CERN - DATA HANDLING DIVISION
DD/73/3
E.S. Gelsema
B.W. Powell
P.L. Berrini*
January, 1973

PROCEDURES FOR THE AUTOMATION OF THE WHITE BLOOD CELL DIFFERENTIAL COUNT

Paper to be presented at the NTG/GI Meeting on "Cognitive Procedures and Systems", Hamburg, April 11th - 13th.

---

* of the Istituto di Biometria e Statistica Medica, Milano

DD-jh

PROCEDURES FOR THE AUTOMATION OF THE WHITE BLOOD CELL DIFFERENTIAL COUNT

E.S. Gelsema, B.W. Powell
CERN - Geneva

P.L. Berrini
Istituto di Biometria e
Statistica Medica - Milano

1. INTRODUCTION

The complete clinical examination of a blood sample comprises a number of tests of which one is the white blood cell differential count. This test consists of examining a suitably prepared blood smear under the microscope to establish the percentage of occurrence of the five normal types of white cell and to see whether any abnormal or immature cells are present. Typical values would be:

Neutrophils	(including band cells)	65%
Lymphocytes		27%
Monocytes		5%
Eosinophils		2%
Basophils		1%

A complete description of the various hematological tests has been given by Wintrobe (cf (11)).

When done by a trained technician, the average time required to find and classify the customary 100 cells is in the region of ten minutes. A large hospital, with some 2000 beds, would do in the region of 100.000 differential counts per year and thus need several technicians for this purpose alone. Given the repetitive nature of such work and the importance of recognising the occasional anomalies, maintaining a uniformly high standard of work is difficult and automation would be desirable if the required level of performance could be reached.

Although other approaches are being tried too by Kamentsky (cf (7)) and Saunders (cf(10)), publications by Bacus (cf(2)), Ingram (cf (5)), Prewitt (cf (9)) and Young (cf (11)) suggest that most of the work on automatic white blood cell recognition has been based on the "pattern recognition" approach, meaning: the analysis by digital computer of a digitised grey-scale image. This is not surprising in that many of the features thought to be used by the eye can in principle be extracted by the computer from such an image and on the other hand there are not many other properties of these cells which could be exploited for recognition purposes

and which would also reveal the anomalies which may occur.

Of the work published to date, only the groups of Bacus (cf (2)) and Ingram (cf (5)) have presented results based on as many as a thousand cells and although some manufacturers have announced their intention to market scanners suitable for routine use, it has still to be demonstrated that one can do as well as a competant technician. This is particularly true with regard to recognising immature cells.

The work described in the present paper has been done in collaboration with the group of Neurath at the New England Medical Center Hospitals (NEMCH) in Boston who are engaged in a project aimed at the development of an instrument suitable for routine use. From NEMCH we have received cell images digitised using their existing CRT scanner PIQUANT which has been described by Neurath (cf (8)). The results presented here have been obtained by analysing this data using the CDC 6600/6500 system at CERN.

2. DATA ACQUISITION

For normal use, a blood smear is prepared by placing a drop of blood on a glass microscope slide and then smearing it out into a thin layer by pressing another slide over it. It is then stained either by hand or with an automatic staining machine using Wright's stain or an equivalent. More uniform preparations can be obtained by using a "slide-spinner" to spread out the drop of blood but these are not in general use. Most of the cells present are red cells since they are about one thousand times more numerous than the white ones. The white cells seldom overlap one another but do sometimes appear in contact with one or more red cells. In a good preparation this is relatively rare.

Staining gives the white cell nucleus a bluish-purple appearance which varies somewhat with cell type. The outer part of the cell (or cytoplasm) stains differently and less heavily and may range from blue-grey through blue to pink in colour. This colouring serves to define the boundaries of both the cell and the nucleus as well as providing important information about the cell type (figures 1a and 1b). Further information about the cell can come from the presence of granules in the cytoplasm which stain differently and thus stand out clearly. Also in the nucleus, in addition to its variation in shape, some internal structure or texture is usually visible and this too contributes to the recognition process. Even from this very brief description it is easy to see that recognition by eye uses some combination of features such as cell area, nuclear area, nuclear shape, contrast between nucleus and cytoplasm, cytoplasm colour, nuclear colour, nuclear texture, presence and colour of granules, etc. In the present work one tries to evaluate similar features with the computer and

to classify a sample of cells using the values obtained.

The PIQUANT scanner measures a 35 mm photograph rather than operating directly through a microscope. Thus each cell is first photographed at high magnification using coloured filters. The resulting black-and-white negatives (one for each filter) are then scanned in turn using PIQUANT. In this way the optical density of the images is digitised into 64 levels over an area which on the slide corresponds to 96 x 64 microns using a nominal resolution of 0.1 microns. Though such a procedure would be unsuitable for routine use, for the purpose of acquiring data with which to develop the recognition programs it is entirely adequate and in fact produces higher quality data than one would have from a scanner optimised to this specific application.

The absorption spectra of the cells after staining do not show a great deal of structure. For the present data two coloured filters (Kodak 44 and 22) have been used corresponding to the regions 440 - 550 nm and 560 - 700 nm respectively. Other tests have indicated that the red cells can be more reliably rejected by using an additional blue filter for the region 400 - 470 nm (Kodak 47B) but it is not yet clear whether this additional complication is necessary when one uses well-prepared slides. Moreover the results of Gelsema (cf (4)) suggest that the additional filter does not contribute much to the colour measurement of the areas within the cell.

3. PREPROCESSING

The raw data for a cell consists of two sequences of grey-scale values (one for each filter image). The work described in the following sections is done on a (3 x 3) reduced image, i.e. only one point out of 9 in the original raster is used. The nominal resolution is therefore 0.3 μm and a typical cell image has 60x60 raster points. Fig.1c is a computer printout of a cell image, where the grey-scale information is represented by 16 different printer characters.

The purpose of preprocessing is to distinguish three areas of interest in such a picture, i.e. background, cytoplasm and nucleus. Also, any touching red cells should be removed from the white cell. They must not be included in the background either, because, as will be seen later, the background density is used as a reference for the colour determination.

In order to achieve this, a grey-scale histogram is constructed, in which the frequency of occurrence of all density values is plotted versus the density. Fig.1d shows such a histogram for the cell in fig.1c.

Ideally such a histogram has three peaks, roughly corresponding to the

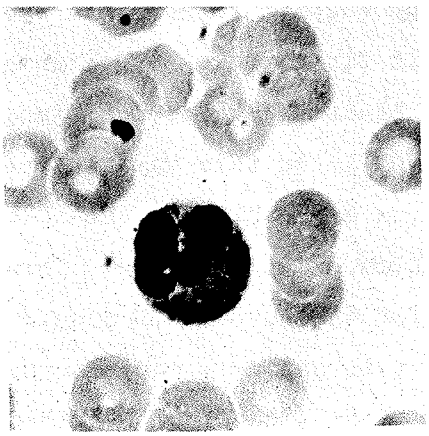

a

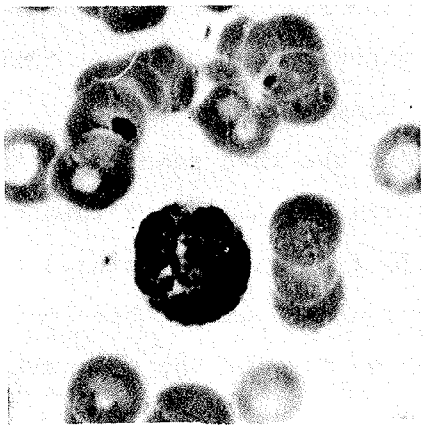

b

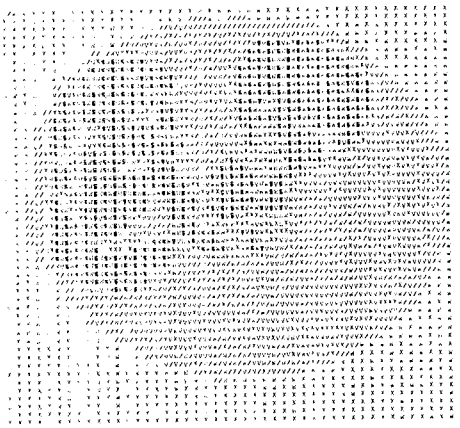

c

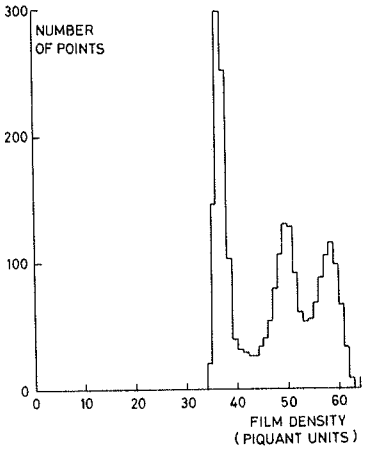

d

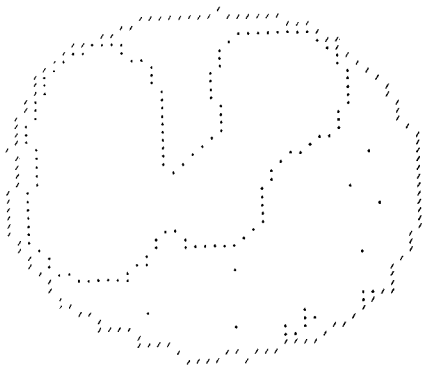

e

Figure 1

a and b: eosinophil as registered on black and white film through a red and a green filter, respectively. c: computer print-out of the red image of the same cell with 16 levels of grey. The horizontal dimension is slightly expanded with respect to the vertical one. d: histogram showing the number of points in the field of fig. 1c with a given density versus the density. From this histogram threshold values for the separation into background, cytoplasm and nucleus are obtained. e: boundaries of cytoplasm and nucleus as obtained in the preprocessing phase.

three areas mentioned above. The background density value d_{bg} is defined as the density corresponding to the maximum in the first peak. The threshold value for the cytoplasm is then defined as the integer value nearest to $d_{bg} + 3$. This has proved to be a simple and adequate estimate for a large number of cells.

The threshold value for the nucleus is more difficult to obtain, because to some extent the shape of the right hand side of the grey-scale histogram depends on the cell type. For a histogram as in fig.1d, where there is a clear minimum between the peaks corresponding to cytoplasm and nucleus a simple and stable estimate for the threshold level is the integer value nearest to the value for which the frequency is minimum. This is stable in that if an error of one unit in the threshold value occurs, it will cause little change to the boundaries computed. The situation is more difficult in cases where the nuclear peak is only a shoulder on the cytoplasm peak or when the cytoplasm peak degenerates to a shoulder on the nuclear peak. It was therefore decided to determine the nuclear threshold not from the grey-scale histogram but from its derivative. The threshold level is then set at the integer value nearest to the minimum in the derivative curve, (having first excluded that due to the maximum of the histogram itself). It was shown on a large sample of cells that this is an adequate estimate. In "shoulder cases" the determination is of course more critical, since a small error in the threshold value may have a large effect when applied to the image.

The procedure described above having been applied to both histograms of a cell, the cytoplasm threshold on the red image t_c is then taken as the final one and the nuclear threshold on whichever image gives the smallest nuclear area, t_n, is retained, this yielding the best approximation to the real morphological nucleus.

A process of contour following is then initiated on the red picture to locate the cell boundary. This is done on the red image because there the red cells are more transparent than on the green image and are therefore less likely to distort the boundary. The process consists of: starting at a given point with density $\geq t_c$ on the cell boundary, moving to a neighbouring point on the scan raster with:

i) density $\geq t_c$ and

ii) colour angle greater than a preset minimum value, keeping all neighbouring points with density $< t_c$ to the left (our units of colour measurement are defined in the next section).

This process is repeated until one returns to the original starting point. In the area enclosed by the cell contour a search for nuclear material (points with density $\geq t_n$ on the appropriate image) is then initiated and when one has found a nuclear point, a nuclear contour is traced in the same way as described above. One must then

make sure that all nuclear material is enclosed by the contour as some cells may have nuclei consisting of apparently distinct fragments. Once all nuclear fragments have been found and contoured in this way, the areas corresponding to cytoplasm and nucleus are easily obtained in the form of segment tables, i.e. tables giving the starting point and end point of the area on each scanline. In this form the areas of interest are in the most suitable form for the process of feature extraction. The cell and nuclear contours for the cell in fig.1c, as obtained in the way described above are given in fig. 1e.

4. FEATURE EXTRACTION

The properties or features that the eye uses in recognising the different cell types may be subdivided into three general categories:

i) geometry
ii) colour
iii) texture.

In this work similar features are extracted. The three categories are described below:

Geometry

Once one has the segment table representation of the image (section 3) parameters such as total cell area, total nuclear area, nuclear perimeter, etc. are readily obtained. Two other parameters, based on these and which are thought to contribute considerably to the separation between the various types are the ratio o. nuclear to cell area and the nuclear shape. The shape factor was defined as the ratio of the square of the circumference to the nuclear area.

It is already clear at this point that there will exist strong correlations between the parameters. They cannot therefore be expected to contribute equally to the separation. This problem will be dealt with in the next section.

Other geometrical features are total extinction, average transmission for both the cytoplasm and the nucleus, and contrast between the cytoplasm and the nucleus. Through the density vs. illumination curve of the film these parameters can be defined on the smear rather than on the film, assuming that one always works in the linear part of the curve and assuming $\gamma = 1$. Total extinction is defined as the area of a completely black disk of the same absorption. These parameters are classified in the category of geometry rather than of colour because they can be extracted from both colour images independently. This is in fact done because it is not known a priori which image will yield the more powerful descriptor. Here again, of course,

strong correlations between the parameters occur. In total 19 geometrical features are extracted.

Colour

The colour properties of a cell (for a given light source) are completely described by its absorption spectrum (fig. 2). Even the eye is unable to exploit all the information contained in such a curve. What the eye sees can be simulated by combining the information contained in three suitably colour-filtered images. In this way each possible colour may be represented as a point in a chromaticity diagram as described e.g. by Judd (cf (6)). Young (cf (11)) and Gelsema (cf (4)) have studied the use of three filters for this application. Using only two filters, it is not possible to represent colours in this way. By measuring the transmission through the object using each of the two filters in turn one can, however, characterise the colour by using the ratio of the two values. This is a degenerate form of colour representation with respect to the chromaticity coordinates but it is still powerful when the two filters are suitably chosen.

Transmission, being directly related to film density, a colour may be represented as a point in a diagram with the two axes corresponding to the two film densities. Shifting the origin in this diagram to the point corresponding to both background densities (the two-filter representation of white), the colour represented by a point in the diagram may then be measured as the angle between the radius vector to that point and one of the axes. The "red axis" was taken as the reference direction. Absorptivity is defined as the absolute value of the radius vector.

In the process of extraction of the colour parameters the colour and absorptivity of a sufficient number of uniformly distributed points in the cytoplasm and in the nucleus were determined. Points near the boundaries were not taken into account in order to minimize the propagation of errors from the preprocessing phase. Average values and standard deviations of the distributions for cytoplasm and nucleus are the final colour parameters. In order to have colour parameters that are completely free from errors in the nuclear boundary definition, averages of colour and absorptivity for the whole cell were also retained. Finally, colour contrast was defined as the vector in the colour diagram joining cytoplasm and nucleus of each cell. A total number of 12 colour parameters were extracted.

Texture

The approach to texture presented here should be regarded as a preliminary one, liable to indicate ways to possibly better solutions.

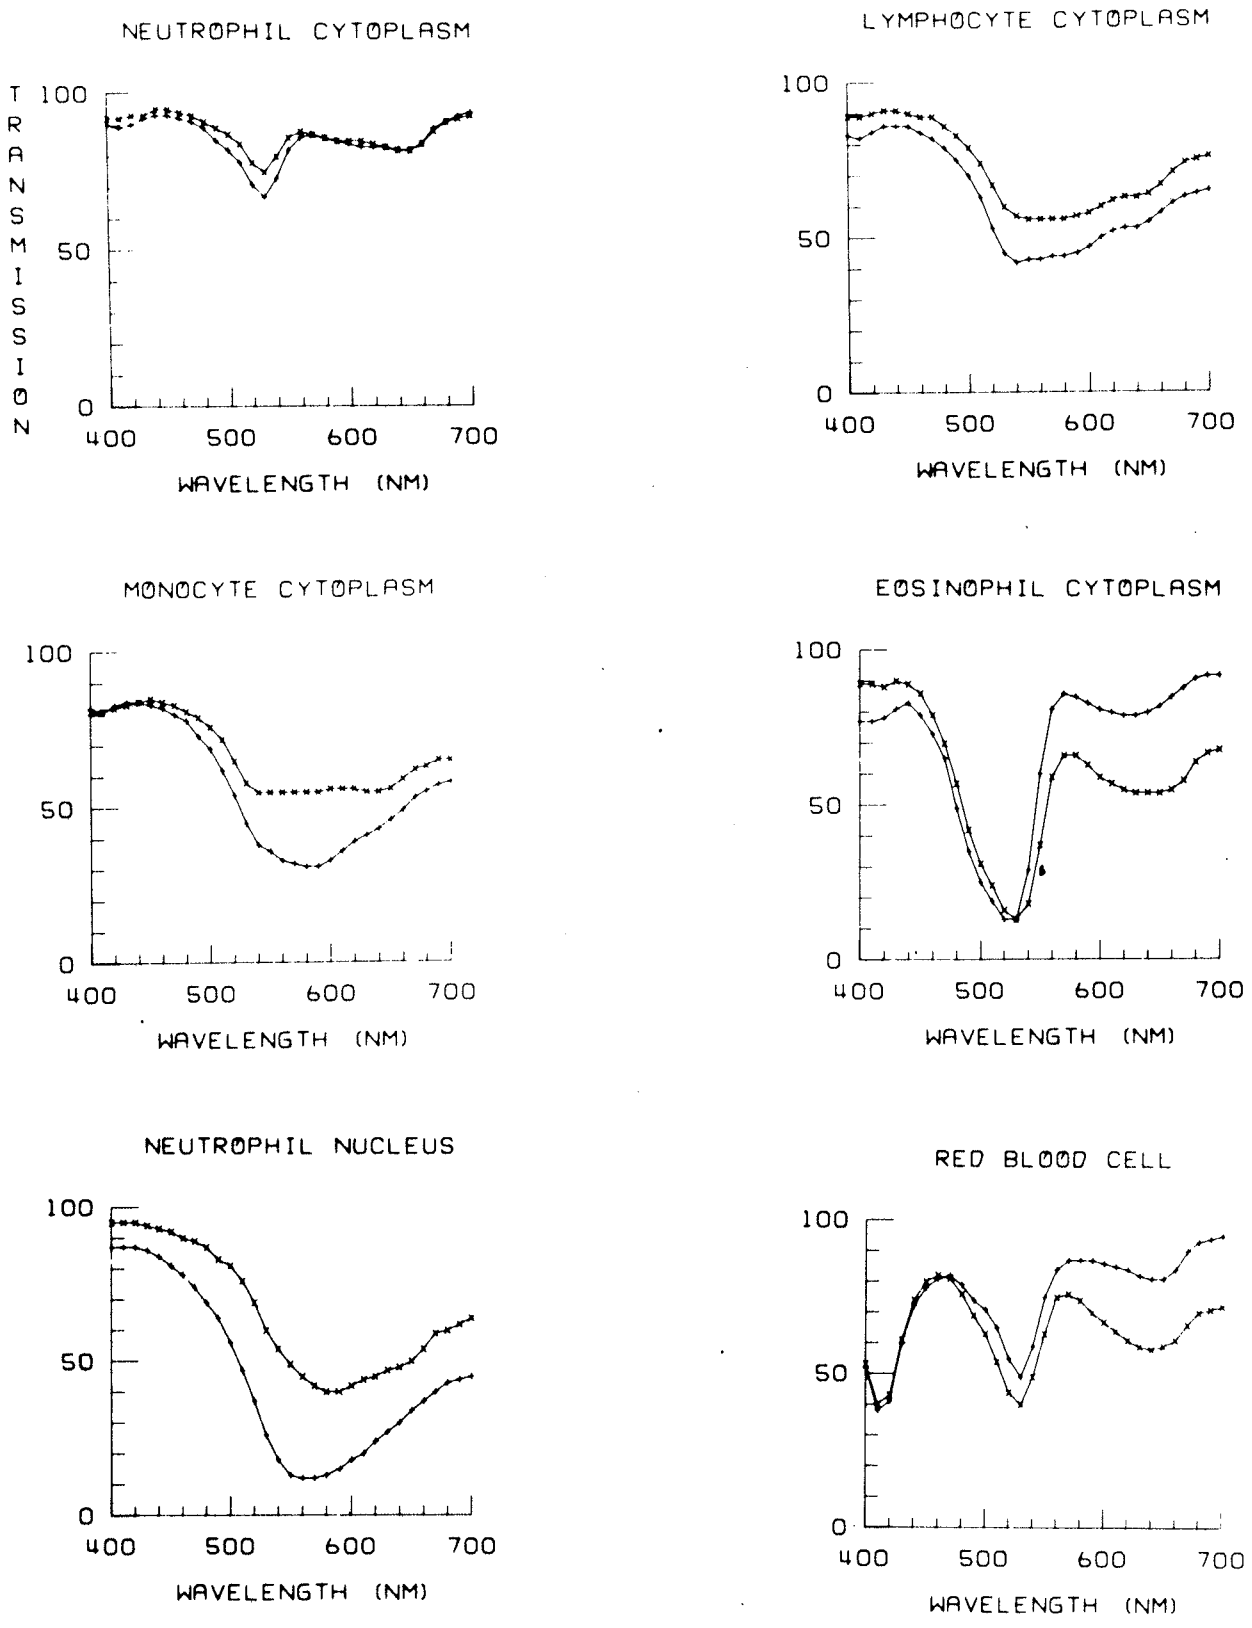

Figure 2

Absorption spectra as obtained using a microspectrophotometer. Each curve shows the transmission as a function of wavelength for a region of the cell approximately 1 μm in diameter.

In the cell three types of object were defined:

1) very high density objects ($d_{max}-4 \leq d \leq d_{max}$, where d_{max} is the highest non-empty bin of the frequency distribution)

2) low density objects (holes) in the cytoplasm ($d < d_c$, where d_c is the cytoplasm threshold level)

3) low density objects (holes) in the nucleus ($d < d_n$, where d_n is the nuclear threshold level).

For these three types of objects the following five parameters were obtained:

a) chance of hit (i.e. following all scan lines the number of transitions from non-object points to object points divided by the total number of non-object points)

b) total area (i.e. the total number of object points)

c) mean length (i.e. the average number of consecutive object points)

d) average colour

e) average absorptivity.

If a certain type of object was not found the last two parameters were set to the colour and absorptivity of the area in which it was defined (i.e the surrounding area being either cytoplasm for type 2 or nucleus for types 1 and 3). Thus a total of 15 texture parameters were transferred to the classification process.

5. CLASSIFICATION

For the classification of the white cells on the basis of the 46 features described in the previous section two methods have been used, one being in a way complementary to the other. Both methods consist of two phases. In the learning phase cells of known type are used to establish for each type mean values and standard deviations of all parameters as well as the correlations between them. In the classification phase these values are then used to assign unknown cells to one or other of the classes thus defined. Textbooks on multivariate analysis are e.g. Anderson (cf (1)) and Cooley (cf (3)).

Linear Discriminant Analysis

With this procedure one finds the linear combination of all parameters which best separates the different classes by optimizing the ratio of differences amongst the different classes to the within group differences. The coefficients for

the discriminant function represent the weights assigned to the different parameters and may be used to list them in order of decreasing separation power. In this way the problem of singular dispersion matrices in the case of completely dependant parameters is avoided. Also, if strong correlations between pairs of parameters exist, the better one will have the larger weight, the weight of the second one decreasing with increasing correlation coefficient. This ranking of parameters will in general depend on the classes being considered. The linear discriminant model used is based on two assumptions about the different populations:

i) all parameters are normally distributed and

ii) have the same dispersion matrix.

For the majority of parameters the first condition is probably fulfilled, the second, however, is not. For this reason only the learning part of the discriminant procedure was used. Classification was then achieved using a least chi-square method, using the best parameters as obtained from the discriminant model.

Least Chi-Square Analysis

In the least chi-square procedure the learning phase is repeated in order to calculate a dispersion matrix for each class. By taking the most powerful parameters from the preceding analysis one is relatively sure that problems with singular matrices will not occur.

In the classification phase the following quadratic form is evaluated for each class j:

$$\chi_j^2 = (\vec{x}-\vec{\bar{x}}_j)^T D_j^{-1} (\vec{x}-\vec{\bar{x}}_j)$$

where $\vec{x}$ is the vector of parameters, $\vec{\bar{x}}_j$ is the vector of group means for class j and D_j is the dispersion matrix for class j (the superscript T indicating transposition of the column vector). This function has a χ^2 distribution with N degrees of freedom, where N is the dimension of $\vec{x}$ (= the number of parameters used).

For each unknown cell the expression:

$$\chi_j^2 + \log_e |D_j|$$

is then evaluated, where $|D_j|$ is the determinant of the dispersion matrix. The cell is assigned to the class for which this expression is minimum. When the dispersion matrices are different for the various classes, the second term serves to minimize the number of misclassifications by drawing the boundaries between classes through points of equal population densities rather than through points of equal χ^2. The only assumption underlying this classification procedure is the assumption of multivariate normal distributions.

A priori probabilities are sometimes included in the expression to be minimized. This is not done here as the objective of the present work is to demonstrate the classification power of a set of parameters, independently of relative sample sizes.

6. RESULTS

When comparing a computed classification with that of a hematologist, disagreements may arise from both machine misclassification and human classification errors. Bacus (cf (2)) has shown that for the five normal cell types a human error rate as low as 1.5% can be expected.

However, when a subdivision of these classes is attempted, the number of disagreements between humans rises rapidly. This is not surprising since different hematologists use slightly different criteria to decide, for example, when a band cell has developed into a neutrophil. A given hematologist may, however, be quite self-consistant in applying his own criteria. For this particular subdivision Bacus finds that for a group of nine people, individuals disagreed with the consensus in at least 15% of the cases. In the comparison of the machine classification against that of a hematologist these systematic differences may not be so important. If there is no clear transition between two successive states, the machine will merely show the same biases as the hematologist who provided the initial classification.

For the present work a total of 1146 cells were available in digitised form. These consisted of 483 normal cells (including band cells) and 663 immature cells. This very large proportion of immature cells was chosen deliberately so that one could study the probability of immature cells being misclassified as normal ones and also attempt the classification of the various immature forms. In table 1 the sample sizes in the training set and in the test set are given according to the hematologist's classification. It should be noted that the total sample consists of white blood cells and other nucleated cells that may occur in a smear. In the present work the cell types indicated with an (N) are considered as normal cells, all other types being referred to as immature cells.

With a training set of 243 cells and a testing set of 240 cells, 15 parameters have been used to classify normal cells into their six types. The confusion matrix for the testing set is shown in table 2. The percentage of correct classifications is 87%. Apparently better results than these can be obtained with more parameters and by using the same cells for both the training and the testing set. Indeed, with 30 parameters and using the whole set, a 99% correct confusion matrix was obtained. This results from using too many parameters for the statistics avail-

TABLE 1

Class Name	Sample size in training set	Sample size in testing set
Myeloblast	33	32
Promyelocyte	31	30
Myelocyte	34	34
Metamyelocyte	34	33
Band cell (N)	44	44
Neutrophil (N)	38	37
Monocyte (N)	36	35
Lymphoblast	34	34
Atypical lymphocyte	41	41
Lymphocyte (N)	37	37
Eosinophil (N)	38	38
Basophil (N)	50	49
Plasma cell	35	35
Nucleated Red cell A	25	24
Nucleated Red cell B	40	39
Nucleated Red cell C	27	27
Total	577	569

Composition of the total sample according to the hematologist's classification.

TABLE 2

		Computer Classification						
Cell Type	Total No.	Bands	Neutrophils	Lymphocytes	Monocytes	Eosinophils	Basophils	% correct
Bands	44	36	7	0	1	0	0	82
Neutrophils	37	5	29	0	1	1	1	78
Lymphocytes	37	0	0	30	2	0	5	81
Monocytes	35	0	1	0	31	0	3	89
Eosinophils	38	0	1	0	0	35	2	92
Basophils	49	0	0	0	1	0	48	98

15 parameters

Percentage correct 87%

Training set: 243 cells

Testing set: 240 cells

Confusion matrix for six class machine classification of the cells of normal type.

TABLE 5

		Computer Classification															
Cell type	No. of cells	Myeloblast	Promyelocyte	Myelocyte	Metamyelocyte	Band cell	Neutrophil	Monocyte	Lymphoblast	Atypical Lympho	Lymphocyte	Eosinophil	Basophil	Plasma cell	Nuc. Red Cell A	Nuc. Red Cell B	Nuc. Red Cell C
Myeloblast	32	13	8	1	2	0	0	1	1	2	0	2	0	2	0	0	0
Promyelocyte	30	2	14	4	2	0	0	0	3	1	0	0	1	3	0	0	0
Myelocyte	34	0	2	14	10	0	1	1	3	0	0	1	2	0	0	0	0
Metamyelocyte	33	1	1	7	15	0	1	5	0	0	0	0	3	0	0	0	0
Band cell	44	0	0	0	0	32	10	2	0	0	0	0	0	0	0	0	0
Neutrophil	37	0	0	0	0	4	29	1	1	0	0	2	0	0	0	0	0
Monocyte	35	0	0	1	6	1	0	17	0	7	0	0	2	1	0	0	0
Lymphoblast	34	3	1	5	0	0	0	0	18	0	0	0	3	4	0	0	0
Atypical Lymph.	41	2	1	0	1	0	0	0	1	26	3	0	1	6	0	0	0
Lymphocyte	37	2	0	0	0	0	0	1	2	4	26	0	0	1	1	0	0
Eosinophil	38	2	0	0	0	0	0	0	0	0	0	36	0	0	0	0	0
Basophil	49	0	1	0	2	0	0	1	1	1	1	2	40	0	0	0	0
Plasma cell	35	1	0	0	0	0	0	1	2	1	0	0	1	28	1	0	0
Nuc. Red Cell A	24	0	0	0	0	0	1	0	0	4	0	0	0	0	13	6	0
Nuc. Red Cell B	39	0	0	0	0	0	1	0	0	0	0	0	0	0	4	32	2
Nuc. Red Cell C	27	0	0	0	0	0	2	0	0	0	0	0	0	0	2	12	11

15 parameters

Percentage correct for 16 classes: 64%

Percentage correct for 10 classes: 77%

Training set: 577 cells

Testing set: 569 cells

Confusion matrix for machine classification of normal and abnormal cells.

able. With a separate testing set misleading results such as these are unlikely.

In order to make a comparison with the results of Bacus and with an early result obtained by the NEMCH group who used a subset of the present data, classification on the basis of 8 parameters was also performed. With 15 as well as with 8 parameters a testing set different from the training set was used. The results, given in table 3A, are for both five and six types - the five being obtained by adding the band and neutrophil types together. From these results, in view of the sample sizes, it is concluded that comparable performance has been obtained by the three groups. Even when the band cells and neutrophils are combined the performance is not as good as that of a human though with the statistical errors normally present in a differential count the additional errors due to the machine would be small. The effect of using 8 or 15 parameters has been considered (see also table 3B where a greater variety of cell types has been used) and as can be seen the benefit is rather small. The reasons for misclassification have still to be studied in detail but a first look suggests that when errors occur in the location of the cell and nuclear boundaries many of the parameters are affected and the classification will often depend on the particular choice of parameters.

When all cell types are being considered the testing set consists of 569 cells. The confusion matrices shown in tables 4 and 5 show the result of trying to classify these cells using the best 15 parameters chosen by the discriminant analysis program. From table 4 one sees that a total of 9.4% of the immature cells were confused with normal cells. Similarly 13.8% of the normal cells were classified as immature. Since the former figure corresponds to the proportion of "false negatives" it may seem the more important of the two. Though we do not have figures on this, hematologists expect to miss only a few percent of immature cells. However, this is assuming that there will be more than one immature cell present in the sample. With a 90% probability of recognition for each immature cell, the probability of recognising that some are present will normally be high and a greater problem is that of getting rid of the false positives which would be numerous and which could only be checked by visual inspection. These numbers are therefore marginal because they imply a machine which requires too much help rather than because it cannot find the immature cells with sufficient reliability.

In table 5 the complete classification is shown. The order in which the cells are listed has been chosen so that cells of a given type in different stages of evolution occur in the correct sequence. The boxes indicate categories which are not very well separated in nature and which one could reasonably merge together. With the full sixteen types 64% are correctly classified and this rises to 77% when only ten categories are used. Table 3 shows the numbers obtained when using only 8

parameters.

As above, tests using 30 parameters and the whole set of cells have been made giving 91% correct classification for the sixteen cell types. We do not, however, believe that a separate test sample would confirm this result.

At this stage the estimation of performance seems of greater importance than the prediction of the cost and speed of a device to do routine differential counts. However, it is estimated that the extraction of 15 parameters and the subsequent classification of a cell take respectively in the region of 0.9 and 0.04 CP seconds on the CDC 6600. This figure can certainly be improved upon by optimizing those parts of the procedure that are most time consuming. This has not been studied so far. On the other hand it is not clear how much additional computation would be needed to get a significant improvement in performance and this can only be found out by further study.

7. CONCLUSION

The results presented in the previous section confirm that the five commonly occurring cell types can be recognised with a reliability slightly better than 90%. They also show that immature cells and other similar cells which are sometimes present in a sample can be recognised as such with 90% probability, while some 14% of normal cells are incorrectly classed as immature. The classification of the immature cells has also been tried with an overall success of between 64% and 77% depending upon the degree of subdivision attempted.

For the classification of white cells into the five normal types, the present results, together with those of Bacus and the unpublished results of the NEMCH work show that one is approaching the recognition efficiency of the trained technician. The present results also show with good statistics that on the crucial point of recognising whether immature cells are present, the performance is adequate but the rate of false positives, being 14% of the normal cells present, is too high for an automatic device.

Since the normal differential count is based on only 100 cells the statistical accuracy would only be slightly modified by the additional errors of the automatic system (even if the standard sample became 200 cells this would still be true). It is therefore mainly in the reduction of the probability of getting false positives without loss of efficiency in recognising genuine immature cells that there is the most need for improvement.

Further study of the present data is likely to show that some improvement is possible. This study needs to cover the whole sequence starting with the hematologist's classification and then looking for reasons for misclassification by the computer, for systematic effects in the data, for incorrect extraction of features etc. Finally a reconsideration of the statistical methods of classification would be desirable. The choice of parameters also deserves more thought. In particular, it has emerged very clearly that increasing the number of parameters does not necessarily improve performance. Better parameter evaluation may be more important and since colour seems to be one of the most significant areas, detailed study of how to evaluate it more effectively would be desirable. There is also reason to think that a more careful procedure for defining the boundaries would help significantly in the parameter extraction.

The white blood cell differential count may prove to be one of the first successful medical applications of image processing techniques. The results now available show that on statistically significant samples one is coming close to the performance required for a working device. There is a need to improve the recognition performance still further and the studies should be extended to even larger samples so that as far as possible one meets with the full range of biological variability.

Acknowledgements

We would like to thank Dr. G.R. Macleod, head of the Data Handling Division at CERN, for encouraging us to persue this work. We are also very much indebted to Dr. P.W. Neurath and his group at NEMCH in Boston for their help and for providing us with the data on which this work is based.

We would also like to thank Dr. R. Anner of the Hopital cantonal in Geneva, Dr. G. Gallus of the Istituto di Biometria e Statistica Medica in Milan, Dr. J.M. Howie of CERN, Mr. W. Selles of NEMCH and Dr. H.P. Wagner of the Tiefenauspital in Bern for their contributions to this work.

References

1. Anderson, T.W., An Introduction to Multivariate Statistical Analysis (John Wiley, 1958).

2. Bacus, J.W. and Gose, E.E., IEEE Transactions on Systems, Man and Cybernetics, SMC-2, 513, (1972).

3. Cooley, W.W. and Lohnes, P.R., Multivariate Data Analysis, (John Wiley, 1971).

4. Gelsema, E.S. and Powell, B.W., Colour Measurement and White Blood Cell Recognition, CERN Data Handling Division Report DD/72/24 (1972).

5. Ingram, M. and Preston Jr, K., Scientific American, 223, 5, p. 72 (November 1970).

6. Judd, D.B., Colour in Business, Science and Industry (John Wiley, 1963).

7. Kamentsky, L.A. and Melamed, M.R., Proceedings IEEE, 57, 2007 (1969).

8. Neurath, P.W., Brand, D.H. and Schreiner, E.D., Annals of New York Academy of Sciences, 157, 324, (1969).

9. Prewitt, J.M.S. and Mendelsohn, M.L., Annals of the New York Academy of Sciences, 128, 1035, (1966).

10. Saunders, A.M., Groner, W. and Kusnetz, J., A Rapid Automatic System for Differentiating and Counting White Blood Cells. Paper presented at the Technicon International Congress in New York, (November 1970).

11. Wintrobe, M.M., Clinical Hematology (Lea and Febiger, 1967).