The Application of Intelligent Compression to Telepathology†

M.J.Varga[1], P. G. Ducksbury[1], F.J. Leong[2], J. McGee[2]

[1] Defence Evaluation & Research Agency, St Andrews Road, Malvern, Worcs, WR14 3PS, UK.

[2] University of Oxford, Nuffield Department of Clinical Laboratory Science, Level 4, Academic Block, John Radcliffe Hospital, Headington, Oxford, OX3 9DU, U.K.

ABSTRACT

Telepatholgy is a means of practicing pathology at a distance, viewing images on a computer display rather than through a light microscope. To date the use of compressed images in pathology remains controversial. This is because commercial image compression algorithms such as JPEG achieve data compression without knowledge of the diagnostic content. Often images are lossily compressed at the expense of corrupting informative content/features.

None of the currently available lossy compression techniques are concerned with what information has been preserved and what data has been discarded. Their sole objective is to compress. By contrast, this paper presents a novel image compression technique, which exploits knowledge of the slide's, diagnostic content.

This 'intelligent' approach combines different compression ratios of a lossy compression technique to different image locations, applying each in the appropriate context across an image so as to maintain 'diagnostic' information while still maximising the possible compression. A standard compression algorithm is used, but it is used in a context sensitive manner that can offer high compression ratios and preservation of diagnostically important information. The precise gain in the compression performance depends on the application (*e.g. initial diagnosis or second opinion consultation*) and the diagnostic content of the images. This 'intelligent' approach could improve the performance of the pathologists as well as the patient's prognosis.

† This Defence Technology for Healthcare project is funded by the UK Treasury under the Capital Modernisation Funding Programme

Introduction

This is the Histopathology Image Management work-package that is part of a Defence Technology for Healthcare project. The aims of the work-package include:

> Developing the capability for processing, analysing and annotating diagnostic features in histopathological images from a virtual microscope environment for cancer management.
>
> Developing an intelligent 'content based' compression approach in telemedicine for consultation and second opinion diagnosis which exploits domain knowledge and diagnostic information of the image
>
> Integration of annotated images in electronic patient records (EPR)
>
> To clinically trial the software and system, comparing them under controlled conditions with existing approaches.

This presentation is concerned with 'intelligent compression' and how techniques developed for defence purposes are also applicable in the medical domain. It is recognized that there is a substantial military and clinical need for high levels of image compression so as to transmit the required imagery across the available bandwidth. At present, in order to achieve high levels of compression it is necessary to use the so called 'lossy' coders (in which some data/information is discarded or lost). Unfortunately, as has been noted in the survey carried as part of the defence project, commercial coders which are developed largely for broadcasting or telecommunications, rely on there being an intelligent (i.e. human) receiver who can tolerate the resulting quality degradation or corruption. However, the same corruption makes processes such as target/tumor detection difficult and at times impossible. It can also have a significant impact on the critical aspects of the human based image interpretation function.

None of the currently available lossy compression techniques address the issues of what information has been discarded. Their objective is to compress and not to acknowledge/address the need to preserve relevant information through the compression or the need to identify irrelevant data to discard.

Techniques for Content Based Image Compression

The aim is to develop both a conceptual framework and a software realisation of the framework, which incorporates the cueing process, the compression and decompression and finally performance evaluation. The aim is that the framework should, as far as practicable, be application independent.

Transmission

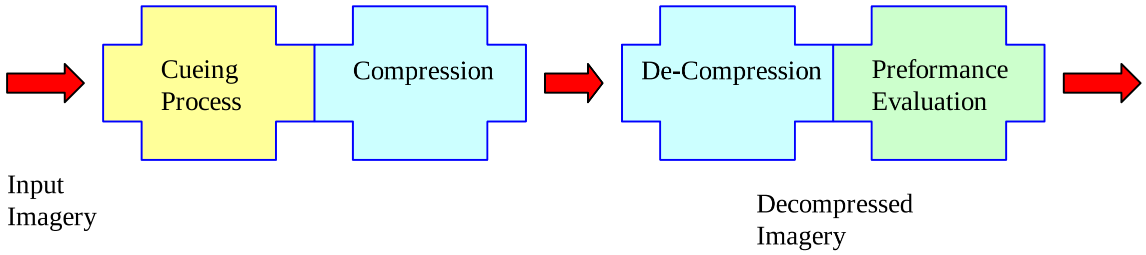

Cueing Process

The aim of the cueing process is to provide essential and reliable guidance for the application of lossless and lossy compression techniques intelligently on the regions of diagnostic interests and redundant areas in the slide respectively.

The objective of cueing is to :

- Maximise the detection of 'True +', i.e. diagnostic areas for which lossless (or near lossless) compression must be used;
- Maximise the detection of 'True - ', i.e. real redundant areas for which lossy compression can be used;
- Minimise all the 'False +' and 'False -' , i.e. mistaken targets and background.

Cueing process:

- Manual annotation
- Automated
- Interactive

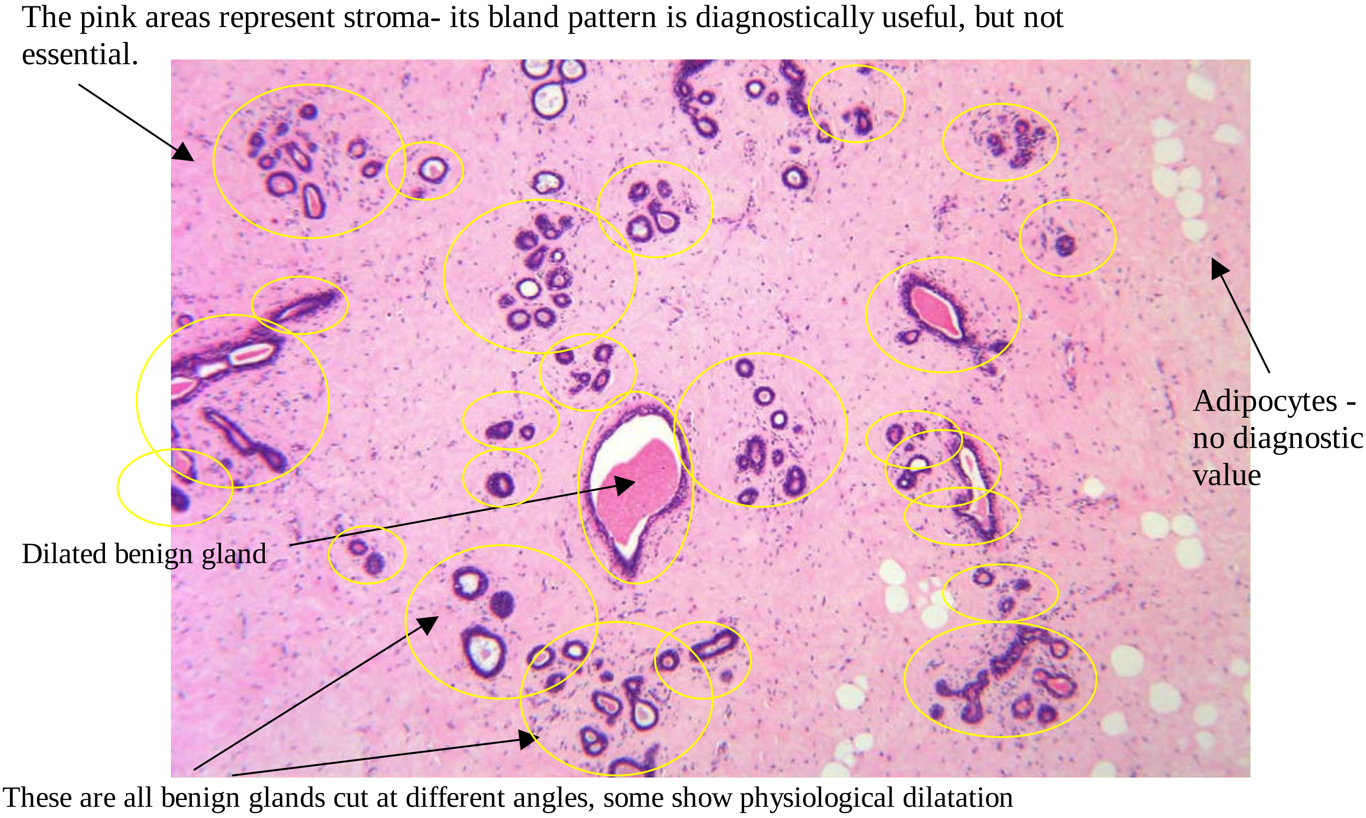

Figure 1 : Manual Annotation of Benign Breast tissue by histopathologists (JFL)

Figure 2: Manual Annotation of a Military Image by a Photographic Interpreter

Compression and Decompression

The aim is to develop techniques that exploit standard image compression algorithms both in combination and in isolation. In the former case, standard lossless compression algorithms may be used to represent diagnostic regions and standard lossy compression algorithms to represent redundant regions. In the latter case, single algorithms (such as wavelet techniques) can be modified to concentrate the compression in different parts of an image.

Standard Compression Approaches

Figure 3 provides an overview of the effects of applying a number of different standard compression algorithms verses different compression ratios. It can be seen that the standard JPEG suffers from the noticeable blocky artifacts as the compression is increased. The fractal appears to degrade more gracefully but notice that it is not just the background that is 'smoothed' out, eg the boundary of the light structure in the lower left corner. The wavelet based approach however appears to achieve a better overall compression when compared to the other two approaches.

JPEG

Fractal

Wavelet

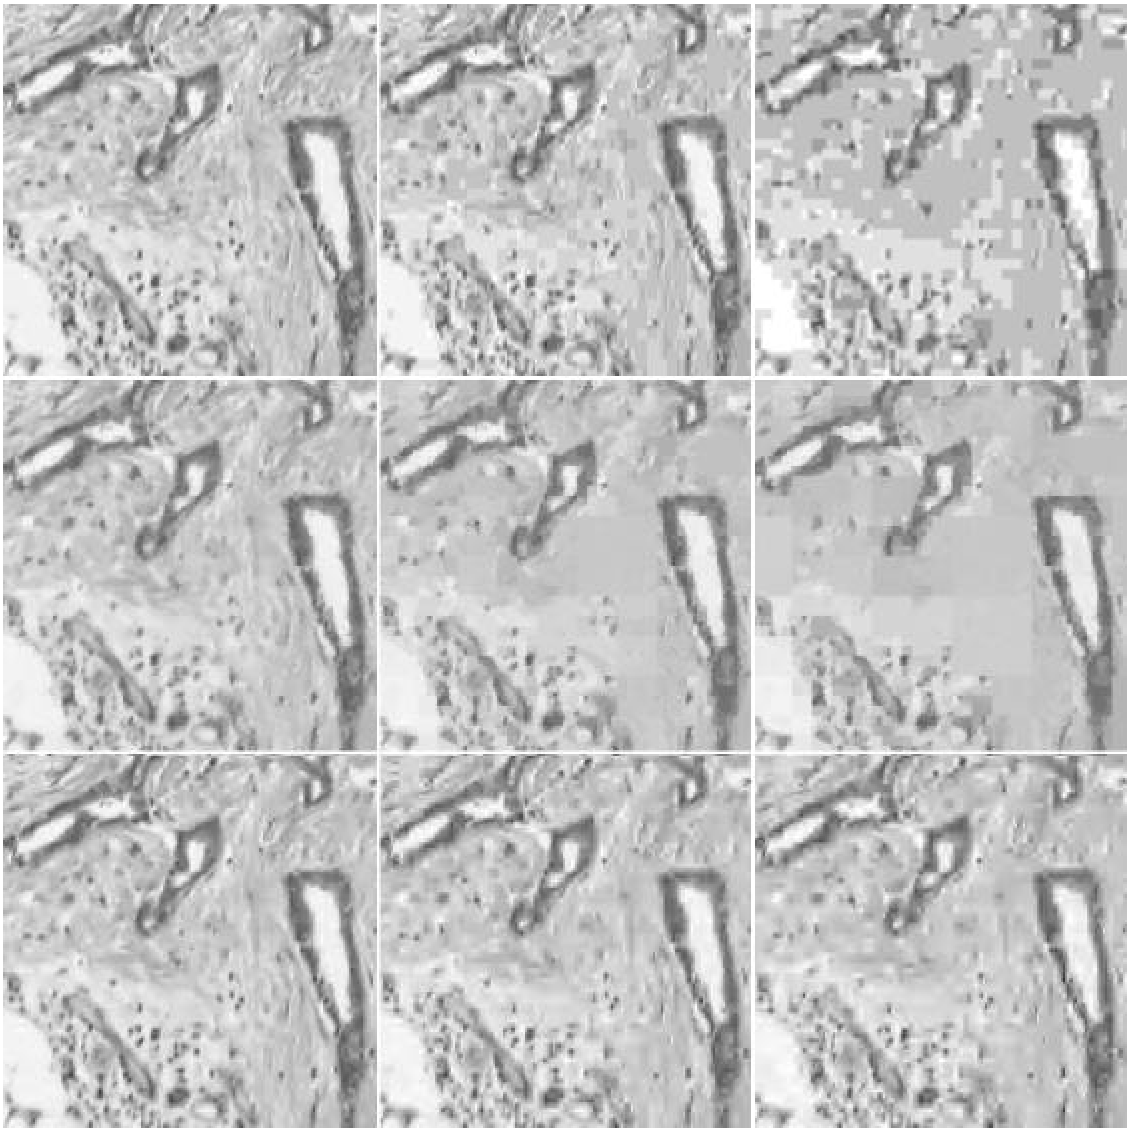

Figure 3, Algorithm comparison. The columns represent the different compression ratios (*left to right*) 20:1, 40:1 and 60:1.

Wavelet Analysis
'Cued' Content Based Image Compression using novel wavelet based approach

A standard wavelet decomposition is used together with the simplest (1-D *and most localised*) Daubechies-4 filter kernel (refer to references [22,23] for the mathematics). The image below-right illustrates this for a 3 level decomposition in which each scale is decomposed into three detail images, namely:

Low-horizontal and High-vertical (LH)
High-horizontal and Low-vertical (HL)
High-horizontal and High-vertical (HH)

Low-horizontal and Low-vertical (LL) is not shown, as this is decomposed into the next level of the pyramid. The final LL image being the subsampled gray level image shown in the top left. The inverse wavelet transform can be used to reconstruct the original image. At this stage no compression has been achieved, we have simply transformed the information in a more suitable form for compression to be applied.

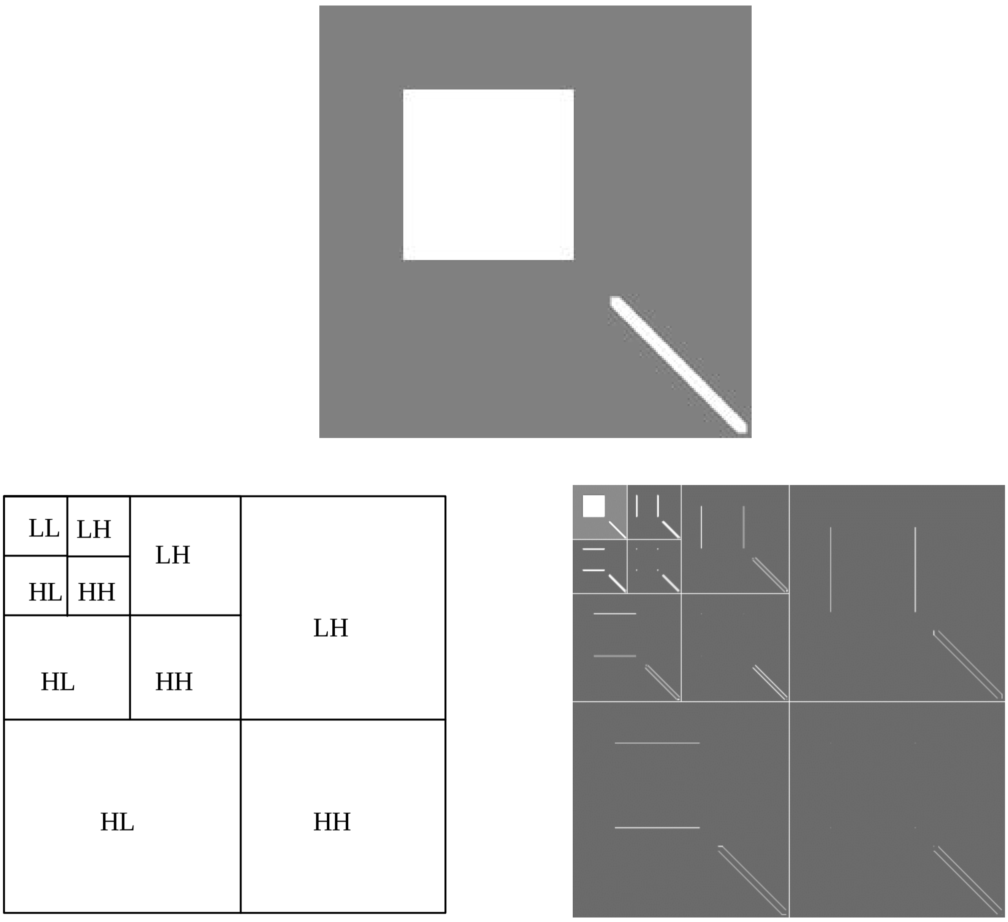

Figure 4: Test image and wavelet decomposition

Intelligent Image Compression

An important and fundamental point behind the wavelet scheme is that once the original image has been transformed into the wavelet domain not all of the data is required. Many wavelet coefficients that are of insignificant magnitude may be discarded without noticeable effects to the human visual system. Figure 5 illustrates that a large percentage of the wavelet coefficients in the wavelet scale space image have been yet still producing a good visual reconstruction. It is apparent therefore that a relatively small percentage of the wavelet coefficients are concentrated in the highly salient areas of the image.

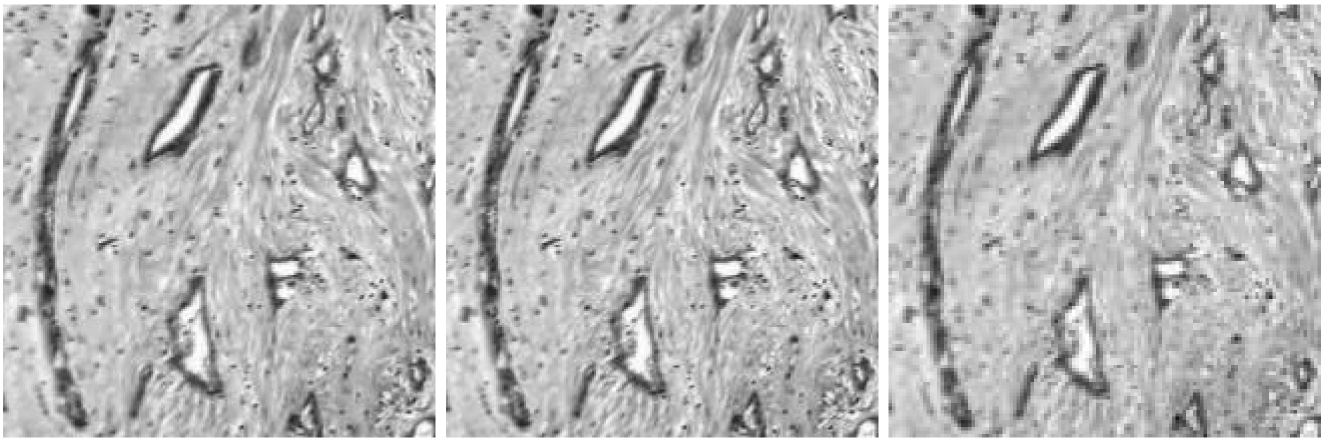

Figure 5: Wavelet reconstruction after removal of 30%, 60% and 95% of weakest coefficients.

Mask Based Wavelet Compression

This fact, that wavelet information is concentrated in important areas is itself the main principle behind the technique used here. The next logical step is to take target masks that have been generated by some saliency operator to prevent removal of coefficients in the diagnostic areas of the image. When this is combined with a suitable encoder an efficient progressive target based compression system can be produced. Currently this cueing process is performed manually with a simple graphical interface but this will in future be replaced with automatic or semi-automatic feature detection being developed within the project. A typical mask[1] is shown in figure 6, here the wavelet coefficients are protected within the two chosen areas and destruction allowed to take place outside, this is noticeable from the image on the right in terms of background and other structure within the image.

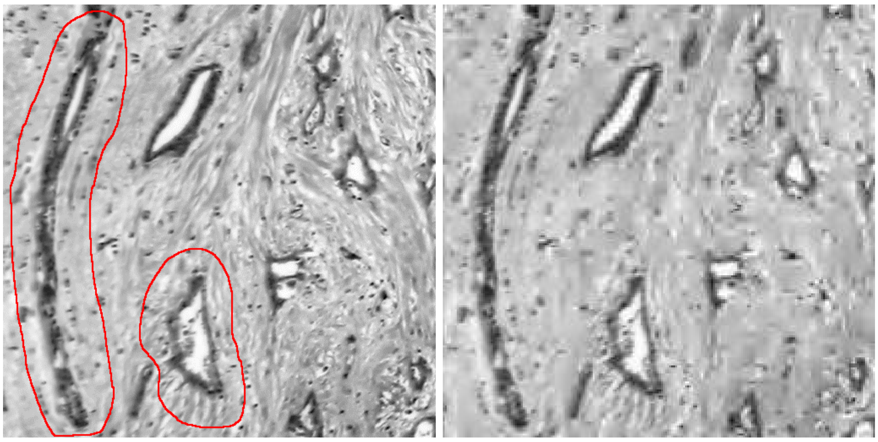

Figure 6: Mask based coefficient removal.

[1] These are somewhat crude masks but they serve to illustrate the technique of intelligent content-based compression.

Progressive Encoder

The SPIHT (Set Partitioning In Hierarchical Trees) algorithm[2] is a wavelet based compression scheme offering good image quality with high PSNR[3]. It is designed for progressive transmission and produces a fully embedded coded file. SPIHT [10,11] is based upon an extension to the embedded zerotree wavelet (EZW) work of Shapiro [12]. It is a coding method so any artifacts produced may be due to the wavelet transform and not the coding process. The process differs from conventional wavelet compression in the way it encodes the wavelet coefficients. It can also be used for lossless compression as it codes the bits of the image wavelet transform coefficients in a bit-plane sequence, thus if all bits are encoded the image can be fully recovered[4]. As with other techniques error protection is important, but using SPIHT's embedded coding scheme means that information is sorted according to importance hence the need for error protection reduces towards the end of the file.

The technique is based upon three ideas:

- A tree based representation of the wavelet coefficients (figure 7),
- Partial magnitude *(so that the most important information may be transmitted first*) ordering of the coefficients
- Ordered bit plane transmission of refinement bits for the coefficients.

As a result of this, both the encoder and decoder know the structure of the data and the algorithm therefore no explicit ordering information needs to be transmitted. Hence if a coefficient at a given level in the hierarchy is insignificant (*relative to some given threshold T*) then those children of that node are also likely to be insignificant. The algorithm can also be efficiently implemented using a few list structures.

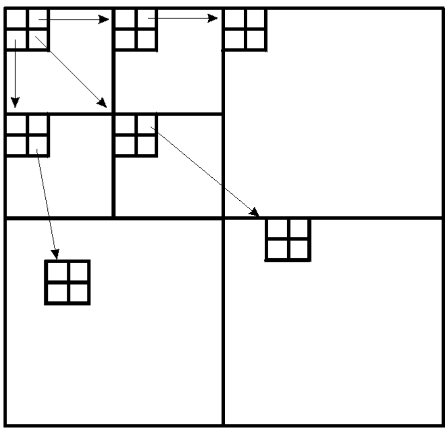

Figure 7: SPIHT - tree structure.

---

2 The implementation of the SPIHT algorithm used here does not contain any arithmetic coding.

3 $psnr = 20\log_{10}\left(2^b - 1/_{rmse}\right)$ where b is the bit depth and *rmse* is the root mean square error.

4 There are issues of precision here, as theoretically lossless compression is only possible if the coefficients are encoded using infinite precision. The remedy to this is to use an integer multiresolution transformation refered to as the S+P transform. This is claimed to be as efficient as most lossless encoders.

Experimental Results

Typical performance of this cued wavelet based compression and encoding approach is illustrated in terms of the signal to noise ratio (PSNR) verses compression factor, figure 8. PSNR is not necessarily the best measure to use for intelligent compression as it does not take into account the fundamental information in the image, however it does provide a useful generic measure for comparison purposes.

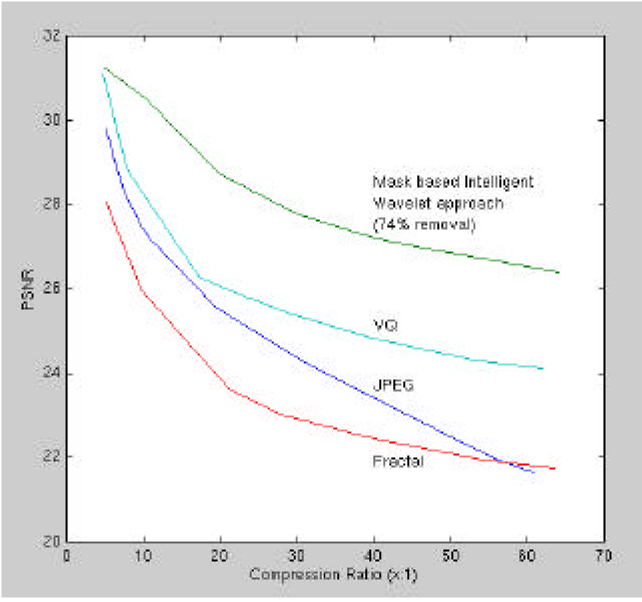

VQ - Vector Quantisation

Figure 8: Typical PSNR verses compression ratio

Figures 9 - 12 show the progressive nature of the algorithm with example cued areas marked. Examples are shown for compression ratios[5] of 150, 100, 100 and 50:1 respectively. At the same compression ratio the preservation of information in the intelligently compressed images is evidential. The cued areas are restored more quickly as the available bandwidth increases. A pathologist could mark the cued regions in figure 9 - 12 or they could be generated automatically. This is used to demonstrate the need of a cueing performance measure as it has an impact on the diagnostic value of the decompressed images.

Performance Evaluation

Performance evaluation is an important requirement and suitable metrication methods must be identified and implemented in the clinical trials. Evaluation is a complex and many-sided issue.

Region Cueing Performance

The performance of the cueing process will be evaluated by the Receiver Operating Characteristic (ROC); here the percentage of true positive detection of diagnostic regions is assessed against the percentage of false negative detection of non-diagnostic regions. The regions of interests will be identified by a number of pathologists and will be used as ground truth.

Compression Performance

The first and most obvious metric for compression performance is ‘compression ratio'; i.e. the reduction in the number of bits used to represent the image. The term compression ratio is useful in relation to lossless compression because it is concerned with the reduction in the number of bits used to

[5] resulting in file sizes of approximately 39, 58, 49 and 97 Kbytes

represent all the data/information. In relation to lossy compression, however, compression ratio conveys no information regarding what information has been preserved and what data/information has been discarded. In general such a metric would need to be related to the relevant criteria for the particular clinical task.

In the case of the developing Content Based Image Compression approach compression ratio is perhaps justifiably used as it is simple and it quantifies the removal of data defined (by content description) as redundant. Thus, the compression achieved is dependent on the overall dimension of all the regions of interests in an image. Imagery with few small regions of interests can be compressed significantly more than imagery with large areas of interests. The achievable compression can be seen to be dependent on the amount of information that must be preserved losslessly in order to ensure value of the transmitted (compressed) image[6].

An algorithmic comparison can be made between the Content Based and conventional approaches using an analysis of the performance of the region cueing processes on images compressed by conventional lossy techniques. Using the annotated imagery as ground truth, the results of the automatic cueing process with and without compression can be compared.

The pathologists involved in the project will assess the 'diagnostic' value of the decompressed images. Their assessment will be used for guiding the development of the Content Based compression approach. A Mean-Square-Error (MSE) or Peak Signal-to-Noise Ratio (PSNR) is used to measure the distortion introduced by the lossy system in the regions defined as informative and the distortion in the non-target areas. The latter is important as the surrounding parts of the image can also play a role by supplying contextual information. The assessment based on the MSE/PSNR and the photographic interpreter/intelligence analyst will be evaluated to determine if there is in fact a corresponding relationship between the objective and subjective measures.

Another important metric is computational load, or the time required to carry out the compression and decompression**.**

Performance Visualisation

It can be seen that there are many important dimensions of performance, which must be taken into account when evaluating compression systems. Reduction of this complex space to single figures of merit quite simply destroys the necessary information. However, understanding and assimilating this complex space is a significant problem for the human and a multi-dimensional graphical representation is necessary. An interactive performance evaluation visualisation tool is therefore being developed to facilitate comparison of the performance of conventional and content-based compression approaches using, among others, the following criteria:

Region cueing in raw/decompressed images; Diagnostic information preservation; Compression ratio; Computation load; Mean-square-error/Peak Signal-to-noise ratio; Pathologist’s assessments

The work on visualisation is thus concerned with the presentation of the performance evaluation measures of the efficiency and effectiveness of different compression approaches at different compression ratios in different circumstances. It will facilitate the assessment and selection of appropriate compression approaches and compression ratios for the image/image type in question. It will, for example, provide a means to understand the degree of invariance of the diagnostic information cueing and preservation to the presence of a compression-decompression step.

[6] It will also depend upon the type of medical imagery to be compressed, eg mitosis imagery such as those in figure 11 and 12 are far more complex and a greater level of detail may be required to satisfy the clinician.

Conclusions

This content-based image compression approach utilises the selective preservation of image features with the inherent benefits of wavelet compression. This has advantage of offering a higher degree of compression compared with applying a compression algorithm uniformly across the whole image. The application of this technique is in telepathology, where images have to be transmitted over narrow bandwidth conditions. Unlike commercial image compression methods, where the aim is using moderate compression in order to preserve image appearance for viewing, the aim of this project is to offer higher compression ratios with preservation of diagnostic data and reduction of redundant data.

It is apparent from studies in telepathology that a diagnosis can be rendered on images that are substandard in comparison to those viewed through microscope eyepieces. The inference is therefore that not all of the data present in a normal microscope view is needed to form an interpretation. This approach takes advantage of this fact and aims to retain salient (diagnostic) features to render an accurate diagnosis.

The region cueing process of the intelligent compression approach will benefit from the quality assurance programme as it will indicate the inter- and intra-pathologist variations in the annotation process. It is also of clinical importance to assess the maximum compression ratio that can be achieved without degradation in the diagnosis. In particular, if we remove a certain amount of data (*e.g. the wavelet coefficients*) then it is more intuitive to remove them from undesirable and less salient parts of an image rather than removing them indiscriminately. Despite the fact that the use of a state-of-the-art codec is being employed it is still advantageous to attempt to remove this non-essential data prior to the encoding process. Typically in experiments carried out, to use a SPIHT codec directly without any masks and removal of coefficients would have resulted in file sizes which would be approximately 3 to 4 times as large.

An issue with medical imagery is that some of it can have highly complex background structure, this structure can lead to problems. After the wavelet analysis has been completed and the spare coefficients have been discarded in the background there can still remain significant wavelet information awaiting encoding and this can compete for the available bandwidth to the detriment of the region of interest. This can be seen on the mitosis image of figure 11 (and to a lesser extent on figure 12). JPEG2000 takes the approach of performing a maxshift operation on the wavelets within the region of interest to ensure that they get encoded first[7], a similar but modified strategy is being developed.

Admittedly, the manual annotation of the image is a labour intensive step. Future work on the project will be to investigate techniques for performing automatic or semi-automatic feature detection, which will then form a cohesive mask that can cue the compression stage. Work will also focus on the development of a performance measure matrix for the clinical trials.

References

1. Ducksbury P.G., 'Target detection and intelligent image compression', SPIE Aerosense 2000, April 24-28, 2000
2. Ducksbury P.G., Varga M.J., 'Feature detection and fusion for intelligent image compression', Submitted to SPIE Aerosense 2001, April , 2001.
3. Samet H., Tamminen M., 'Computing geometric properties of images represented by linear quadtrees', IEEE Trans PAMI, vol. 7, no 2. March 85.

[7] Initial evaluations of the JPEG2000 software achieved less than desirable results as it appeared to corrupt the colour spectrum.

4. Gilles S., ‘Robust description and marching of images’, PhD Thesis, Department of Engineering Science, Oxford University, 1998
5. Pincus S.M., ‘Approximate entropy : statistical properties and applications’, Commun. Statistical Theory Meth., 21:3061-3077, 1992.
6. Mortimer V., ‘Content–based Image Compression’, First year report, Department of Engineering Science, Oxford University, Oct 1999.
7. Mallat S., 'A wavelet tour of signal processing', Academic Press, 1998.
8. Kovesi P., 'Image features from Phase Congruency', Journal of computer vision, MIT Press.
9. Houng-Jyh Wang et al. ‘Multi-Threshold Wavelet Codec (MTWC)’, ISO/IEC JTC1/SC29/WG1 N819, Convergence Phase Meeting, Geneva, March 98.
10. Said A., Pearlman W.A., ‘A new fast and efficient image codec based on set partitioning in hierarchical trees’, IEE Trans on Circuits and Systems for video technology, Vol. 6, June 1996,
11. Manduca A., Said A., ‘Wavelet compression of medical images with set partitioning in hierarchical trees’, SPIE Symposium on Medical Imaging, Cambridge, MA, March 96,
12. Shapiro J.M., ‘Embedded image coding using zerotrees of wavelets coefficients’, IEEE Trans Signal Processing, vol. 41, pp 3445-3462, Dec 93.
13. Boliek M., et al., ‘Next generation image compression and manipulation using CREW’, ICIP-97, Santa Barbara, CA, Oct 97.
14. Rizvi S.A., Nasrabadi N.M., ‘Predictive residual vector quantisation’, IEEE Trans Image Processing, vol. 4, no 11, pp 1482-1495, Nov 95.
15. Oehler K.L., Cossman P.C., Gray R.M. May J., ‘Classification using vector quantisation’, In Proc 25th Annual Asilomar Conf on signals and computers, pages 439-445, Pacific Grove, CA, Nov 91
16. Fisher Y., ‘Fractal Image Compression’, SIGGRAPH 92 course notes.
17. http://www.hpl.hp.com/loco/ - Lossless JPEG download page
18. Preuss D., ‘Comparison of 2-D facsimile coding schemes’, IEEE Press, 1976.
19. ISO Final Draft standard, ‘Lossless and near-lossless coding of continuous tone still images, JPEG-LS’, FCD14495, ISO/IEC JTC1/SC29
20. -,'An analytical study of JPEG 2000 functionalities', ISO/IEC JTC1/SC29/WG1 N1816, July 2000, To be published by Proc SPIE, vol. 4115, 45th annual SPIE meeting, Applications of Digital Image Porcessing XXIII.
21. Christopoulos C., Askelof J., Larson M., 'Efficient methods for encoding regions of interest in the upcoming JPEG2000 still image coding standard', IEEE Signal processing letters, vol. 7, no 9, pp 247-249, September 2000.
22. Daubechies I., ‘Wavelets’, S.I.A.M., Philadelphia, 1992.
23. Clarke R.J., ‘Digital compression of still images and video’, Academic Press, 1995.
24. Bruckmann A., Uhl A., 'Selective medical image compression techniques for telemedical and archiving applications', Computers in Biology and Medicine 30, pp 153-169, 2000.
25. Perlmutter S.M., et al., 'Image quality in lossy compressed digital mammograms', Signal Processing 59, pp 189-210, 1997.

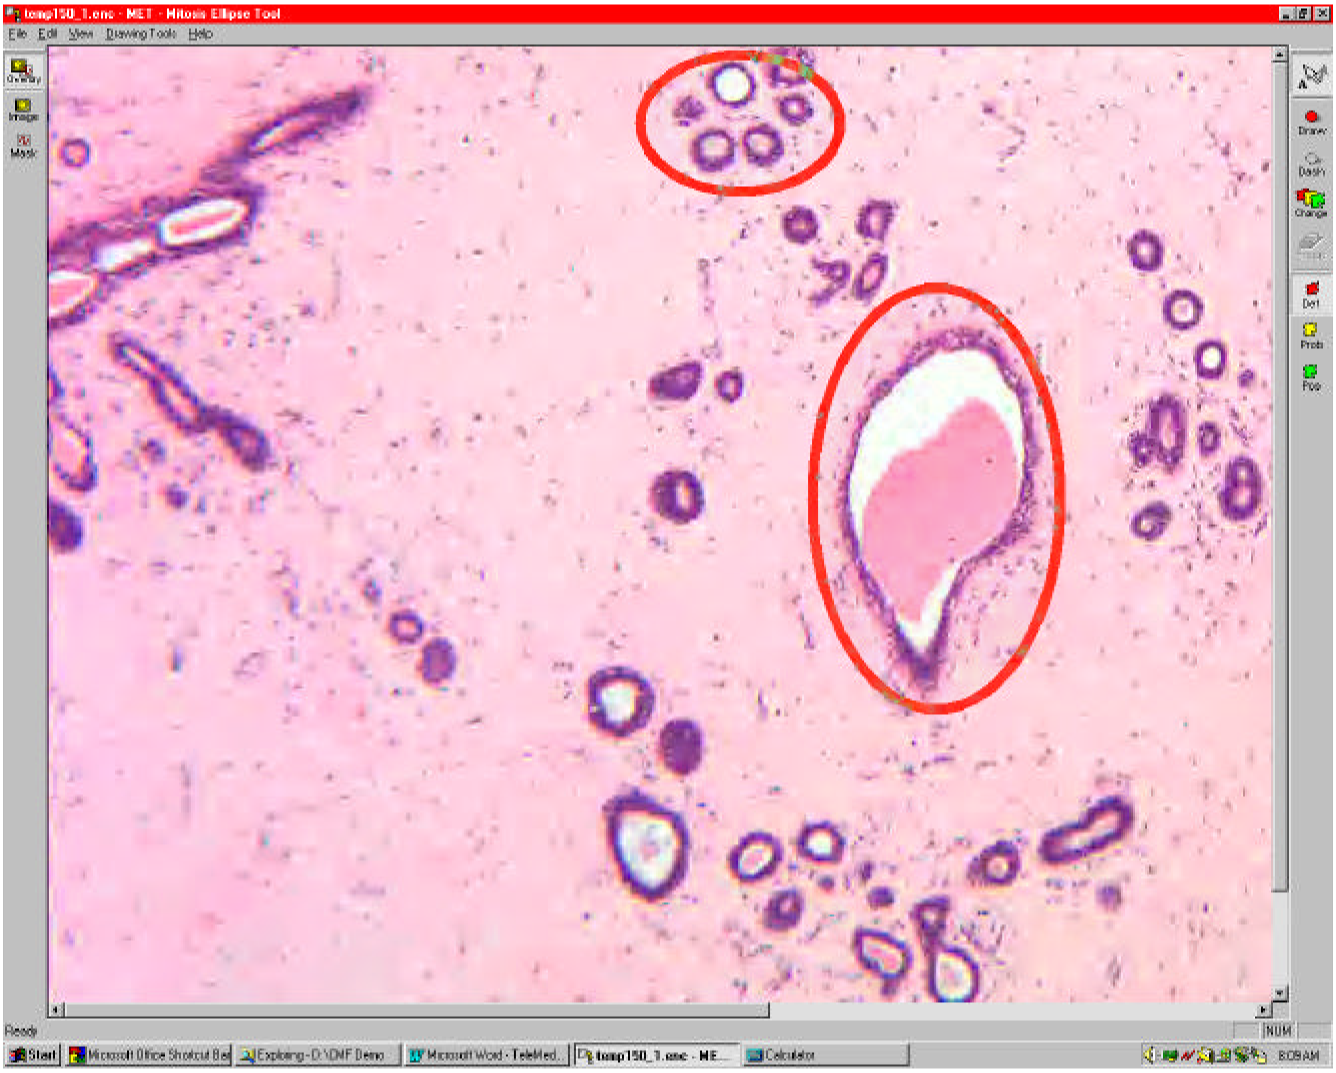

Figure 9 – Compression at approximately 150:1 (95% background removal)

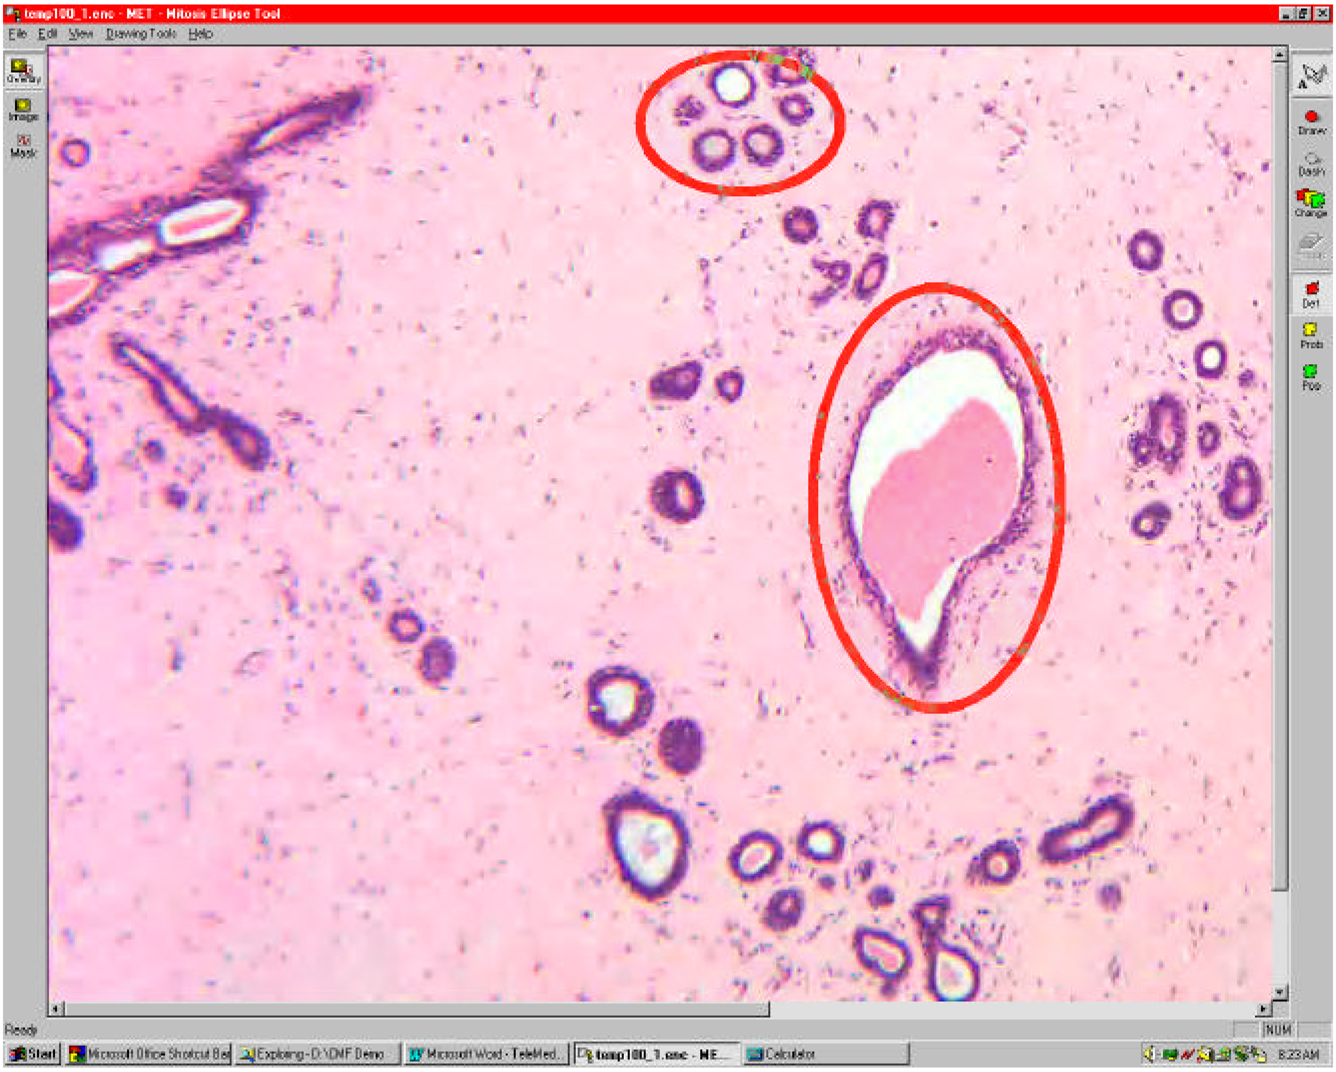

Figure 10 – Compression at approximately 100:1 (95% background removal)

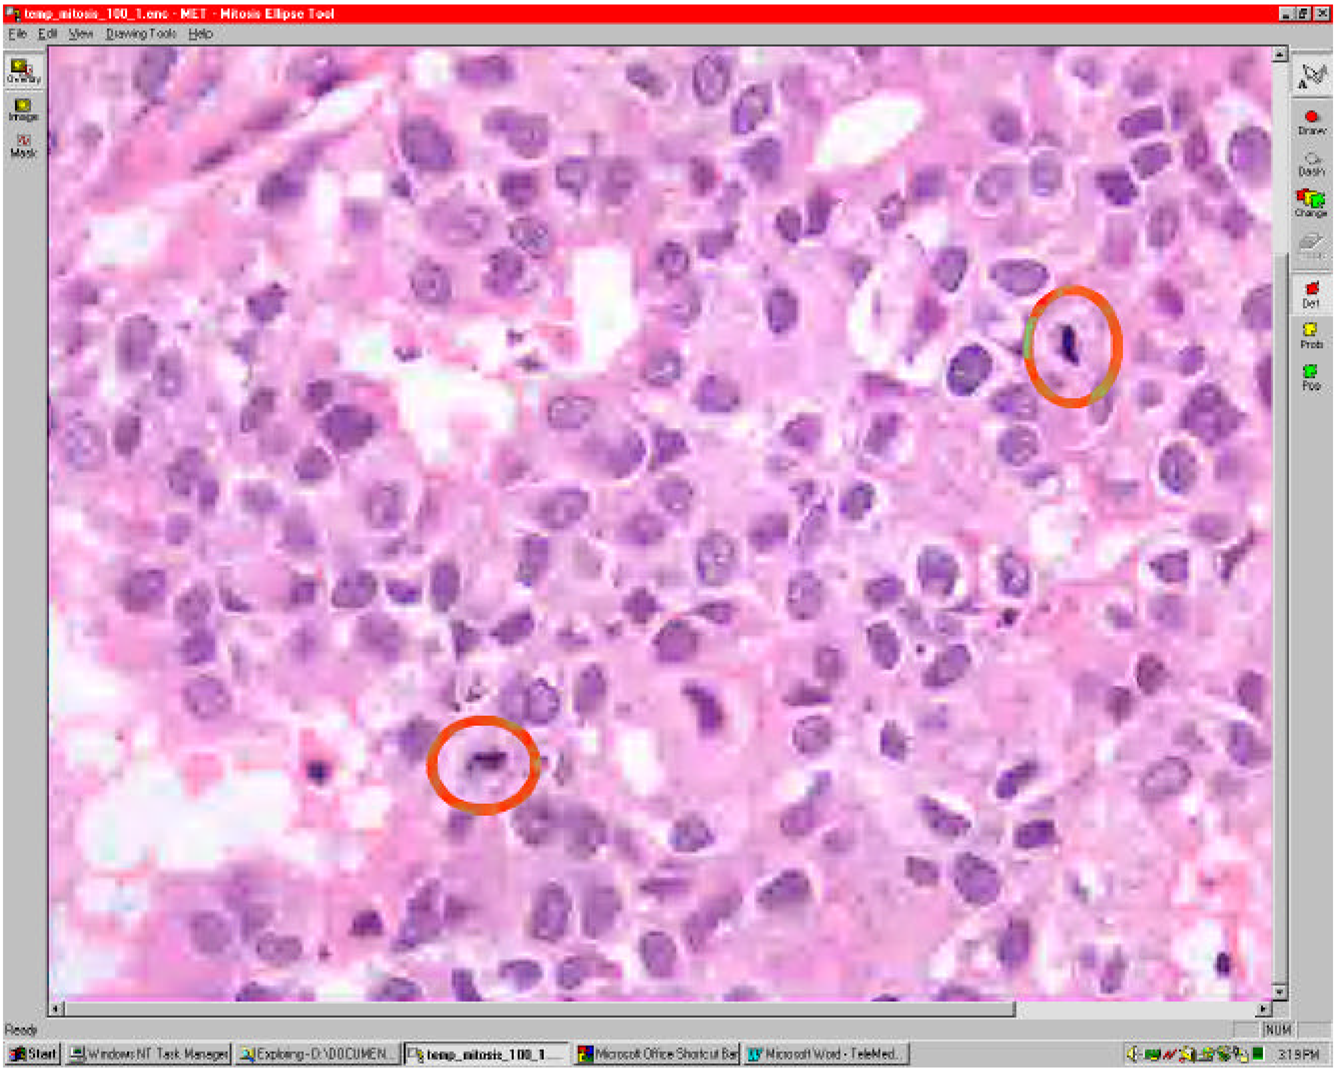

Figure 11 ; Compression at 100:1 Mitosis image (95% background removal)

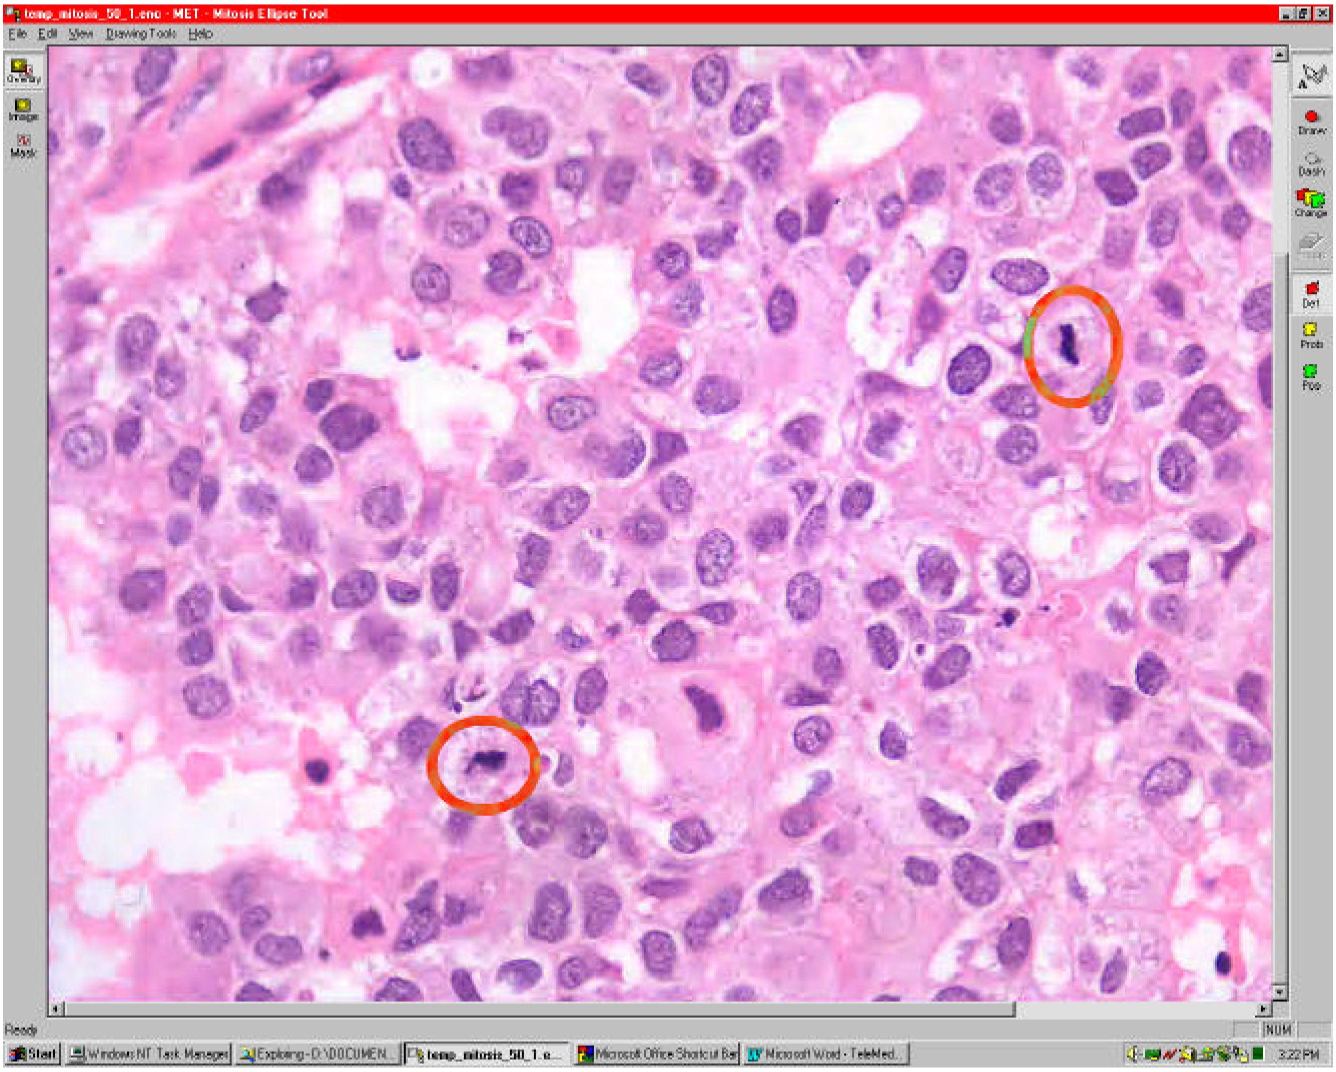

Figure 12 ; Compression at 50:1 Mitosis image (95% background removal)

Based on work presented at ‘National Corrections Telemedicine Conference, 19-21 Nov 2000, Tucson, Arizona’.