These include the use of fluorescent microscopy, rapid antigen detection methods and polymerase chain reaction (PCR)-based techniques that detect specific nucleic acid sequences [8]. Despite these advances, malaria diagnosis by means of light microscopy remains the most widely and commonly used method [1].

Microscopic diagnosis entails examining thick and thin blood smears for the presence of *Plasmodia*. It is the most efficient and reliable diagnostic technique and is very sensitive and highly specific [14]. The advantages of microscopy are numerous: it is possible to differentiate between species, quantify parasitaemia and observe asexual stages of the parasites [1, 14]. Low material costs [11] mean that the marginal costs of tests are very low.

Unfortunately, there are also disadvantages to the method: substantial costs are incurred purchasing and maintaining microscopes and training technicians [11], the technique is labour intensive and time-consuming [1] and the accuracy of the final diagnosis ultimately depends on the skill and experience of the technician and the time spent studying each slide [14]. Compared

with expert microscopy, standard laboratory microscopy has a sensitivity of approximately 90%, a figure which drops dramatically in the field [1]. Variable smear quality and slide degeneration with time [12] are also problematic.

The aim of this research is to improve upon microscopy by removing its most serious limitation: reliance on the performance of a human operator for diagnostic accuracy. This is pursued by developing a digital image processing system to automate the examination of blood smears. The system must provide a positive or negative diagnosis of malaria with similar sensitivity and specificity to normal microscopy; and differentiate parasites by species.

2 Materials and methods

Malaria samples were obtained from the Department of Clinical Microbiology and Infectious Diseases at the University of the Witwatersrand Medical School. Due to the prevalence of *Plasmodium falciparum* in southern Africa, the samples are almost all of that species, so teaching slides of the other species (*P. vivax*, *P. ovale* and *P. malariae*) are also used.

The slides are routinely stained using the quick Giemsa protocol, which is similar to but faster than the conventional Giemsa stain. All slides had already been examined and verified by expert microscopists, who had given a species-specific diagnosis.

Thin blood smears of the samples, whereby the material is fixed and parasites can be visualised inside the cells [14], are used for the automated diagnosis. The advantage of thin films is that there is no loss of parasites during staining; the erythrocytes (red blood cells) are fixed; artifacts are uncommon and, most importantly, parasite morphology is not distorted [2].

Slides are examined under oil-immersion using a light microscope with 1,000× magnification, and images are captured using a charge-coupled device (CCD) camera connected to the microscope. In total 12 patient samples and two teaching samples of *P. falciparum*, one patient sample and one teaching sample of *P. vivax* and one teaching sample each of *P. ovale* and *P. malariae* are examined.

3 Algorithm development

The automated image processing algorithm is designed to diagnose malaria in much the same way as a human operator performing microscopy. To do this, the algorithm finds and identifies erythrocytes and malaria parasites present in a microscopic field of a thin blood smear. Based on the parasites and erythrocytes found, the program makes a diagnosis as to whether or not malaria is present, and if present, it determines the species of the infection. Low parasitaemia does not imply a good outcome for the disease [15], so the system must have a high degree of sensitivity. It must also have good specificity to be useful as a clinical tool [11].

The algorithm design is essentially an image classification problem, and thus takes the form of a standard pattern recognition and classification system. It consists of four stages: image acquisition, pre-processing, feature generation and classification [17], and the performance of the system is then evaluated. A morphological method used to identify malaria in Giemsa-stained blood slides [3] is used as a starting point for the algorithm, from which many of the pre-processing and image segmentation steps are derived.

3.1 Image acquisition

Images are acquired using a 3.34 megapixel Nikon Coolpix 995 digital camera (Nikon Corporation, Tokyo, Japan). The camera, using full 4× optical zoom, is connected to a light microscope with 1,000× magnification. Malaria slides are examined under oil immersion. Images are captured in the JPEG format at the maximum resolution of the camera, 2,048×1,536 pixels.

3.2 Pre-processing

The purpose of the pre-processing stage is to remove unwanted effects such as noise from the image, and transform or adjust the image as necessary for further processing. The resolution of the image is reduced by a factor of four to 512×384 to speed up performance of the system.

The complemented, green component of the true colour original is primarily used by the system, since it has the least noise [3] and the parasites, which stain a purple colour, are most visible.

Following Di Ruberto et al. [3], the image is filtered using a 5×5 median filter, followed by a morphological area closing filter using a disk-shaped structuring element (SE) of radius 6 pixels. The morphological filter removes some parasite detail from the image, and so the morphologically-filtered image is only used for functions where parasite detail is not important, such as erythrocyte segmentation. In all other cases, the median-filtered image is used.

3.3 Feature generation

The goal of the feature generation stage is to develop a set of quantitative features from which the objects present in the image can be classified. Objects present in thin blood smears include erythrocytes, a variety of white blood cells, artifacts in the blood, and parasites (in infected samples). These must be identified and segmented from the background before differentiating features can be generated and used to classify the objects. The feature generation procedure therefore requires

image analysis, image segmentation and feature calculation (Fig. 1).

3.3.1 *Image analysis*

In order to use morphological methods for image segmentation, the shape and size of the objects in the image must be known. The size and eccentricity of the erythrocytes are also required for the calculation of some feature values (as these can be indicative of infection). The shape of the objects (circular erythrocytes) is known a priori, but the image must be analysed to determine the size distribution of objects in the image and to find the average eccentricity of erythrocytes present.

A pattern spectrum showing the size distribution of objects in a sample image (Fig. 2) can be calculated using granulometry [3]. Granulometry is computed from the difference in morphological openings with SE's of increasing size [9]. In this case, a disk-shaped SE with increasing radius is used, so the pattern spectrum will indicate the radii that most commonly occur in the image.

From the sample pattern spectrum (Fig. 3), it is possible to identify the principal object radius present in the sample image (~30 pixels). This corresponds to the mean external radius of erythrocytes in the image. A peak in the pattern spectrum is also commonly found for the internal radius (~12 pixels) of the erythrocytes, which are not prevalent in this sample, and for parasites (~7 pixels), as is the case in Fig. 3.

The external radius, which is the dimension of interest, and an estimate of the standard deviation of the external radius, are calculated from the principal mode of the pattern spectrum (Fig. 3) as 30 and 3 pixels (2.56 pixels rounded to the nearest integer), respectively.

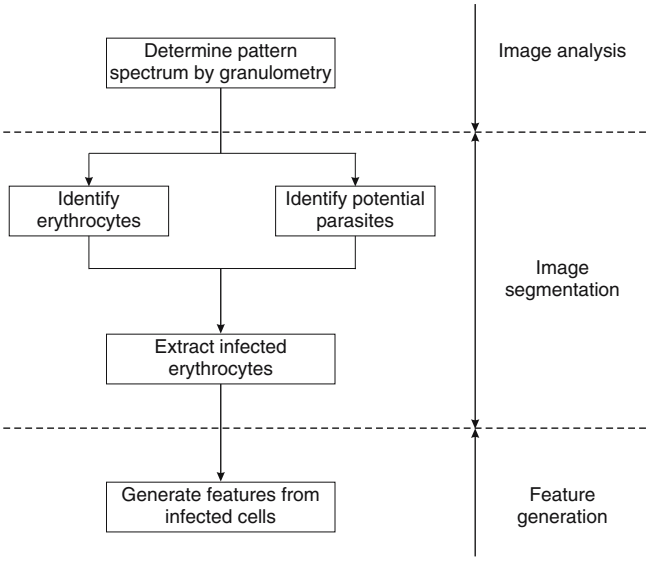

Fig. 1 Feature generation stage

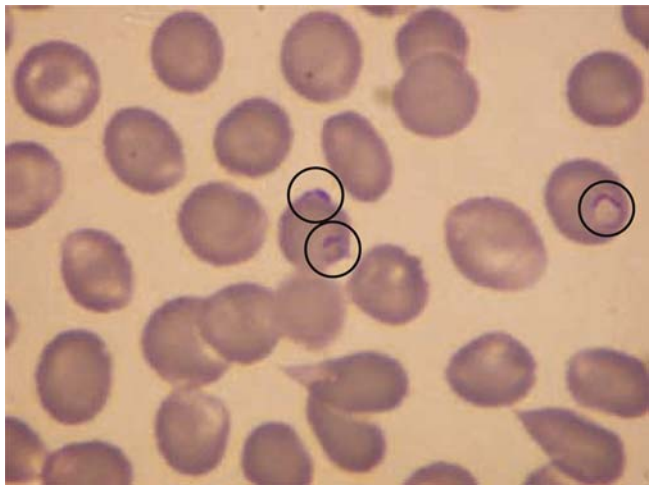

Fig. 2 Sample image (*Plasmodium falciparum* parasites *circled*)

The average eccentricity of the erythrocytes is determined from a binary image (obtained by thresholding following the method of Otsu [13]). The eccentricity is expressed as a value between 0 and 1, with 0 indicating a circle and 1 a line segment. Free-standing erythrocytes are differentiated from overlapping cells by their area. The area of a circle, with a radius equal to the sum of the mean and the standard deviation of the erythrocyte radius determined by granulometry, is used as the threshold. Any cell with an area less than that is adjudged to be free-standing. Eccentricity is calculated using free-standing erythrocytes only.

Granulometry is a very computationally intensive routine due to its use of a series of morphological openings. In order to reduce the computational time of the algorithm, it is assumed that the size distribution of erythrocytes is constant throughout a sample. Image analysis is only performed on four images from every sample (which typically consist of about twenty images),

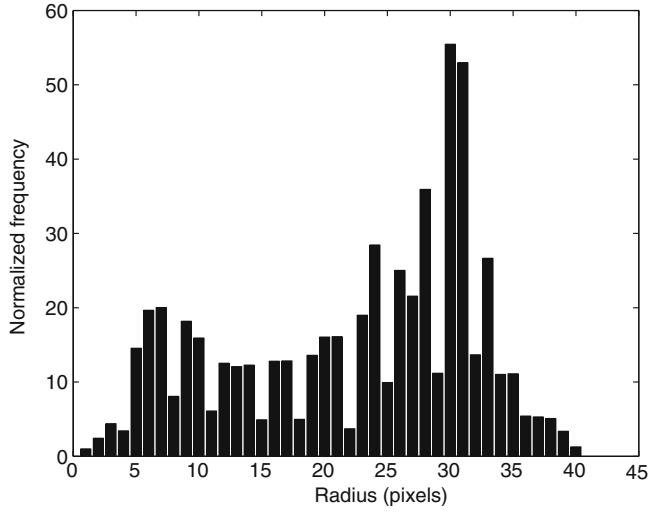

Fig. 3 Sample pattern spectrum determined by granulometry

and an average pattern spectrum is used to determine the sample magnitude and standard deviation of the erythrocyte radii. The average cell eccentricity of the sample is also taken as the average from these four images. Ideally, should computational resources allow it, image analysis should be performed on every image considered.

3.3.2 Image segmentation

The next stage of the algorithm, image segmentation, identifies and segments potential parasites and erythrocytes from the image background. To extract the infected erythrocytes, it is first necessary to identify them from the combination of parasites and erythrocytes in the image, and then segment them from the background (Fig. 1).

Although techniques such as edge detection and the watershed algorithm are also commonly used for image segmentation, this algorithm relies primarily on thresholding. The key to successfully segmenting an image using thresholding is threshold selection. A method to find thresholds based on the image histogram was developed for this algorithm.

The histogram of the complemented, green component of the sample image (Fig. 4) is a bimodal distribution typical of all the images considered. The principle mode is due to the greyscale intensities of the image background, and the second mode is due to those of the erythrocytes in the image. Two threshold levels need to be determined from the histogram: one for erythrocytes, and one for parasites.

The first threshold is selected to separate the erythrocytes from the background of the image. This essentially means separating the two modes of the image histogram. The threshold level is automatically selected using a method that maximises the separability of the resultant classes of the grey-level histogram [13].

The resulting thresholded binary mask of erythrocytes (Fig. 5) then has all holes with an equivalent radius less than an empirically determined 62.7% of the average erythrocyte radius removed (Fig. 5). A morphological opening using a disk-shaped SE with a radius 40% of the mean erythrocyte radius is applied to smooth the objects in the image, and any objects with an equivalent radius of less than half the mean erythrocyte radius are removed (Fig. 5). The problem with this binary image of erythrocytes is that clusters of cells are not separated.

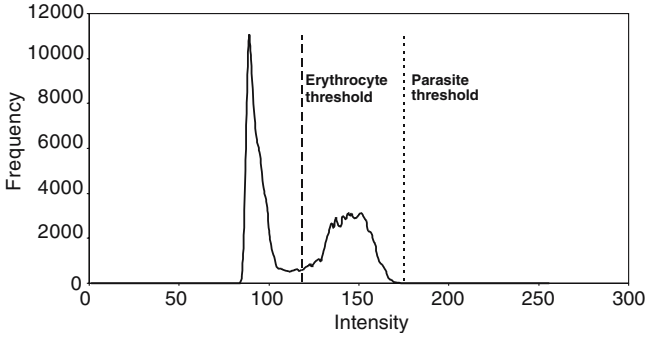

Fig. 4 Histogram of sample image

The next step is to select the second threshold to find parasites present in the image. A global threshold level, taking the threshold as the first local minimum in the histogram after the mode due to erythrocytes (Fig. 4), is not sensitive enough (Fig. 6). This is a common problem experienced with global threshold selection, caused by inconsistent intensities in the image.

The solution is to find local threshold levels. The erythrocytes, having already been identified, provide excellent image regions in which to find these, especially since valid parasites are only found inside erythrocytes. The threshold is then found by taking the first minimum after the principal mode of the histogram incorporating only the erythrocytes.

While this method has greater sensitivity, it is at the expense of a reduced specificity (Fig. 6). There are also cases in this study, particularly with *P. ovale*, where the global threshold is able to detect parasites that are missed by the local thresholds. This is due to colourisation of the infected cells, which shifts the principle mode of the local histograms of the affected cells. For this reason, both local and global thresholds are used, and the union of the two binary images is used as the parasite marker image.

Invalid objects in the marker image (objects detected with the global threshold that lie outside any erythrocyte) are removed by taking the intersection of the parasite marker image with the binary mask of erythrocytes. The erythrocyte mask is dilated first, to ensure that 'blister' forms of the parasites, that appear to bulge out of the edge of the cells, are not removed.

Other artifacts in the blood containing nucleic acid, particularly white blood cell nuclei, are also detected by this thresholding. They are removed by excluding all objects greater than an empirically determined size (chosen to exclude all objects greater than the largest trophozoite that one would expect to find.)

The infected cells are identified by morphologically reconstructing the erythrocyte mask with the valid parasite marker (Fig. 7). Binary reconstruction simply involves extracting the connected components of an image (the mask) that are marked by another image (the marker). Where cells are clustered together, if an infected cell forms part of the group, then the entire aggregation is reconstructed.

To separate these clusters so that the infected cell can be isolated and extracted, a modification of the morphological technique used in Di Ruberto et al. [3] is used. A morphological opening filter, using a disk-shaped SE with radius equal to the mean erythrocyte radius less the standard deviation, is applied to the greyscale, morphologically filtered green component of the image to remove any objects smaller than an erythrocyte. The morphological gradient—the difference between a dilation and erosion of the image—is then calculated using a diamond-shaped SE with unity length.

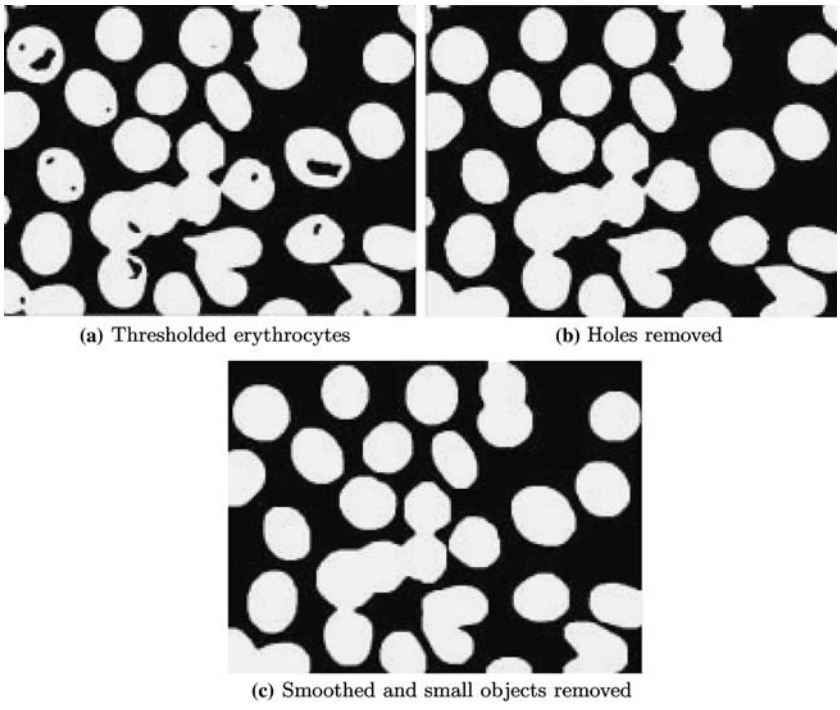

Fig. 5 Binary mask of erythrocytes

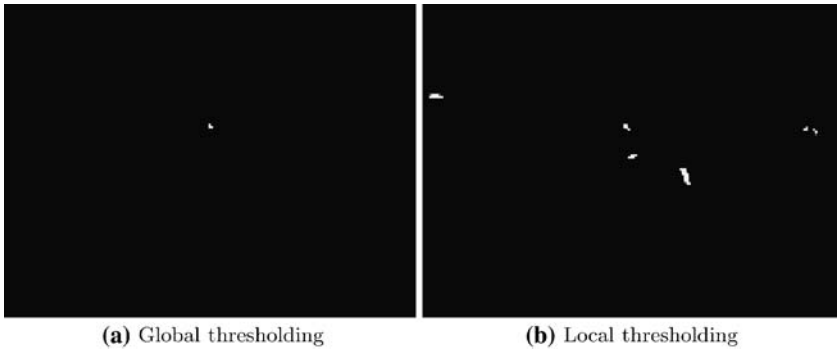

Fig. 6 Binary parasite markers

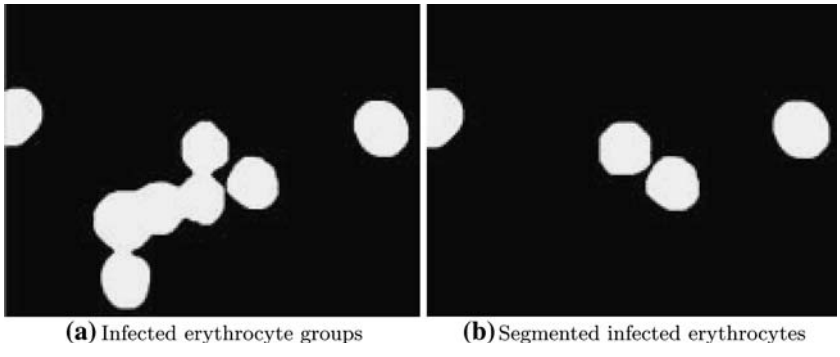

Fig. 7 Morphologically reconstructed erythrocyte mask

The segmentation method is applied to each object in the reconstructed binary image of erythrocytes individually. Those objects that do not exceed the area of a circle with radius equal to the mean erythrocyte radius plus the standard deviation are regarded as being single cells, and are unmodified.

Unlike the method in Di Ruberto et al. [3], where the morphological gradients are used to generate marker images for the watershed algorithm, the objects deemed to be overlapping erythrocytes are segmented as follows. First, the intersection of the morphological gradient image and the dilated cell cluster is taken. This image is

then transformed to a binary image by thresholding any value greater than zero. A series of morphological operations, namely a closing operation, thinning, and spur-removal are then applied to generate a contour of the segmented erythrocytes. The contours are filled, and the segmented mask is again reconstructed with the valid parasite marker image to result in a segmented mask of infected cells (Fig. 7).

The erythrocytes that have been identified as possibly infected are then extracted from the image (Fig. 8) and passed to the next stage of the algorithm for feature generation. The binary mask of the erythrocyte, as well as a binary mask (obtained by local threshold selection based on the image histogram as detailed above) of parasite-like objects present in the cell, are also passed to the next stage (Fig. 9).

3.3.3 Feature generation

The purpose of feature generation is to compute new variables from the image array that concentrate information to separate classes [16]. The classifier has two functions: it must determine whether or not a detected cell is truly positive for malaria, and what the species of the infection is. Features are created with these functions in mind. They must provide information with which the classifier can distinguish between parasites and other artifacts in the blood, and information which will allow the classifier to differentiate between parasites of different species. The final performance of the classifier directly depends on the success of the feature generation stage.

Two sets of features are developed. The first set is based on image characteristics that have been used previously in biological cell classifiers, which include geometric features (shape and size), colour attributes and grey-level textures (6, 10]. Texture features are generated from the greyscale image matrices of the red, green and blue components, as well as the intensity component from the hue-saturation-intensity image space. First order features, based on the image histograms, and second order features, based on co-occurrence and run-length matrices (as described in Theodoris and Koutroumbas [16]) are used.

Colour features are derived from the red, green, blue, hue and saturation components, and include measures such as the peak intensity, average intensity, skewness, kurtosis and entropy of the component histograms.

Geometric features include the roundness ratio and bending energy of binary masks [16]; boundary analysis performed on chain codes [10]; and size information such as the area and equivalent radius.

It is advantageous to apply expert, a priori knowledge to a classification problem [6]. This is done with the second set of features, where measures of parasite and infected erythrocyte morphology that are commonly used by technicians for manual microscopic diagnosis are used. It is desirable to focus on these features, because it is already known that they are able to differentiate between species of malaria.

These features, based both on the parasites and the infected erythrocytes, include: the relative size of the infected erythrocytes; the relative eccentricity of the infected erythrocytes; smoothness of the cell margin (i.e. crenellated or not); the relative colour of infected erythrocytes (in some cases decolourised); texture information of infected erythrocytes (i.e. the presence of stippling—Schüffner's dots—or Maurer's dots); the number of parasites per erythrocyte; the number of chromatin dots per parasite; morphology of the rings (large and coarse, or small and fine); the position of the

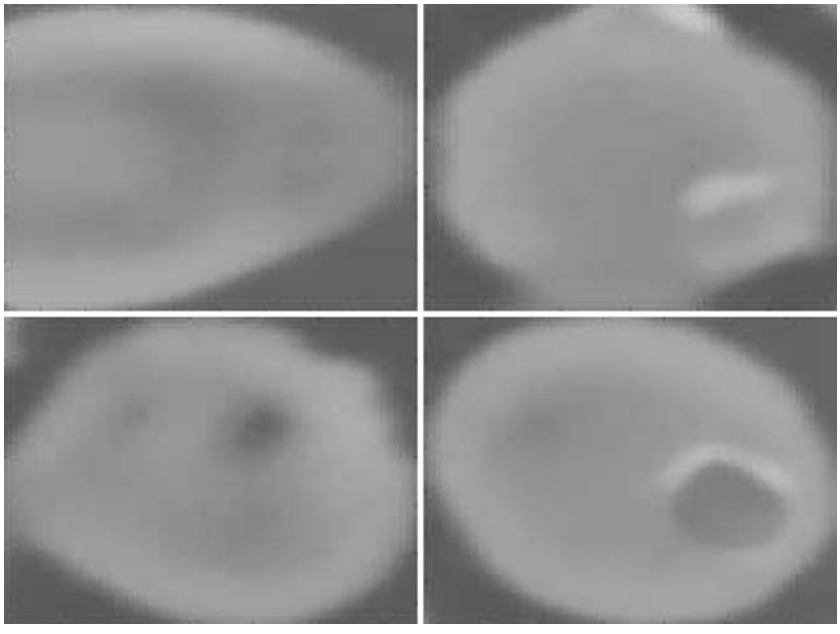

Fig. 8 Extracted infected erythrocytes

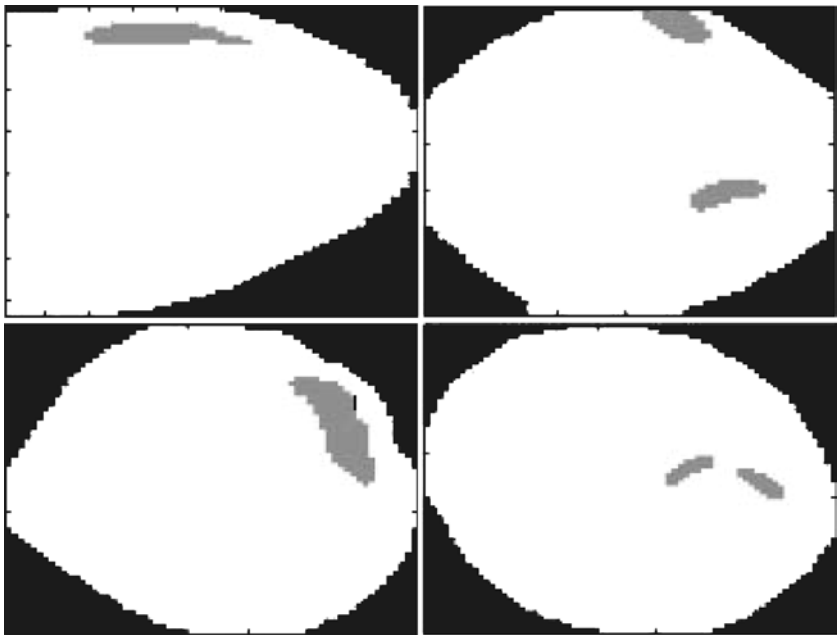

Fig. 9 Erythrocyte masks with parasite masks (*grey*) superimposed

parasite in the cell (i.e. accolé forms); and the eccentricity, or the ratio of major axis over minor axis, of the parasite (i.e. band forms) [11]. The size and shape of the parasite [18] is also a differentiating factor.

Not all of these features can be quantitatively assessed using an automated method, particularly given the limitations of the image segmentation stage of the algorithm. Those that can include measures of the relative size and eccentricity of infected erythrocytes (relative to the averages found during image analysis), the relative colour of infected cells, the number of parasites per cell, the texture of infected cells, the number of chromatin dots per parasite, the size and shape of the parasite and the distance of the parasite from the edge of the cell.

3.4 Classification

The final classification of an erythrocyte as infected with malaria or not, and if so, the species of the parasite, falls to the classifier. The classifier takes the form of a two-stage tree classifier (Fig. 10), with an infection classified as positive or negative at the first node, and the species assigned at the second node.

There are many advantages to using a tree structure for the classifier: it is more flexible, allowing only the most relevant features to be used at each stage, accuracy is higher with fewer features, and the division of classes follows a more natural, logical form [10].

This is especially true of this problem, which lends itself to a tree classifier since the two classifications have significant differences and utilise different features. By requiring fewer features, the dimensionality problem is lessened. The dimensionality problem is that with a fixed sample size the classification accuracy can decrease when the number of features is increased [7]. This means that the more features that are used, the more training data required to give reliable results [4].

The design of a tree classifier follows three steps: the design of a tree structure (which has already been assigned), the selection of features to be used at every node, and the choice of decision rule at each node [10]. The same type of classifier is used at both nodes, although this need not necessarily be the case.

Biological cell classification problems have previously used classifiers utilising the quadratic decision rule [10], minimum Bayes error [6] or scoring systems [17]. Taking into account the fact that there is no guarantee that the classes are linearly separable, backpropogation feedforward (BFF) neural networks are used here.

Backpropogation feedforward networks are designed by specifying an architecture, and then using a training rule to train the network using training data to set the

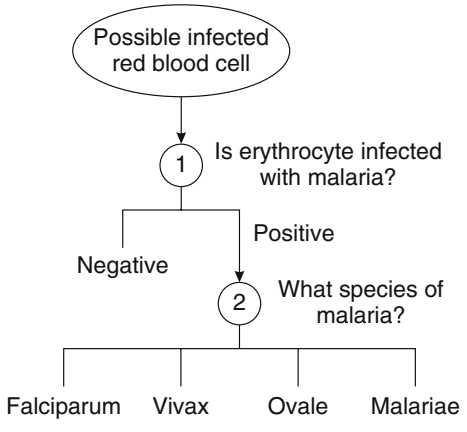

Fig. 10 Structure of the tree classifier

synaptic weights. An important consideration is to ensure the generalisation of the network, i.e. its ability to classify data that were not presented to it during training. The generalisation of the network can be compromised by overfitting, which occurs when there are too many free parameters present in the network, causing it to adapt specifically to the training data used. To avoid this problem, the smallest possible network that can differentiate between the classes must be used. Overtraining also reduces generalisation by adapting the synaptic weights to the particulars of training data [16].

The first classifier is an 8-5-1 three layer network (two hidden layers and one output layer), with tan-sigmoid transfer functions. Three layer perceptron networks are theoretically able to classify any combination of features [16], and the BFF network is a differentiable approximation of a perceptron network.

The output layer is determined by the number of output classes (two), while the hidden layers were designed by trial and error. The second classifier is a 10-10-4 three layer network. There are four output classes, and the hidden layers were again designed by trial and error.

The features selected for the first classifier are those that describe the colour and texture of the possible parasites. The features used by microscopists to differentiate malaria species are selected for the second classifier. Principal component analysis is used to remove any elements of the feature vectors that account for less than 0.1% variation in the data set. This decreases the first feature vector length from 75 to 37, and the second from 117 to 38.

The networks were trained using the resilient backpropogation training algorithm, commonly used in pattern classification problems. The training goal was to minimise squared errors, and training was stopped when the error of a validation set increased. This was done to avoid overtraining.

Due to the poor availability of samples of *P. vivax*, *P. ovale* and *P. malariae*, half the slides from these samples had to be used for training. Two samples of *P. falciparum* were also used for training, giving a total training set of 350 images (120 *P. falciparum*, 120 *P. vivax*, 70 *P. ovale* and 40 *P. malariae*) with approximately 950 possiblly infected erythrocytes. Since different slides from the same samples of *P. vivax*, *P. ovale* and *P. malariae* are used for training and assessment, only the performance of *P. falciparum* samples that were not used for training are an indication of the generalisation of the network.

4 Results

The performance of three stages of the algorithm is analysed: the image segmentation stage that detects possible infected blood cells, the first classifier stage that confirms infections and the second classifier stage that differentiates species. Two measures of algorithm performance and accuracy are used: sensitivity, the ability of the algorithm to detect a parasite present; and positive predictive value (PPV), the success of the algorithm at excluding non-infected cells . These values are expressed in terms of true positives (TP), false positives (FP) and false negatives (FN):

$$\text{Sensitivity} = \frac{\text{TP}}{\text{TP} + \text{FN}} \tag{1}$$

$$\text{PPV} = \frac{\text{TP}}{\text{TP} + \text{FP}} \tag{2}$$

The results of the algorithm with regards to segmenting possible parasites from the images, and after the first neural network classifier, as compared to a human operator, are summarised in Table 1 (their outputs from processing the sample image can be seen in Fig. 11a, b, respectively). Note that the performance of the image segmentation stage indicates initial parasite detection for all samples, while the training samples were excluded from the assessment of the first classifier.

The presence of infected erythrocytes in images of a blood smear can then be used to determine, using a suitable decision rule, whether the sample is positive. This decision will require that the number of detected parasites sufficiently exceeds the amount of false positives that will normally be found in negative samples. To compare the system to other diagnostic methods, the success of the system at correctly identifying full samples and not individual infected erythrocytes, must be ascertained.

A Bayesian argument can be used to estimate the the sensitivity and positive predictive value of the algorithm on a full sample. The sensitivity of the system after classification (as shown in Table 1) can be represented as the conditional probability of measuring a cell as infected (m), given that it is infected (p), i.e. $\mathbb{P}(m|p)$.

Assuming that it is necessary and sufficient to find a single infected cell to diagnose the sample as positive, then the probability of diagnosing a positive sample as positive [or the sample sensitivity (SS)] is

$$\text{SS} = 1 - \mathbb{P}(\overline{m}|p)^N \tag{3}$$

where the probability of misdiagnosing an infected cell is given by

$$\mathbb{P}(\overline{m}|p) = 1 - \mathbb{P}(m|p) \tag{4}$$

Table 1 Results for detecting and confirming individual erythrocyte infections

Measure	Initial parasite detection (all samples)	After first classification (excludes training samples)
True positives	905	481
False positives	1,378	114
False negatives	78	84
Sensitivity (%)	92.07	85.13
Positive predictive value (PPV) (%)	39.64	80.84

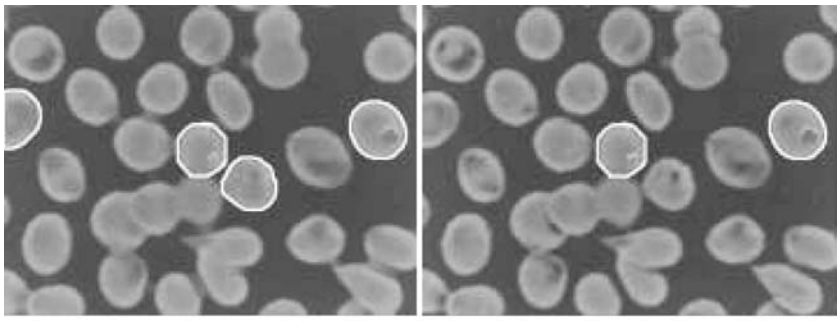

(a) Initial detection **(b)** After first neural network classification

Fig. 11 Infected erythrocytes (*circled*) detected by the algorithm in the sample image

and N is the number of infected cells examined by the algorithm.

Clearly from (3), as $N \rightarrow \infty$, SS $\rightarrow$ 1. Since the number of infected cells examined is proportional to the product of the parasitaemia level of the sample and the number of views examined, it is possible to achieve high sensitivities by considering a large number of views. A similar argument can be used to estimate the sample positive predictive value.

This is a fundamental limitation of microscopy: the ability to detect malaria depends on the number of microscope fields examined [11]. The higher the parasitaemia of the infection, the greater the chance of detecting a parasite in a fixed number of slides.

A simulation can be performed to assess the ability of the system to diagnose full slides by assuming that there are 40 erythrocytes per slide and simulating the occurrence of infections in the sample at a low parasitaemia level (0.1%) and the reported sensitivity of the program (85%) using uniformly distributed random numbers. In a simulation where only 200 microscope fields are examined, 10 out of 11 infected cells would be detected by the algorithm.

The advantage of the automated technique is that coupled with an automated image acquisition system, it is possible to examine an entire slide, whereas technicians generally examine a slide for a set time or a fixed number of fields (e.g. 200) before declaring a sample negative. Furthermore, once the algorithm has identified infected erythrocytes, they are easily available for technicians to confirm the diagnosis (by viewing Fig. 11, for example) without having to search through fields of the slide themselves.

The results of the second classifier stage, which performs species differentiation per infected erythrocyte, are shown in Table 2. In this case, it is possible to determine the accuracy of the system at classifying complete samples, because the diagnosis of each sample, performed and validated by experts, is known.

Table 2 Results for species differentiation of infected erythrocytes

Species	Sensitivity (%)	PPV (%)
Plasmodium falciparum	57	81
P. vivax	64	54
P. ovale	85	56
P. malariae	29	28

The species for every infected erythrocyte is determined, and the sample species is the species of the highest number of parasites in the sample. Eleven of the fifteen samples tested were correctly diagnosed, giving an accuracy of 73% for the second classifier.

Although the sample accuracy is reasonable, clearly the success of the system at evaluating individual cells (and hence the ability to detect mixed infections) is problematic. The threshold selection technique currently used to find the parasites in the cells is primarily to blame. This technique does not identify enough of the entire parasite, but only fragments of it (see Fig. 9). Species classification is largely determined from the parasite shape, so a better thresholding technique that creates more accurate binary masks is required.

However, given that many methods of malaria detection currently in use, such as rapid antigen tests, have only limited success at differentiating between species, or are even specific only to certain species, the shortcomings of the algorithm at differentiating species does not invalidate the potential for this method as a diagnostic tool.

Another problem experienced relates to the detection of *P. malariae*. Unlike the other species, which are detected by the presence of trophozoites in the erythrocytes, *P. malariae* parasites form characteristic 'band forms'. Due to significant differences between these forms and the trophozoites of the other species, the features used are not well-suited to classifying them. A separate algorithm should be used to detect *P. malariae*, and the current system should be limited to the other three species.

5 Summary

An automated image processing and classification technique is able to detect erythrocytes infected with malaria parasites, and differentiate between the species of the infection. Images of thin blood smears are acquired using a CCD camera connected to a light microscope.

Following pre-processing, images are analysed to determine the size distribution of objects by granulom-

etry. Parameters generated in this way are processed using morphological techniques. This is combined with a novel local and global thresholding algorithm to segment possible infected erythrocytes from the image. The method of Di Ruberto et al. [3], modified to avoid the intermediate step involving the watershed transform to make it simpler and more efficient, separates clustered erythrocytes.

Features based on image characteristics such as colour, texture and geometry, as well as original features that mimic the qualities used by microscopists when diagnosing malaria, are generated from the erythrocytes which are candidates for infection. A tree classifier with two nodes using BFF neural networks determines whether or not a cell is infected, and if so, the species of the malaria.

Potential parasites are segmented from the image with a sensitivity of 92% and a PPV of 40%. The classifier is able to positively identify malaria parasites with a sensitivity of 85% and a PPV of 81%. A Bayesian argument that relates these results to examination of a full sample shows that a high sample sensitivity would be achieved by examining a large number of slides. The species was correctly determined for 11 out of 15 samples (a classification accuracy of 73%).

Problems arise due to inherent limitations of microscopy such as the degradation of slide quality with time, and the dependence on the number of microscopic fields of the sample. The thresholding used to identify parasite morphology is also a limiting factor, and *P. malariae* parasites should be diagnosed using a separate algorithm.

6 Conclusion

Compared to other diagnostic techniques, there are many advantages to using the proposed algorithm. It avoids the problems associated with rapid methods, such as being species-specific and having high per-test costs, while retaining many of the traditional advantages of microscopy, viz. species differentiation, determination of parasite density, explicit diagnosis and low per-test costs.

The biggest detraction of microscopy, namely its dependence on the skill, experience and motivation of a human technician, is removed. Used with an automated digital microscope, which would allow entire slides to be examined, it would allow the system to make diagnoses with a high degree of certainty. High capital costs and the necessity of electricity are some of the principle obstacles to using the system in the field.

There are, however, many possible applications where it would be suitable, such as performing an initial assessment for validation in a clinical laboratory. It would also constitute a diagnostic aid for the increasing number of cases of imported malaria in traditionally malaria-free areas, where practitioners lack experience of the disease.

Acknowledgements The financial assistance of the Department of Labour (DoL) towards this research is hereby acknowledged. Opinions expressed and conclusions arrived at are those of the author and are not necessarily to be attributed to the DoL.

References

1. Bloland PB (2001) Drug resistance in malaria, WHO/CDS/CSR/DRS/ 2001.4. World Health Organization, Switzerland
2. Bruce-Chwatt LJ (1985) Essential malariology, 2nd edn. William Heinemann Medical Books, London
3. Di Ruberto C, Dempster A, Khan S, Jarra B (2002) Analysis of infected blood cell images using morphological operators. Image Vis Comput 20(2):133–146
4. Foley DH (1972) Consideration of sample and feature size. IEEE Trans Inform Theory IT-18 618–626
5. Foster S, Phillips M (1998) Economics and its contribution to the fight against malaria. Ann Trop Med Parasitol 92:391–398
6. Hibbard LS, McCasland JS, Brunstrom JE, Pearlman AL (1996) Automated recognition and mapping of immmunolabelled neurons in the developing brain. J Microsc 183(8):241–256
7. Hughes GF (1968) On the mean accuracy of statistical pattern recognition. IEEE Trans Inform Theory IT-14, 55–63
8. Makler MT, Palmer CJ, Alger AL (1998) A review of practical techniques for the diagnosis of malaria. Ann Trop Med Parasitol 92(4):419–433
9. Mohana Rao KNR, Dempster AG (2001) Area-granulometry: an improved estimator of size distribution of image objects. Electron Lett 37(15):950–951
10. Mui JK, Fu K-S (1980) Automated classification of nucleated blood cells using a binary tree classifier. IEEE Trans Pattern Anal Machine Intell 2(5):429–443
11. Oaks SCJ, Mitchell VS, Pearson GW, Carpenter CCJ (eds) (1991) Malaria: obstacles and opportunities, National Academy Press, Washington, DC. A report of the committee for the study on malaria prevention and control: Status review and alternative strategies
12. Ohrt C, Purnomo, Sutamihardja MA, Tang D, Kain KC (2002) Impact of microscopy error on estimates of protective efficacy in malaria-prevention trials. J Infect Dis 186(4):540–546
13. Otsu N (1979) A threshold selection method from gray level histograms. IEEE Trans Syst Man Cybern SMC-9(1):62–66
14. Pammenter MD (1988) Techniques for the diagnosis of malaria. S Afr Med J 74(2):55–57
15. Soni PN, Gouws E (1996) Severe and complicated malaria in KwaZulu-Natal. S Afr Med J 86(6):653–656
16. Theodoridis S, Koutroumbas K (1999) Pattern Recognition. Academic, San Diego
17. Thiran J-P, Macq B (1996) Morphological feature extraction for the classification of digitalim ages of cancerous tissues. IEEE Trans Biomed Eng 43(10):1011–1019
18. World Health Organisation (1991) Basic malaria microscopy. World Health Organisation, Geneva