Machine vision for automated optical recognition and classification of pollen grains or other singulated microscopic objects.

G.P.Allen[1], R.M.Hodgson[1], S.R.Marsland[1] J.R.Flenley[2]
1 Massey University, School of Engineering and Advanced Technology.
2 Massey University, Geography Programme, School of People, Environment and Planning.

Email: r.m.hodgson@massey.ac.nz

Abstract – The location and identification of singulated objects on microscope slides is a problem that is common to many applications, including recognition of pollen. In this paper, we describe a working system to solve this problem and demonstrate that it can be used to effectively locate pollen grains on slides, focus on them, photograph them, and then identify them based on a trained neural network. Our system aims to remove the need for laborious, time-consuming, and inaccurate counting of pollen grains by humans with a low-cost machine solution. It can deal with slides obtained using different preparation techniques and media. As well as describing the system, we present positive test results, including a comparision with human experts on the classification and counting of pollen on slides.

I. POLLEN CLASSIFICATION

Palynology, the study of pollen and spores, is an important tool for many fields of study including climate change, plant radiation and evolution, air pollen counts for inhalant allergy sufferers, honey typing, archaeology, and forensics. As is reported in [1], an automated pollen identification and counting system would reduce the laborious counting required by highly skilled people involved in palynological endeavours (there is an estimated 30 months microscope work in a PhD).

The broad requirements of such a system are to locate pollen grains on a microscope slide, photograph them and identify features that will enable them to be classified into taxonomic categories at reasonable cost, and with a success rate at least that of a skilled person. The saving is labour, and time consumed by people with skills that could be better applied to less mundane tasks. The difficulties are the variety in appearance of pollen, the requirement that the system generalise over different preparations and slide media, and the fact that other foreign matter may also be present on the slides.

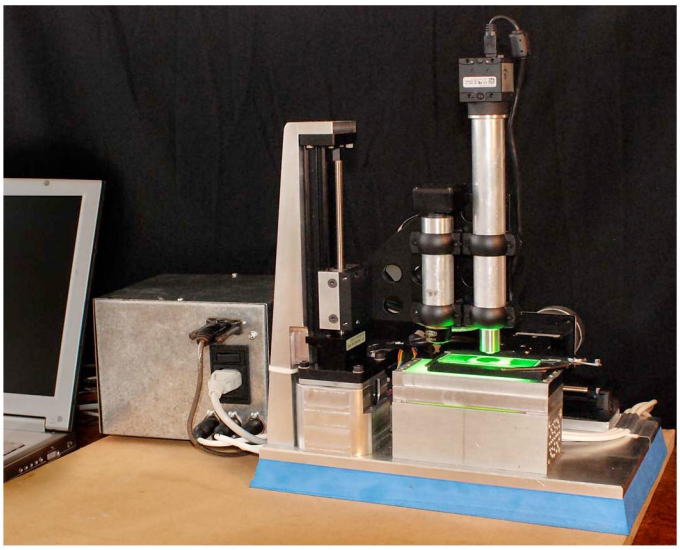

Figure 1: The System; Second Prototype

Different pollen vary in size and appearance, depending upon factors such as whether they are air dispersed or animal dispersed. Typically, they have diameters from around 5µm to 120µm, with variation within a taxon of up to 20µm. We have selected a mid-range for the sizes of pollen we will focus on, of 10µm to 100µm diameter. While there is a wide variety of pollen appearances, there are various related taxa where it is very difficult to distinguish them by eye using morphological clues. For example, see the images of different grasses shown in fFigure 9-11.

Pollen can be captured in many different ways, from core samples to air capture. Samples will include sediments and/or detritus. Further processing can be from dilution of samples in water and centrifuging to chemical treatments to remove silica and plant material, filtration at filter mesh sizes above and below 100µm. The pollen are usually placed in suspension on a slide in the manner already standardised by palynologists. No attempt has been made to automate this, rather our system uses the same slides as are used by humans, although there are

two requirements that we have found to be important, namely the cleaning and filtering of samples to minimise detritus, and diluting so that the pollen grains are unlikely to form clumps.

II. Automated System Description

Our system locates pollen grains on a slide and captures high magnification images of them, together with their location on the slide. A set of image features are extracted from the high resolution images and are used for classification of pollen types, enabling a count of the number of grains of each pollen type.

Once the user has processed slides in the same way as they would for human inspection, they simply need to place a slide on the holder and, using the computer control, indicate the general area of the slide containing pollen, together with an initial focus. The machine will then process the entire slide, locating individual pollen grains on the slide, saving a tightly cropped image of each, and extracting a set of image features that are used to identify the pollen in each image, based on a pre-trained neural network. In this way, the slide is ultimately sorted by pollen taxa, with the quantity of each shown pollen found on the slide being counted.

The individual components of the system are:

- a 'machine' to capture the images (Figure 2)
  - Two microscopes (§*Two Microscopes*)
  - Lighting (§*The Lighting*)
  - Movements (§*Movements*)
- auto-focus algorithms (§*Auto-Focus*)
- segmentation algorithms (§*Segmentation*)
- classification algorithms (§*Classification*)
- a computer to run the algorithms and control hardware (§*The Computer*)

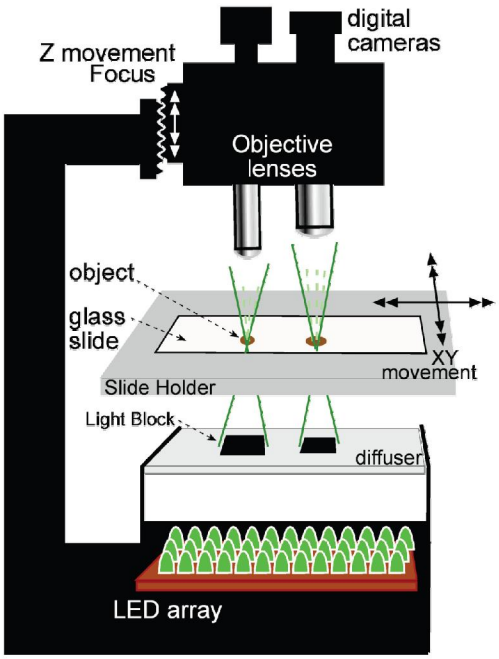

Figure 2: System Elements

The 'machine', is an XY stage with attached slide holder. Two purpose-designed digital microscopes are solidly mounted above a filtered light source. As transmission lighting is used, the slide sits on an aperture in the XY stage positioned between the cameras and light source, as indicated in Figure 2.

There are power supplies for lighting and stepper motors. Two motors move the XY stage to locate pollen under the microscope and a third motor adjusts the relative height of the cameras for focussing.

Two Microscopes

Two microscopes are used, rather than one with variable magnification, because it is the simplest mechanical option. A wide angle image is required to locate pollen by viewing as large an area of the slide as possible while maintaining enough magnification to identify pollen and reject detritus. A high magnification microscope is required to obtain all the information available given the resolution of the optics.

Much of the complexity of optical microscopes is to obtain an image suitable for the human eye. By using a digital sensor instead, the magnification required is reduced and the depth of field increased. Image resolution is limited in theory by the wavelength of light, and any diffraction (caused by the lens diameters and any other aperture restrictions in lighting and lenses) reduces the resolution. Maximal resolution is required to capture as much detail of surface texture as possible in the images, as texture is an important part of pollen classification.

The resolution for the low resolution camera was calculated to be 0.9µm, hence the Nyquist sampling value should be 0.45µm or less. The sample size of the sensor is 4.65µm and this is magnified optically. For the high resolution camera, 11.2x magnification was achieved by using a 10x objective with a slightly longer tube length than standard. This gave 4.65/11.2 = 0.415µm sampling, which is just less than the required upper limit. For the high resolution camera, a higher numerical aperture lens is used with 16x magnification and the correct tube length (to attain as high quality as possible of all the characteristics that the lens was designed for) giving higher resolution and improved sampling rate.

The wide angle, lower magnification microscope with its large field of view, locates pollen grains quickly, since it requires few movements and image captures. As this microscope has less magnification and therefore more depth of field, the image will remain in focus as the microscope moves across the entire slide. This, of course, requires the stage to be orthogonal to the optical axis to within about 30µm.

The image captured by at low resolution is used to identify the location of possible pollen grains on the slide. The high magnification microscope is then moved to each of these possible locations, and on finding an acceptable object, an image slightly larger than the object bounding rectangle is produced. The image (an example of which is shown in Figure 4) is stored for feature extraction and classification, and the microscope is moved to the next location identified.

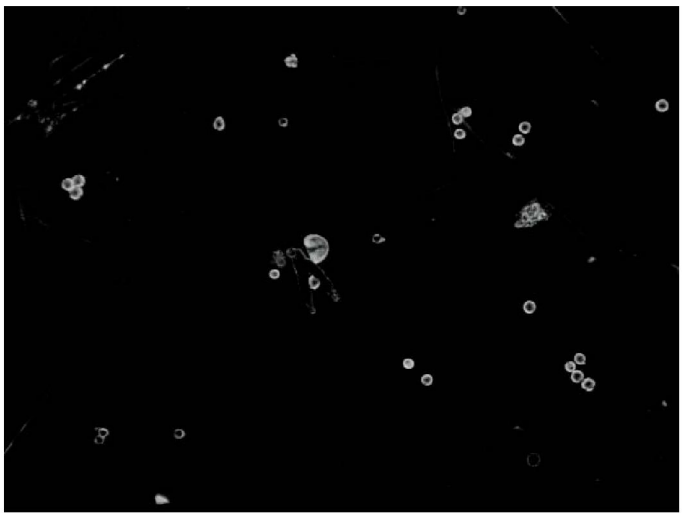

Figure 3: Wide angle image of pollen on a slide with a Pinus pollen shown near the centre.

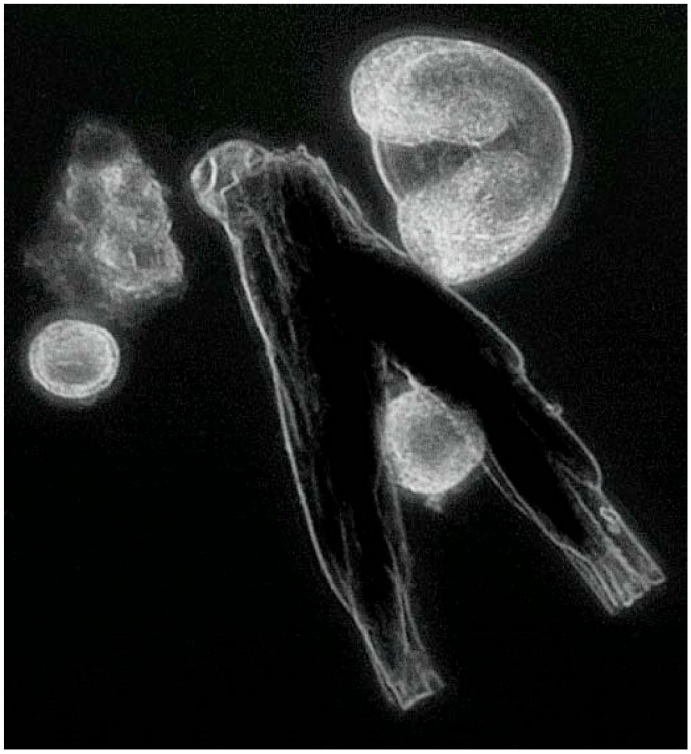

Figure 4: High magnification image of the centre (cropped) of Figure 3. There are four pollen grains and detritus in this image.

THE LIGHTING

Lighting is provided by a matrix of 60 high-intensity, narrow-beam LEDs, arranged in a plane perpendicular to the optical axes of the microscopes. Two diffusion filters even up the distribution of light over the area of the final filter (see Figure 5, Figure 6, & Figure 7). One filter is almost directly on the LEDs and the second about 7mm above. Spectrally limited LEDs are used to minimise the bandwidth of the light, thus reducing chromatic distortion in the objective lens. Green was chosen because the microscope sensors have a peak sensitivity at about 550nm; the wavelength of yellow/green light. A similar arrangement using incandescent lamps and a yellow/green filter was used for P1, but even though cooled with a fan, heat transfer to the slide could cause convection currents in some less viscous suspensions used, and thus slight movements of pollen on the slide. The light is placed in-line with the optical axis of the microscope objectives for simplicity, ease of manufacture and cost reduction.

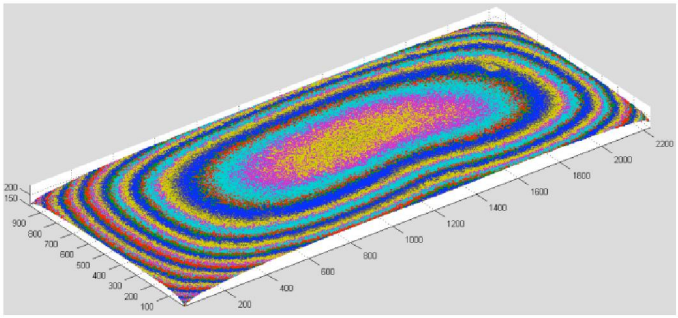

Figure 5: Intensity variation map of lighting

Figure 6: Pixel values across length at centre of lighting.

Figure 7: Pixel values across length at lower 1/3 of lighting

The diffusion filter with a light blocking object below each objective lens produces a simple form of "dark field" illumination. The angle of light allowed by the light block into each lens is greater than the acceptance angle of each lens so that no direct light reaches the sensor. Only light reflected, refracted or diffracted by the object into the lens is 'seen' by the sensor. Thus the name "dark-field" and the resulting black images with white objects, as seen in Figure 3 and Figure 4.

Dark-field imaging is known for producing greater contrast in an image [2]. Another reason for choosing the dark-field lighting was that translucent pollen showed up brighter and in greater contrast against the more opaque detritus as demonstrated in Figure 8. This should improve segmentation.

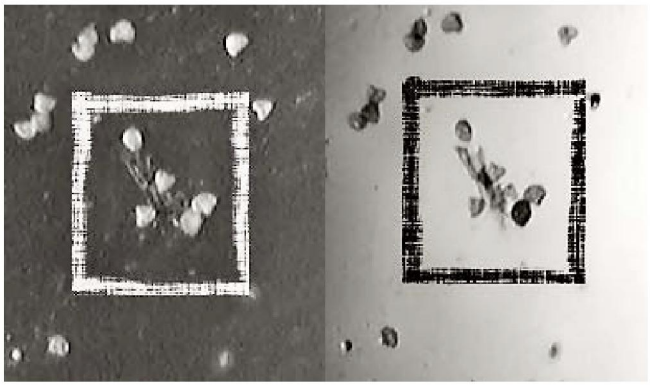

Figure 8: An image shown under dark-field (left) & bright-field lighting, showing visual differentiation of pollen against detritus greater with the dark-field lighting.

MOVEMENTS

The slide is held onto, and moved by, a slide holder mounted onto two commercial precision linear movements, which are driven by two stepper motors. The micro-stepped motors allow linear movements of 0.6µm per step in our current machine (the smallest pollen of interest is about 10µm across). The field of view of the high magnification camera is 465 x 348 steps.

We originally used a stepper motor driven directly from the shaft of a rack and pinion mechanism taken from a commercial microscope. This afforded a barely acceptable, approximately 10µm step size. However, this has since been improved by using the same linear movement and stepper motor as the XY stage with the same minimum 0.6µm step length, with significant improvements.

Precision and repeatability in the movements is reduced by stepper motor micro-stepping inaccuracies, Mechanical 'stickiness' and other repeatability limitations inherent in the linear movements. The horizontal movements require about 20 steps to overcome backlash. The focus has gravity to help and movements to final locations are always upward. A predetermined final direction could also be used for horizontal movements to improve repeatability.

AUTO-FOCUS

The user sets an initial focus for reference and the system sets the low magnification microscope, based on that setting, for the one focus setting required to capture images across the slide with the wide angle microscope for object location (such as the example shown in Figure 3).

The high magnification images are automatically focussed by the system for each pollen grain. At present, a central portion of the image is used for focus and after segmentation, an area bounding the object to be captured is used for focus. The mixture of text commands is sent via RS232 to motor drivers; together with the use of the interpreted Matlab programming language make this focussing very slow, so development is underway to capture images rapidly as the motor drives though focus, storing them all and selecting the image in which each segmented images is best focussed.

We investigated several varieties of focus measure, including normalised variance, power, histogram, entropy and Fourier transform, some of which are described in [3-5]. After experimentation with the various methods on multiple pollen slides, we settled on the standard deviation for wide angle images and maximum-gradient-squared for high magnification images. This latter function calculates the gray values gradient squared of the image in two orthogonal directions and returns the maximum value. This often performed as well as the Fourier transform, but was faster and more consistent over a variety of images.

SEGMENTATION

Segmentation is performed in both the wide angle and high magnification images. Segmentation in wide angle images finds objects on the slide and determines their locations on the slide, while ignoring objects that are likely not pollen grains. Segmentation in the high magnification image is a second opportunity to rule out unwanted objects using images of higher resolution, and for final image capture within the object bounding rectangle.

The initial segmentation (for images as in Figure 3) subtracts a background image (taken without a slide in place to account for lighting variation), calculates a gray level threshold from the background image and finds edges using a Canny edge detector with the background image gray level threshold as a parameter. All end points of lines are found and if they are less than a certain distance apart they are joined. Morphological 'close', then 'dilate' are performed and any areas totally enclosed by lines are filled with white. Morphological 'open' then 'close' are performed. The image is now black with white blobs where objects were.

Objects are now removed by a series of tests. Objects too large or small (as measured using a pixel count of blobs) are removed. The area of bounding rectangles is tested for ratio of smallest side to largest side and the object rejected depending on value. The ratio of area of a blob to area of its bounding rectangle is calculated to remove further un-pollen-like objects. The ratio of the area of a blob to the area of its convex hull is calculated and the blob rejected depending on the value. Remaining blobs are listed according to their centre of bounding rectangle location within the image. The image is one of a series of overlapping images so the locations are converted from pixel distances within the image to motor step distances within the whole series of images. The list of locations for the entire series of images is then stored for sending the high magnification microscope to each.

When the high magnification microscope arrives at a location, a pollen grain is expected near its centre. From an accumulation of errors, there is an area of uncertainty of about 500x500 pixels (out of 1024x 768). A similar segmentation algorithm to the wide angle microscope is performed and a number of objects may be discovered. Their locations are

recorded and they are compared to locations of already captured images and if found to be the same within a certain tolerance, they are rejected as already captured. All remaining objects are captured as close-cropped images and saved onto disk, and their locations added to the list of captured objects. The images are then ready for feature extraction and presentation to a trained neural network for recognition, sorting and counting. Locations of each image are stored so it is possible to go back to a particular pollen image if necessary.

CLASSIFICATION

There are many and varied methods of classification in the literature. We currently use a standard multi-layer Perceptron neural network trained by back-propagation of error [6], in line with studies done by Li and Flenley [7]. This requires training using images of expertly identified pollen grains. A set of 43 image features are extracted from each high resolution image and these are used as input to the neural network. The features comprise measures of pollen size and shape, and most significantly various texture measures. They were first suggested in [8], and have proved to give excellent results when tested with fresh pollen acquired directly from known plants.

THE COMPUTER

The computer used is a PC with a 2.6GHz processor and 1Gbytes of RAM running Windows XP professional. All of the code is written in Matlab.

III. EXPERIMENTS AND RESULTS

The tests performed were of two major categories:

- Classification module accuracy
- Whole-system compared to human classifiers

For the classification accuracy, two image databases were compiled:

1. images captured using a conventional microscope
2. images captured using the system developed here

The images for these two sets were captured using the same reference pollen slides.

Images of each pollen type were required for training the neural network, and a separate set for testing the trained neural network for accuracy of classification. A random selection from within each database of images was made of 50% for training, 25% for validation (repeated testing to improve neural net parameters until optimal for the dataset) and 25% for a test set *used only once* for the final test. The final tests were each a series of 5 tests where image features were presented in varying random order to the neural network with the results over the 5 tests being averaged. The validation set was used with the training set to adjust neural net parameters for optimum results and verify the system working. The training and validation sets were then combined for training, and the test set used for the final test. The feature sets extracted from the images, were presented in random order to the classification software.

Percentage success rates given as results are the number of correctly classified images divided by the total number of images multiplied by 100.

CLASSIFICATION OF STAINED VS UNSTAINED IMAGES

Staining is often used in pollen preparation to add contrast and colour to the pollen grains in order to aid identification by humans. To determine whether staining is important to the classification module, tests were conducted resulting in 96% accuracy for stained pollen and 93% for unstained pollen, indicating that stained pollen should give slightly better results.

COMPARE THE SYSTEM ACQUIRED IMAGES WITH CONVENTIONAL MICROSCOPE IMAGES

When images obtained from our system and images from the same slides obtained from a high quality conventional microscope were tested using the classification module, in 5 tests, the average correctly identified for the conventional microscope was 94% while the system presented achieved 98%.

CLASSIFICATION OF GRASS POLLENS

The aim of this experiment was to check the performance of the system when classifying grass pollens that are commonly counted as one type as many are very difficult to distinguish manually under a light microscope.

Using three pollen types with images captured on the system, the classifier identified 90% correctly.

Test description: the three grasses tested were,

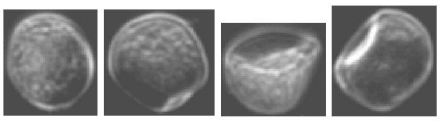

Figure 9: Brown-top grass pollen

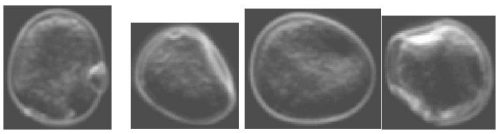

Figure 10: Cocksfoot grass pollen

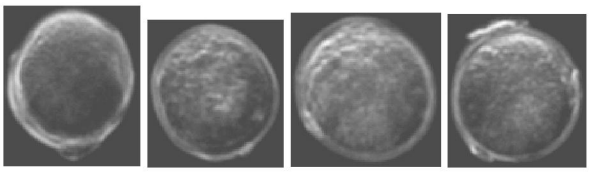

Figure 11: Philaris grass pollen

THE SYSTEM COMPARED TO EXPERTS

The aim of this experiment was to compare the result of the total process of counting pollen on slides by the system, with the resultant count of the same slides by 5 experts.

Test description: Four slides, each with the same 6 pollen types are prepared. Five 'experts' including two professors, a

post-doctoral researcher, a technician working in palynology and an honours student, count the slides. The system is then set to count each of the 4 slides, four times.

Result. The graphs below show the mean and range of P = person compared to M = machine for each pollen type on all four slides. For counts, on the 'Y' axis, of pollen-type for person and machine on the 'X' axis.

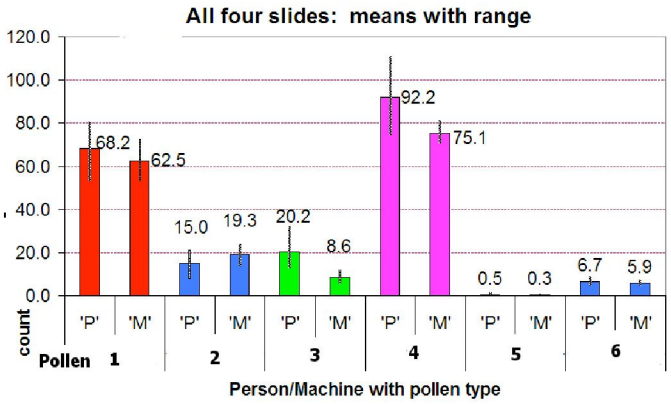

Figure 12: Means with range. Pollen 1, for example, in red: persons counted 68.2 and the machine counted 62.5.

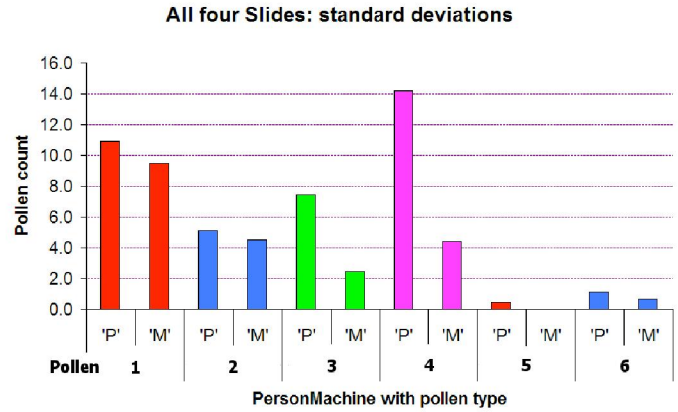

Figure 13: Standard deviations: showing variation in results between counts of the same pollen type to be lower for machines compared to humans for all pollen types.

IV. Discussion:

We have presented a system for the automatic imaging, identification, and counting of pollen on convential microscope slides. Our system takes slides as input and produces as output a collection of images of pollen, together with a count of each identified type. In a comparison with human experts on the same images, our system produces comparable output, with significantly less variation that different humans. In addition, the system produces classification results that are at least as good as known published results.

We thus believe that the system promises to be useful to palynologists in the laboratory.

The various components of the system all appear to work well both in isolation and together; for example the XY stage, with movement limits larger than a slide, a repeatability of position of 20 microns, speed in excess of 10mm per second, and a spatial resolution of 2.6 microns, would be satisfactory for a manufactured product.

The system functions well, and promises to meet the requirements to be useful to a palynologist in the laboratory, all the more impressive since the component costs of the system were under $NZ 7,000 including the computer.

V. Acknowledgements

Many thanks to Steve Denby and crew at the mechanical workshop of the Institute of Fundamental Sciences, Massey University, who built the machine. Thanks also to, Xiuying Zou for conventional image capture of reference pollens.

VI. References

[1] E. C. Stillman and J. R. Flenley, "The needs and prospects for automation in palynology," *Quaternary Science Reviews*, vol. 15, pp. 1-5, 1996.

[2] P. C. Montgomery and J. P. Fillard, "Study of microdefects in near-surface and interior of III-V compound wafers by dark-field transmission microscopy," *Electronics Letters*, vol. 24, pp. 789-790, 1988.

[3] F. C. A. Groen, I. T. Young, and G. Ligthart, "A comparison of different focus functions for use in autofocus algorithms," *Cytometry*, vol. 6, pp. 81-91, 1985.

[4] J.-M. Geusebroek, F. Cornelissen, W. M. Arnold, and H. G. Smeulders, "Robust autofocusing in microscopy," *Cytometry*, vol. 39, pp. 1-9, 2000.

[5] A. Santos, C. O. De Solorzano, J. J. Vaquero, J. M. Pena, N. Malpica, and F. Del Pozo, "Evaluation of autofocus functions in molecular cytogenetic analysis," *Journal of Microscopy-Oxford*, vol. 188, pp. 264-272, 1997.

[6] I. T. Nabney, "NETLAB," 2003.

[7] P. Li and J. R. Flenley, "Pollen texture identification using neural networks," *Grana*, vol. 38, pp. 59-64, 1999.

[8] Y. Zhang, "Pollen Discrimination Using Image Analysis," Massey University, Palmerston North report, 2001-2003 2003.