Paper presented at the Third Irish Neural Networks Conference,
The Queen's University of Belfast, Northern Ireland, 1-2 September 1993

Cell searching with a neural net

Stephen J. McKenna, Ian W. Ricketts, Alistair Y. Cairns and Kudair A. Hussein*

Microcentre, Department of Mathematics and Computer Science, The University, Dundee, DD1 4HN, Scotland. *Department of Pathology, Ninewells Hospital, Dundee, Scotland.

Abstract. Cervical cancer screening programmes are currently based upon the microscopic examination of Papanicolaou smears. These contain a wide variety of cell types and non-cell artefacts. Extensive research effort has been devoted to developing automated cytology systems to analyse samples such as these. A useful step in any image analysis system for automated cytology is to detect interesting objects for detailed analysis as most of the area covered by a slide does not contain objects of diagnostic significance. Previous approaches to this specimen enrichment task have included methods based on thresholding, mathematical morphology and linear combinations of handcrafted image filters. An lternative approach is proposed which employs artificial neural networks applied directly to low-resolution images of conventionally prepared Papanicolaou smears. Training sets of labelled image pixels were formed using a previously developed high-resolution segmentation algorithm. Feed-forward networks with localised connectivity and hard weight-sharing were used to ensure that the features extracted were invariant under both translation and rotation, and to limit the number of free parameters in the network, thus promoting generalisation. In preliminary experiments, feed-forward neural networks were able to learn to locate abnormal cell nuclei and to enrich the specimen by rejecting uninteresting cells and background clutter.

1. Cervical smear inspection

The microscopic examination of a cytological sample, such as a cervical smear, often involves assessing in excess of 50,000 cells. This is a tedious and labour-intensive undertaking and one not without error. It is for these reasons that extensive research effort has been devoted to the development of partially or fully automated cytology systems. Previous research in this area has been reviewed elsewhere (Banda-Gamboa et al. 1992).

Papanicolaou smears typically contain between 10,000 and 00,000 cells (Tucker 1976). These cells comprise intermediate and superficial cells from the squamous epithelium of the cervix as well as leukocytes (white blood cells), other cells such as parabasal and metaplastic cells and various other artefacts. All the cellular material on a slide must be inspected under a microscope in the search for any abnormalities. Approximately 5% of smears contain some visually abnormal cells which might indicate a pre-cancerous or cancerous condition.

Figure 1 shows images of abnormal and normal cells taken from several different cervical smears. The abnormal cells have enlarged nuclei and relatively small cytoplasms. In contrast, normal superficial cells have small, dark nuclei and relatively large, polygonal cytoplasms. Normal intermediate cells have larger nuclei than superficial cells. Leukocytes are small highly contrasted objects.

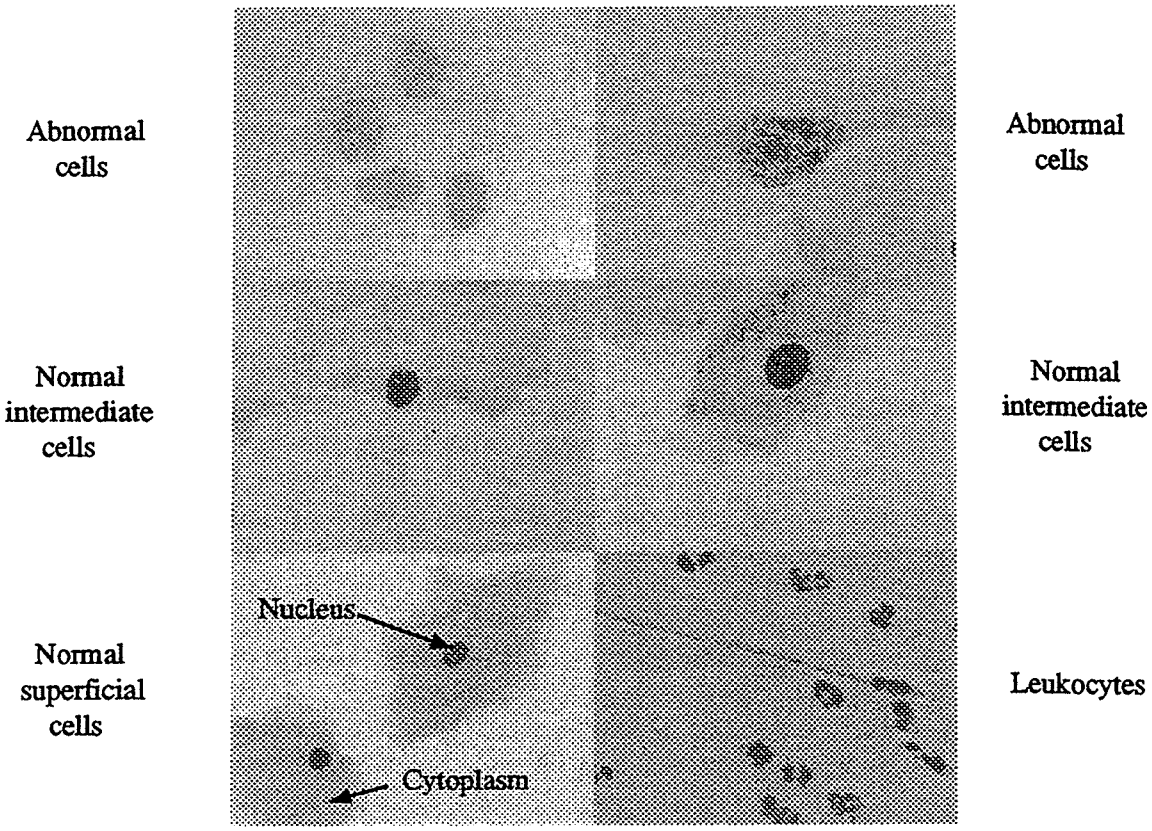

Figure 1. Intensity images of cellular material from conventional cervical smears.

The conventional Papanicolaou slide preparation used for manual screening is not well suited to automated analysis. Staining solutions and procedures vary widely. In addition, conventional smears are multi-layered with cells clustered together and overlapping one another. Alternative monolayer (or `thin') preparations have been developed in which cells appear isolated and within the same optical plane. Alternative stains, such as Acriflavine-Feulgen which renders cytoplasmic material invisible and nuclei highly contrasted, have also been used (e.g. Meyer 1979). Such preparations simplify image analysis but their introduction into a well established screening programme would be costly and disruptive. A system capable of inspecting conventional Papanicolaou smears would be more easily integrated into the programmes. The experiments described in this paper have been performed using routinely prepared smears.

Many automated systems adopt a dual-resolution approach to the inspection of a slide. An initial low-resolution scan is used to locate cellular objects and to discard those cells which are clearly of no diagnostic importance. This results in a much smaller, enriched cell specimen from which a reliable diagnosis can be made. The enriched specimen is subjected to more detailed analysis at increased resolution (see e.g. McKenna et al . 1993). This involves either searching for individual abnormal cells (`rare event detection') or analysing the enriched cell population as a whole in order to detect shifts in the cell population (`malignancy associated changes').

This paper is concerned with the low-resolution scan for locating cells and performing specimen enrichment. It is organised as follows. The specimen enrichment problem and previously published methods are described in more detail in the next section. In section 3,

the proposed method is introduced, suitable weight-sharing networks are discussed and a procedure for the collection of labelled training data is described. Section 4 presents results from preliminary experiments and section 5 contains concluding remarks and suggestions for further work.

2. Previous approaches to the specimen enrichment problem

The specimen enrichment process aims to discard material not useful for making a diagnosis such as background, normal superficial cells and leukocytes. Enough suspicious material should be retained to allow abnormalities to be detected. Suspicious material does not necessarily mean only visually abnormal cells: subvisual markers of abnormality have also been identified in intermediate cells (Zahniser et al. 1991).

Previous approaches to object location and specimen enrichment have been based upon thresholding techniques, iterative image transformations from mathematical morphology and combinations of handcrafted image filters. This section briefly reviews these methods.

2.1 Thresholding techniques

Early attempts at cell location and specimen enrichment employed thresholding techniques. These were typically applied to images with a pixel resolution of 4 microns. Cell nuclei have been detected using thresholds determined from peaks in the grey-level frequency-of-occurrence and transition histograms (Taylor et al. 1975). An adaptive thresholding method has been used to detect interesting objects while compensating for variations in intensity due to shading and staining. In this method, the threshold value was set equal to the mean intensity of neighbouring background pixels plus some user-defined constant (Tucker 1976). The CYBEST systems employed dynamic thresholding using differential histograms. Errors were due to irregular staining, clumped cells, leukocytes and a few missed abnormal cells (Watanabe et al. 1976, Tanaka et al. 1982, 1987).

The use of an Acriflavine-Feulgen stain results in highly contrasted images and allows almost all nuclei to be detected using a fixed intensity threshold. An erosion function can then be used to reject nuclei too small to be suspicious (Al and Ploem 1979). Pycock and Taylor performed cell searching on specially prepared monodisperse hematoxylin stained samples (Pycock and Taylor 1980). Thresholds were selected using syntactic histogram analysis and were related to the dominant peak and to an estimate of mean cytoplasmic intensity in order to accommodate variations in illumination and cytoplasmic staining respectively. A constant width of grey-levels around the peak was classified as background. Darker pixels to a fixed fraction of the mean background grey-level were classified as cytoplasm. A nuclei detection threshold was then set to a fixed fraction of the mean cytoplasm grey-level. Finally, contours were traced around the detected nuclear regions.

2.2 Mathematical morphology

Iterative image transformations based on mathematical morphology (Serra 1989) have also been used to detect suspicious cells but at a greater resolution than other methods. Meyer used samples stained using Acriflavine-Feulgen and analysed them at a resolution of 1 micron

(Meyer 1979). The method relied upon `top-hat' functions for detecting contrasted objects of various sizes. One such function detected suspicious cells along with granulocytes and clumped leukocytes, a second function discarded highly contrasted, isolated leukocytes and a third discarded clumped leukocytes. Iterative image transformations for the segmentation and separation of overlapping nuclei have also been developed (Meyer 1986). The use of mathematical morphology with conventional Papanicolaou stained smears has been reported recently (Lee et al. 1992). The images used contained 8-bit pixels at an increased resolution of 0.55 microns. Transformations for small object recognition and hole filling were described. They were implemented using custom-built gate array processors.

A two-stage erosion process has been used to eliminate small objects and objects which have a lower grey-level than their background (Erhardt et al. 1980). This suppressed dark structures with a width of less than 6 microns in any direction. Nearly all leukocytes were eliminated. The images used had 0.5 micron, 8-bit pixels.

2.3 Handcrafted image filters

Poulsen developed a parallel algorithm for locating free-lying suspicious cells in conventionally prepared Papanicolaou smears at a resolution of 4 microns (Poulsen 1973). The algorithm was parallel in the sense that each pixel was analysed in the same way and without influence from the result of processing adjacent pixels. A set of filters was applied to each pixel. Each of these filters was an annular ring, either filled (for locating cell nuclei) or unfilled (for locating cytoplasm). Linear combinations of the filter results were then computed by taking weighted sums and applying thresholds. The resulting features were then combined using logical operators. The algorithm's parameters were selected interactively by the user through a series of trial and error experiments. Those objects detected consisted of a dark nucleus surrounded by a medium dense area (cytoplasm) on a clean background. The algorithm was further developed by a group at Uppsala University in Sweden (Bengtsson 1987, Nordin 1989). A working prototype which kept up with "continuous scan rates of several megahertz" was reported.

A similar approach to that of Poulsen used templates consisting of horizontal, vertical and diagonal spokes radiating from a central pixel (Read et al. 1979). A template specification defined whether or not a pixel at a set distance from the central pixel was required to contrast with the central pixel. An additional parameter controlled the number of template spokes which had to match in order for the template to match. This parameter allowed templates to be developed for detecting highly irregular objects. A logical combination of as many as 15 templates determined whether or not the central pixel was of interest. The method was applied to images at a TMsolution of 3 microns. The PAPNET system (Neuromedical Systems, Inc. 1991) used morphologic techniques to locate potentially suspicious objects in routinely prepared Papanicolaou smears at a resolution of 1 micron. Between 1,000 and 10,000 such objects were detected. These operations were implemented on a parallel pipelined computer.

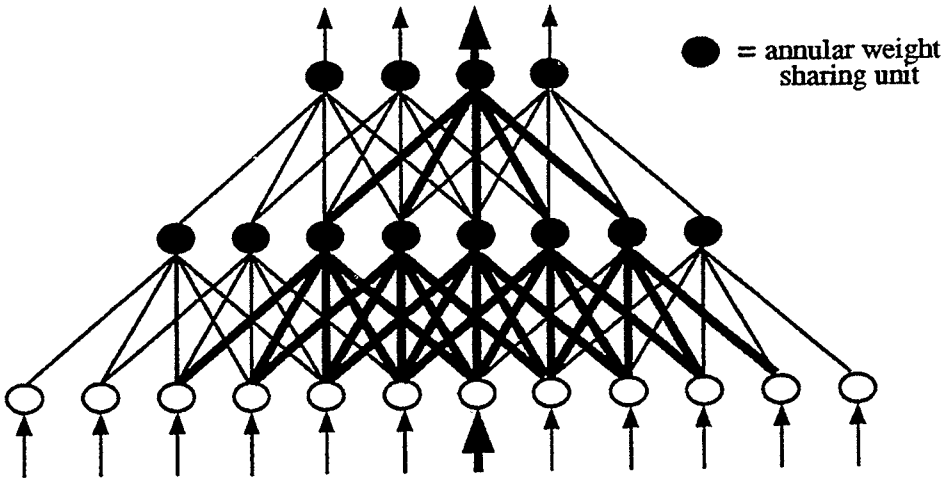

Figure 2. A slice through a network. Dark lines indicate those parts of the network involved in processing a particular pixel.

3. Artificial neural networks for specimen enrichment

The aim of this research is to develop a system for specimen enrichment using artificial neural networks. The learning capabilities of neural networks should enable them to compute appropriate parameters for modelling the appearance of cells. Their ability to generalise makes them good candidates for analysing noisy images. Once trained they are capable of delivering fast results (especially if their parallel structure is exploited in implementation) and when dealing with slides containing as many as 200,000 cells speed is important. The proposed method has more in common with Poulsen's method (see section 2.3) than with the morphological techniques published by Meyer (see section 2.2) and should not require images of the relatively high resolution used by the latter.

3.1 Neural networks with annular weight-sharing

The proposed specimen enrichment method employs feed-forward artificial neural networks trained using error-backpropagation (Rumelhart *et al.* 1986). Each network receives as its input an intensity image of a scene from a cervical smear and outputs a binary image indicating where any areas of interest lie within the input image. The networks were tailored to the application using hard weight-sharing and localised connectivity. All hidden layer and output layer units had circular receptive fields and all connections at an equal radius from a receptive field's centre were shared. This ensured that the features extracted by a network possessed rotational invariance and it forced the units to learn weight masks consisting of annular rings. Translation invariance was obtained by forcing all units in a given layer to share their weights. Figure 2 depicts a slice through such a network. Darker connections indicate which units exert influence upon the output for a particular input pixel. Figure 3 shows a similar network in three dimensions. This diagram illustrates how to extend the network of figure 2 by adding extra slabs of hidden units. All units within a slab share their weights but units in different slabs learn different weights.

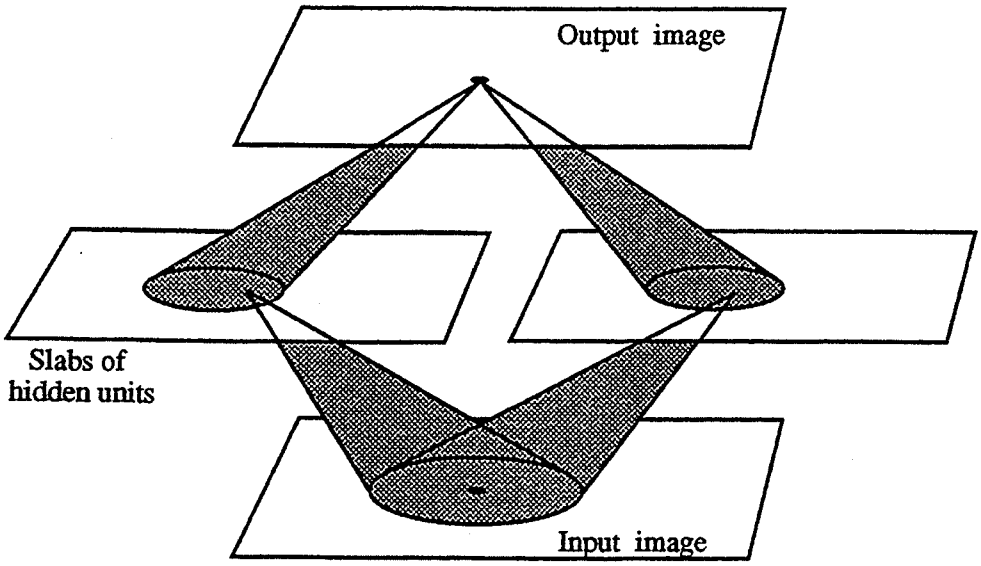

Figure 3. A weight-sharing network with two slabs of hidden units.

3.2 Data sets

The images used were obtained using a CCD/RGB camera mounted on a microscope fitted with a x100 oil immersion objective. Although colour digitised images were obtained, only the intensity frame was used in these experiments. Images originally consisted of 512x512 8-bit pixels at a resolution of 0.15μm. All images were resampled to a resolution of 2.4μm using pixel averaging. In order for supervised learning to be possible labelled data was required to form a training set. Labelled images were obtained using a previously developed segmentation technique for high-resolution images of isolated cervical cells (Parianos 1991). This technique was based upon the converging squares algorithm which searches for points of maximum value within regions of maximum intensity in an image (O'Gorman and Sanderson 1983). Binary segmentation masks were thus computed for images of single, isolated abnormal and normal cells. Both the original images and their corresponding segmentation masks were then resampled at low resolution and pixels were selected to form the training pairs.

There were many more uninteresting pixels of background and cytoplasm in these images than there were of cell nuclei. All classes of pixel, however, should have equal representation in the training set. Initially, an appropriately sized sample of the over-represented class was selected at random for inclusion in the training set. This approach was sufficient for the initial experiments on abnormal nuclei detection which will be described in section 4.1. A more sophisticated approach was needed, however, for the specimen enrichment experiments described in section 4.2. This involved training a network on a randomly selected training set as before. After training, the network was tested on the remaining data from the over-represented class. A proportion of the misclassified patterns were then brought into the training set to replace some of the previously used patterns and training was continued. This technique resulted in improved performance by moving hard-to-classify background and cytoplasm pixels lying near the true decision boundaries into the training set. The effect was to minimise the number of erroneously detected objects.

Table 1. Results for nuclei detection networks

Network	Weights (excl. biases)	Connections per pixel	Nuclei detected	False detections
69 - 1	5	69	108 (98%)	2 (4 pixels)
69 - 9x1 - 1	6	45	108 (98%)	2 (3 pixels)
Ensemble	11	114	108 (98%)	1 (1 pixel)
Ensemble + post-process	11	114	108 (98%)	0

4. Preliminary experiments

4.1 Abnormal nuclei detection

The first task given to the neural network was to scan images containing only abnormmal cells and to detect the cells' nuclei. The training set was compiled using 52 images of isolated abnormal cells taken from 6 smears. It contained 1115 nuclear pattems and 1115 non-nuclear pattems. The test set consisted of 50 images taken from 5 smears not used in the training set and contained 110 abnormmal cells. Test set results are given in table 1. The notation used shows the number of units in each layer of a network involved in classifying an image pixel. Thus a network arrangement of '69-9x2-1' describes a network with 69 inputs, 1 output and 2 hidden layer slabs of 9 units each.

The first network trained had no hidden layer. Its units had receptive fields with diameters of 9 pixels corresponding to 69 inputs and 1 output per image pixel. A second network with a hidden layer slab of 9 units was also trained. Both these networks detected 98% of the abnormal nuclei as well as 2 other erroneous objects corresponding to small areas of dark cytoplasm. Fortunately, the networks' errors did not always coincide so that by combining their outputs only a single pixel object was wrongly detected. A simple post-processing step which involved discarding single pixel objects resulted in no erroneous objects being detected. The 2 abnormal nuclei missed by the networks were poorly stained.

4.2 Specimen enrichment

The second task was to scan images containing both normal and abnormal cells and to perform a specimen enrichment. This involved detecting abnormal nuclei whilst discarding as many normal cells and other objects as possible.

Consider training a single network on the two-class problem to differentiate abnormal nuclei from all uninteresting material. A problem with this is deciding in what proportions to represent the various objects (e.g. superficial nuclei, cytoplasm, background) in the training set. The task could be recast as a multi-class problem with, for example, an abnormal nuclei class, a superficial nuclei class and a background class. This, however, would unnecessarily require the network to differentiate between superficial nuclei and background. Rather than train a single network to perform specimen enrichment, the problem was divided into two sub-problems performed by two different networks.

Table 2. Results for the specimen enrichment networks

Networks	Weights (excl. biases)	Connect/ pixel	Abnormal nuclei	Superficial nuclei	Cytoplasm
A	12	92	66 (60%)	35 (36%)	8
A & B	17	161	44 (40%)	3 (3%)	2
A & B + post-processing	17	161	31 (28%)	0 (0%)	0

(Table 2 gives the number and type of objects detected by the specimen enrichment networks. Network A (69-9x2-1) was trained on abnormal nuclei vs. background/cytoplasm from images of superficial cells. Network B (69-1) was trained on abnormal nuclei vs. superficial nuclei. Post-processing involved discarding isolated output pixels.)

Network A was trained to differentiate between abnormal nuclei and cytoplasmiclbackground material. Network B was trained to differentiate between abnormal nuclei and superficial cell nuclei. The output images from networks A and B were then combined using the logical AND operator. A pixel was thus judged to belong to a suspicious nucleus only if both networks judged this to be the case.

Network A was trained using 1115 abnormal nuclear pixels and 1115 iteratively selected cytoplasm/background pixels from images of superficial cells. Network B was trained on 279 abnormal nuclear pixels and 279 superficial nuclear pixels. These training data were gathered from 11 different cervical smears. The networks were tested on 98 images taken from 7 smears not used in the training set and containing 110 abnormal cells, 97 superficial cells, 6 intermediate cells and 5 leukocytes. Test set results are shown in table 2.

Network A alone was able to detect 60% of abnormal nuclei. In addition it also detected 36% of the superficial cell nuclei and 8 areas of dark cytoplasm. Combining networks A and B resulted in 40% of abnormal nuclei and just 5 erroneous objects being detected. The post-processing step which discarded single pixel objects eliminated false detections whilst detecting 28% of abnormal nuclei.

4.3 Examining the weights

It is useful to be able to examine the weights learned in order to understand the criteria used by the network. Figure 4 shows two examples of weight masks learned by first-layer units in the specimen enrichment networks A and B. A dark pixel indicates a positive weight and a light pixel a negative weight. The border of uniform shade has been added to each mask to indicate the grey-level which would denote a weight of zero. The central pixel indicates the strength of the weight to the pixel to be classified. The surrounding pixels represent weights on the weight-sharing connections approximating annular rings.

The mask from network A is clearly used to search for dark objects a few pixels in diameter surrounded by a lighter background (i.e. cell nuclei). The mask from network B is used to differentiate between superficial cell nuclei and abnormal nuclei. The output from this unit is positive when an abnormal nucleus is detected. The bright central pixel indicates that

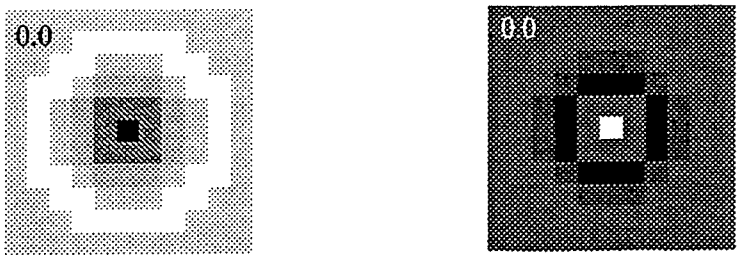

Mask from network A Mask from network B

Figure 4. Example weight masks from specimen enrichment networks A and B.

the small nuclei of superficial cells are usually darker than the nuclei of abnormal cells. The dark ring two pixels from the centre can be interpreted as encoding the information that abnormal nuclei have at least this diameter.

5. Concluding remarks

The proposed method has much in common with the approach introduced by Poulsen. His algorithm relied upon a user to decide how to combine a series of handcrafted image filters consisting of filled and unfilled annular rings (Poulsen 1973, Nordin 1989). The use of a back-propagation neural network has alleviated the need for trial-and-error experimentation by automating the learning of appropriate filters and filter combinations.

The networks coped well with the nuclei detection task. The fact that two abnormal nuclei were missed was not significant. The main concern was to minimise the false detection rate. This was because the vast majority of material on a smear is not of interest. If the false detection rate is not kept very low then the resulting specimen will be too large for detailed analysis.

The networks trained to perform specimen enrichment were able to discard all superficial cells, cytoplasmic material and background clutter as well as the few intermediate cells and leukocytes appearing in the test set. The networks' ability to keep the false detection rate very low has demonstrated the promise of this approach to specimen enrichment. Although it is acceptable to discard some of the abnormal nuclei during specimen enrichment, the detection rate of 28% is perhaps too low to be reliable. These experiments were, however, preliminary and it is hoped that future research will increase this rate.

There are several ways in which performance might be improved. Experimentation with various network topologies has so far been limited. In particular, increasing the receptive field size and using hidden layer slabs with various sizes of receptive field could yield improvements. The images were acquired in colour and as yet only the intensity information has been utilised. Future work will involve expanding the training set to better represent the wide range of material occurring in cervical smears. The method also needs to be tested on a wider range of images.

References

Al I and Ploem J S 1979 Detection of suspicious cells and rejection of artefacts in cervical cytology using the Leydentelevision analysis system. *J.Histochem. Cytochem. 27* 629-634.

Banda-Gamboa H, Ricketts I W, Cairns A Y, Hussein K A 1992 Experimental prescreening systems for automated cervical cytology - a review, Analyt. *Cell. Pathology 4* 25-48

BengtssonE1987 Themeasuringofcellfeatures.Analyt. Quant. *Cytol. 9* 212-217

Erhardt R, Reinhardt E R, Scblipf W, Bloss W H 1980 FAZYTAN: a system for fast automated cell segmentation, cell image analysis and feature extraction based on TV-image pickup and parallel processing. Analyt. Quant. *Cytol. 2* 25-40

Neuromedical Systems, Inc. 1991 The PAPNETTM cytological screening system. In: Data on automated cytology systems as submitted by tbeir developers. Analyt. Quant. *Cytol. Histol.* 13 300-306

Lee J S J, Bannister W I, Kuan L C, Baltels P H, Nelson A C 1992 A processing strategy or automated Papanicolaou smear screening. Analyt. Quant. *Cytol. Histol. 14* 415-425

McKenna S J, Ricketts I W, Cairns A Y, Hussein K A 1993 A comparison of neural network arcbitectures for cervical cell classification. *Proc. 3rd I.CANN. (Brighton, U.K.)* 105-109

Meyer F 1979 Iterative image transformations for an automatic screening of cervical smears, *J. Histochem.Cytochem. 27* 128-135

Meyer F 1986 Automatic screening of cytological specimens. *Computer Vision, Graphics, and Image Processing* 35 356-369

Nordin B 1989 The development of an automatic prescreener for the early detection of cervical cancer: algorithms and implementation. Ph.D. Thesis, University of Uppsala, Uppsala, Sweden

O'Gorman L and Sanderson A C 1983 The converging squares algorithm: an efficient multidimensional peak picking method. I.E.E.E. Int. Conf. on acoustics, speech and signal processing 1 112-115

Parianos E 1991 Automatic segmentation of the nuclei in cervical cell images. M.Sc. Thesis, University of Dundee, Scotland

Poulsen R S 1973 Automated prescreening of cervical cytology specimens. Ph.D. Thesis, McGill University, Montreal

Pycock D and Taylor C J 1980 Use of the MAGISCAN image analyser in automated uterine cancer cytology. Analyt. Quant. Cytol. 2 195-202

Read J S , Borovec R T, Bartels P H, Bibbo M, Puls J H, Reale F R, Taylor J, Wied G L 1979 A fast image processor for locating cell nuclei in uterine specimens. In: Pressman N J and Wied G. L. (eds.) The automation of cancer cytology and cell image analysis, Tutorials of Cytology (Chicago) 143-155

Rumelhart D E, Hinton G E and Williams R J 1986 Learning internal representations by error propagation. In: Rumelhart D E, McClelland J L & the PDP Research Group, Parallel Distributed Processing 1 318-362

Serra J 1989 Image analysis and mathematical morphology 1, Academic Press

Tanaka N , Ueno T, Ikeda H, Ishikawa A, Konoike K, Shimaoka Y, Yamauchi K, Hosoi S, Okamoto Y, Tsunekawa S 1982 CYBEST Model 3: Automated cytologic screening

system for uterine cancer utilizing image analysis processing. Analyt. Quant. Cytol. 4 279-285

Tanaka N , Ueno T, Ikeda H, Ishikawa A, Yamauchi K, Okamoto Y, Hosoi S 1987 CYBEST Model 4: Automated cytologic screening system for uterine cancer utilizing image analysis processing. Analyt. Quant. Cytol. Histol. 9 449-454

Taylor J, Bahr G F, Bartels P H, Bibbo M, Richards D L, Wied G L 1975 Development and evaluation of automatic nucleus finding routines: thresholding of cervical cytology images. Acta Cytologica 19 289-298

Tucker J H 1976 CERVISCAN - An image analysis system for experiments in automatic cervical smear prescreening. Comp. Biomed. Res. 9 93-107

Watanabe S, Imasato Y, Genchi H, Tanaka N, Kashida R 1976 A Pap smear prescreening system: CYBEST. In: Preston K and Onoe M (eds.),Digital Processing of Biomedical Images, Plenum Press, 227-241

Zahniser D J, Wong K L, Brenner J F, Ball H G, Garcia G L and Hutchinson M L 1991 Contextual analysis and intermediate cell markers enhance high-resolution cell image analysis for automated cervical smear diagnosis, Cytometry 12 10-14