MVA '96 IAPR Workshop on Machine Vision Applications, November. 12-14, 1996, Tokyo, Japan

Design of the Cymet A40 Pap Test Prescreener

Ryan Raz, Morphometrix, Toronto, Ontario, Canada

Abstract

The CYMET A40 is an automated Pap test screening system. This instrument combines image acquisition, image processing, classification, mechanical scanning, and slide loading subsystems with very high speed processing electronics to produce an economically viable instrument. This paper is a brief overview of some of the more interesting aspects of this design.

1. Introduction

The CYMET A40, developed at Morphometrix, is an automated Pap test screening system. By combining a powerful image analysis architecture with a high-resolution interactive review procedure the CYMET A40 delivers labour-saving efficiency without compromising diagnostic accuracy.

The Papanicolaou ("Pap") test is a microscope-based screening procedure for the detection of pre-cancerous lesions of the uterine cervix. In the United States alone, over 70 million of these so-called "well-woman" tests are performed annually and this figure continues to grow by 1% per year. The Pap test is perhaps the most successful anti-cancer effort yet devised! Since its introduction in the United States the age-adjusted cervical cancer mortality has dropped by 70%.

Cytotechnologists face many challenges when screening Papanicalaou smears for cervical carcinoma. These professionals must view thousands of microscopic images each working day, marking significant cells for review by a cytopathologist. Abnormal cells may be few, contributing to the tedium of the task. Another challenge is the biological continuum of change that both normal and abnormal cells undergo. This presents a gray zone, where interpretation of cell types becomes subjective.

Early on in our design we realized that the Cymet A40 could not approach the cognitive powers of a trained cytotechnologist, especially as cells and their arrangements are always different. However, what the instrument could do was to act as a high speed filter which stores and then displays for review images which have suspicious cells or groups of cells. If there are no areas for review the slide would then be passed as normal. This reduces the skilled labour needed by over 50%. It also reduces cytotechnologist fatigue as he or she only has to examine the interesting areas on the slide.

Even this seemingly simpler task presents a very difficult problem in image processing and cell classification. As well, scanning, image acquisition and processing speeds must be matched. And finally the result must be an instrument that can be produced economically. In responding to this challenge, Morphometrix has assembled a diverse and multi-disiplined design team.

II. Mechanical and Imaging System (fig. 1)

Highlights of the CYMET A40 mechanical and imaging system include:

- A camera consisting of three direct to digital high resolution CCD image arrays, each with a 10 Mbyte/sec.

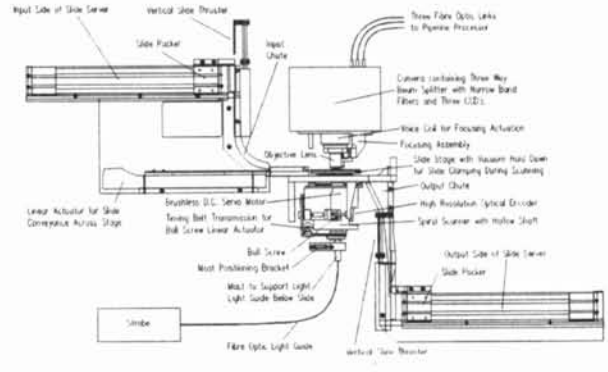

Fig. 1. The Imaging Assembly

fiber optic link. A beam splitting prism combined with narrow band filters provides aligned high resolution multi spectral images to the rest of the system.

- A microscope with a high-speed motion and sub micron focusing system using a voice coil drive. The voice coil moves an infinity corrected microscope objective lens to the image focal plane. The correct focal plane is determined by the electronics using a specialized image pipeline processor. The voice coil is activated with feed forward control. The total time for a point to point move is under 10 milliseconds. This is at a frequency which is is about four times its resonance frequency of 25 Hz.
- A spiral scanning table. This table is light weight with the circular motion being driven by a brushless servo motor which has sinusoidal commutation using the feed back from a 5000 lines per revolution optical encoder. This combination results in ultra smooth motion at speed from as slow as 0.1 revolutions per second up to 15 revolutions per second. It also has a repeatable stop and go position accuracy of better then 20 microns .

Image Processing electronics (fig. 2)

We have developed a special purpose fine-grained pipeline processor for image processing. It is over 300 times faster than a 133 MHz Pentium PC. This pipeline handles all the image processing functions up to the classifier stage. The classifiers are implemented in a traditional general purpose processor environment.

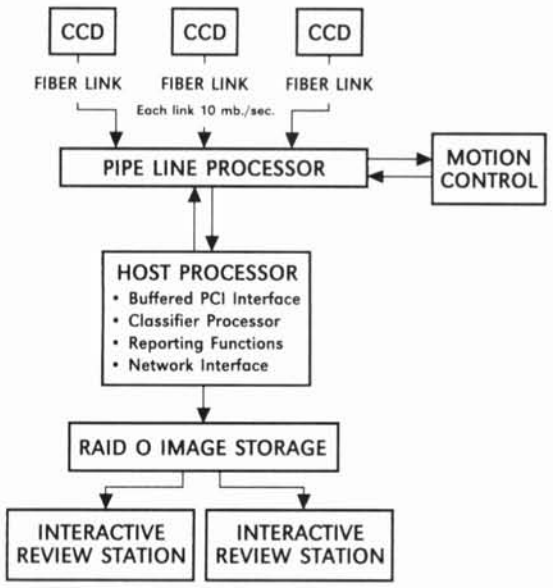

Fig. 2. Image Processing Electronics

The pipeline processor consists of ten different modules. Each module utilizes a few major components: Video Dram for image storage, variable length shift registers as digital delays, flash memory for look up tables and FPGAs for computational and control functions. In two instances, we also use a special function labeling LSI chip from Sumitomo Metals and a histograming LSI chip from Harris Semiconductor.

Figure 3 is a simplified example of a submodule to implement a Sobel filter. All algorithms were first modeled and simulated in software using C. Within the C code, whenever possible, math functions were replaced by look up tables. This C code was used to define the module architecture. Furthermore, the relevant parts of the code were converted to VHDL code and then implemented within the FPGAs.

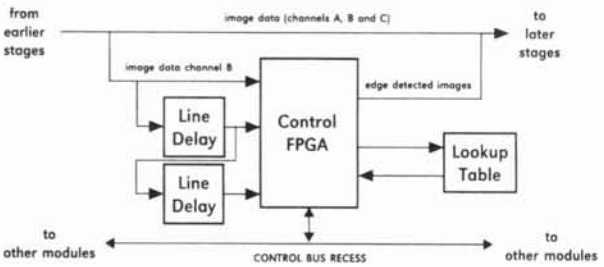

Fig. 3. The Sobel filter

III. Multispectral Image Analysis (fig. 4)

There are four stages to image analysis:

- The images must be acquired and digitized.
- The image must be segmented and labeled to isolate the different cells and differentiate the cytoplasm from the nucleus.
- These isolated regions must be measured and characterized in the form of mathematical features.
- These features are then used to classify the type of cell.

A. Segmentation

Image segmentation is one of the most important and difficult stages of analysis. Humans are naturally good at

segmentation. However, machines have difficulty with the random nature of cell clusters, overlapping cells and the varying cytoplasm and nucleus densities of different cell types. To improve the segmentation and the subsequent classification steps, Morphometrix has developed a high speed multi-spectral camera system. The camera has three channels utilizing narrow band filters. These filters have

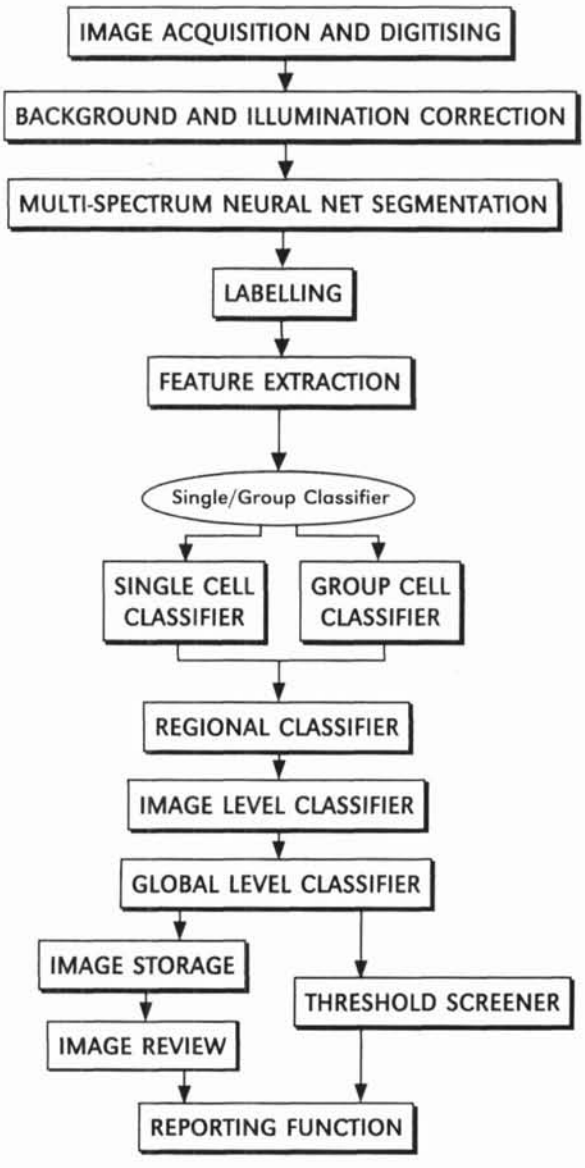

Fig. 4 CYMET A40 Image Analysis

been selected to use the spectral properties of standard Pap stains to differentiate between nucleus and cytoplasm.

As a result of segmentation the multispectral images are divided into two classes: cytoplasmic and nuclear, separated by the multidimensional threshold t (in our case we use 3D space comprised of the 3 channels of 8 bit data from our multi-spectral camera).

Studying the nucleus/cytoplasm pixel distribution in the 3D space we've found that this distribution is extremely complex. It contains many clusters and overlapped regions. Therefore it is obvious that the optimal threshold will have a very complex nonlinear surface in 3D space.

Many studies have been carried out using different approaches, but we've found that the best way to solve this problem is to use the multidimensional Neural Net segmentation technique. In this case the thresholding is performed by the Neural Net Bayesian Confidence Measure calculations. The Neural Net algorithm should manage the following successfully:

- Fast training for large amounts of data (about a million samples)
- Handling of a multimodal non-Gaussian data distribution
- Good generalization simultaneous with greater sensitivity to the small clusters of patterns representing the useful cell subclasses
- Handle unbalanced data sets.

In all the diversity of Neural Net algorithms we haven't found any which successfully satisfy all of these conditions. Therefore, we've designed a special Neural Net algorithm called the Probability Projection Neural Net (PPNN). We've found that the PPNN matches all the above listed conditions and has better classification results than other algorithms. The PPNN belongs to the so called Probability Density Function (PDF) estimator type of Neural Network and therefore it can be used as a PDF estimator or as a general classifier in pattern recognition. The PPNN uses training data to create an N dimensional PDF array which in turn is used to estimate the likelihood of a feature vector being within a given class. To create and train the PPN network we need to partition the input space to three discrete nodes, the equivalent of a three dimensional grid in space. Then we map (project) the influence of each training pattern to the neighbor nodes. For feedforward calculations we use the reverse mapping technique.

For the hardware implementation of this classifier we store the results in a large flash memory which then acts as a high speed associative memory look up table.

We've compared the PPNN, Back propagation (BP), Elliptic Basic Functions (EBF), and Learning Vector Quantization (LQV) neural network performances for the cell segmentation problem. The data contains approximately 765,000 color pixels which belong to two classes:

- Nuclear pixels (100,000 records)
- Cytoplasmic pixels (660,000 records)

The performance results are as follows:

- The training process for the BP, EBF, and LQV nets are approximately 500 times slower than for the PPNN
- The recognition rates are:

TABLE I

NET TYPE	FPOS %	FNEG %	FPOS + FNEG
Backprop. (1 layer, 5 nodes)	5.19	4.52	9.71
EBF (7 clusters)	4.63	2.38	7.01
LQV (200 code vectors)	7.88	4.80	12.68
PPN (1 nuc + 1 cyt clusters)	4.19	1.94	6.13
PPN (4 nuc + 2 cyt clusters)	3.77	2.22	5.99

* FPOS: false positive rate; FNEG: false negative rate

As we can see, the PPNN with multicluster normalization has 2.3% better performance than the simple PPNN and 17-111% better performance than other neural nets.

B. Cell Classification

Cell classification involves essentially two tasks. The first is to determine, with as small an error as possible, the nature (cell type, abnormality, etc.) of each segmented cell in a digitized image. This a purely cellular level classification which does not use any context information from the neighborhood of the cell. The second task is then to exploit the neighborhood information of the cell to make a statistical inference based on the classification results for each individual cell within this relatively large region. This task is particularly relevant for groups of cells which occur very frequently in Pap slides.

It is well known that cell classification, in particular early stage abnormality detection, is a rather difficult task. First of all, the variations between some types of cells are so minor and subtle that even highly trained cytotechnologists and pathologists may draw different conclusions on the same cell. Secondly, due to the quantization error inherent in any digital processing system, the objects to be processed by a computerized cell classification system can be rather noisy. Experiments show that some traditional pattern classification techniques and tools, such as the popular multi-layer feed forward neural networks based on back-propagation learning algorithms can not provide satisfactory classification results for cell classification. This is because the data from different cell classes are heavily overlapped. Thus standard methods which directly use sample data (which is usually unbalanced) seem to be too simplistic. More sophisticated classification algorithms have to be invented to overcome the fuzzy nature of cells and data balance problems.

Morphometrix has developed a novel classification scheme which is capable of recognizing a variety of different types of cells within a very confusing environment. Since the distinction between each class of cells is often blurred, this scheme carefully exploits the morphological changes of cervical cells. Each class of cell is assigned to an individual fuzzy set. Each fuzzy set is then modeled through a statistically supervised learning procedure. To classify an unknown cell it's membership in each set is calculated from the model. Using these membership values as inputs to a subsequent classifier the likely nature of a cell is determined.

Characterization of Classifier (fig. 5)

The graph below shows an example of the actual classification results from our database using an EBF classifier. There are three curves in this graph: curves A represents a false negative curve derived from an abnormal class, curve B represents a false positive curve derived from a normal class, and curve C is a confidence curve derived from curves A and B. The x-axis in the graph represents the membership values coming out of the EBF classifier. Thus, as indicated in the graph, at each given x position (which can be considered as a classification threshold), the corresponding y value of curve A is the false negative rate (i.e., the percentage of abnormal cells classified as normal) and the corresponding y-value of curve B is the false positive rate (i.e., the percentage of normal cells classified as abnormal).

The curve C provides a new way to characterize the classifier. Since it is derived from the curves A and B with all data involved, it provides the confidence level associated with each membership value relative to the whole data set. For any membership value x, we can now transform it into a new value according to this confidence curve. The transformed value provides the true normalized membership in terms of its relative position in the whole data set.

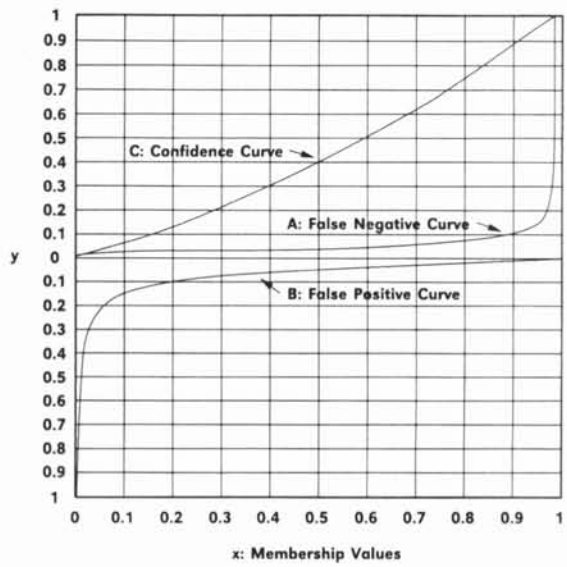

Fig.5 Classification Results