Medical image management system

Margaret Varga*, Paul Ducksbury and
Peter Somol

QinetiQ, St. Andrew's Road, Malvern, WR14 3PS, UK
Fax: +44-1684-894384 E-mail: m.varga@signal.qinetiq.com
E-mail: p.ducksbury@signal.qinetiq.com E-mail: somol@utia.cas.cz
*Corresponding author

Abstract: The success of any healthcare service is dependent on the efficient use and sharing of patient information. A Medical Image Database Management System has been designed to store medical images, textural metafiles containing descriptions of clinical features and their diagnostic results, etc. The PC-based system allows users to share and modify data, to retrieve data for further examination, to browse through the database and make clinical decisions based on the medical images found. It also provides the framework within which an annotation tool can be used to annotate areas of clinical concerns graphically and textually. The annotated images can be used for educational purposes or for seeking second opinion confirmation via the telemedicine system, which is integrated within the database system. This paper discusses the use of the system in dermatology to support primary care service.

Keywords: telemedicine; teledermatology; image compression; diagnosis.

Reference to this paper should be made as follows: Varga, M., Ducksbury, P. and Somol, P. (2006) 'Medical image management system', *Int. J. Healthcare Technology and Management*, Vol. 7, No. 5, pp.387–402.

Biographical notes: Margaret Varga is the Technical Manager of the Biomedical Image Analysis Group in QinetiQ. She obtained her BSc (Hons.) degree from Manchester University and her PhD from Cambridge University. Her PhD was in the field of statistical pattern recognition and image processing applied to medical imagery. She joined the Royal Signals and Radar Establishment in 1988, which is now called QinetiQ. Her current research interests are content-based image compression, medical image processing and analysis, Bayesian reasoning, image understanding and analysis, model-based image processing, IR imagery, database systems, information visualisation as well as information management. She has published over 40 papers on computer vision and pattern recognition and contributed chapters in three books on machine vision.

Paul Ducksbury obtained a BSc (Hons.) in Computational Science from the University of Liverpool in 1980. He then spent a year working for Marconi Radar Systems before carrying out research to obtain a PhD involving the study of Numerical Optimisation algorithms for Parallel Processors from The University of Hertfordshire in 1984. Since then, he has worked for SD-Scicon and the Sowerby Research Centre at British Aerospace in the areas of real-time software and mathematical algorithms for transputer-based systems respectively. In 1989 he joined QinetiQ and has carried on the algorithm work which is now concentrated in the area of mathematical modelling with particular applications in image processing and interpretation. His work has resulted in the publication of over 70 reports and papers including a book on

parallel processing and a chapter in the *Handbook of pattern recognition and computer vision.* He is a fellow of the British Computer Society, a Chartered Engineer and a Referee for the *IEE Electronic Letters journal*.

Peter Somol is with the Department of Pattern Recognition at the Institute of Information Theory, Academy of Sciences of the Czech Republic and also with the Medical Informatics Unit, IPH, Cambridge University. He obtained his MSc and PhD degrees at the Faculty of Mathematics and Physics, Charles University, Prague, both from computer science. His current interests include statistical approach to pattern recognition, combinatorial algorithms, graphics, and modern programming.

1 Introduction

The success of any healthcare service is dependent on the efficient use and sharing of patient information. A Medical Image Database Management System has been designed to store medical images, textural metafiles containing descriptions of clinical features and diagnostic results. It was initially designed for managing pathological data (Varga, *et al.*, 2000; Varga, Ducksbury and Callagy, 2002a; 2002b) but its potential use in dermatology is discussed in this paper. The system allows users to share and modify data, to retrieve data for further examination, to browse through the database and make clinical decisions based on the images found. It also provides the framework within which an annotation tool can be used to annotate areas of clinical concerns graphically and textually for educational purposes or for seeking second opinion via the integrated telemedicine system.

It is clear that there can be many transactions initiated by different users who expect to get the most up-to-date records in a reasonable response time. The data can be, for example, dermatology images of a patient captured in a GP surgery by a digital camera.[1] The size of the image varies from 600 × 800 to 3200 × 2400 at 24 bits (Pak, 1999), i.e., upto 23.04 Mbyte per image, and there can be single or multiple images taken during each examination. These images can be a single view of a lesion or lesions and/or images of the lesions taken from different angles as well as the global view of the location and/or spread of the lesions. Images taken from previous examinations can be retrieved and compared with present images to determine progress of the patient's condition, or response to treatment. Furthermore, it is possible to retrieve images of specific conditions for assisting diagnosis or educational purposes.

The design and implementation of the system fulfils the following objectives:

- the system must be accessible by multiple users at the same time
- the system must be scalable enough to cope with the likely number of transactions
- the system should allow future integration of records into electronic patient records.

A SQL database has been designed to meet the above objectives. The database is accessed by a Client/Viewer software system described in the following sections.

1.1 Patient data management

A digital camera is used to capture images. There are important technical issues to consider when taking images, e.g., consistent and good exposure and calibration of size and colour control. Each image in the database belongs logically to a patient taken during an examination and is stored logically within a patient examination record. Each patient examination record contains the following, see Figure 1.

- global view image (if available),
- detail lesion image possibly taken from different angles and its location on the body can be marked on the schematic figure if the global view is not available,
- images may be taken of different parts of the body,
- images may overlap when multiple lesions are involved (not recommended).

A new patient may be set-up and images may be inserted into the database. Patient's demographics, medical history and lesion specific information can be included in the patient data.[2] Images can then be selected either from a textual list or from a patient record preview by a single mouse-click. The manual annotation tool can be called and used from the database system and the diagnosis can be carried out or the telemedicine function used.

Figure 1 Image archive

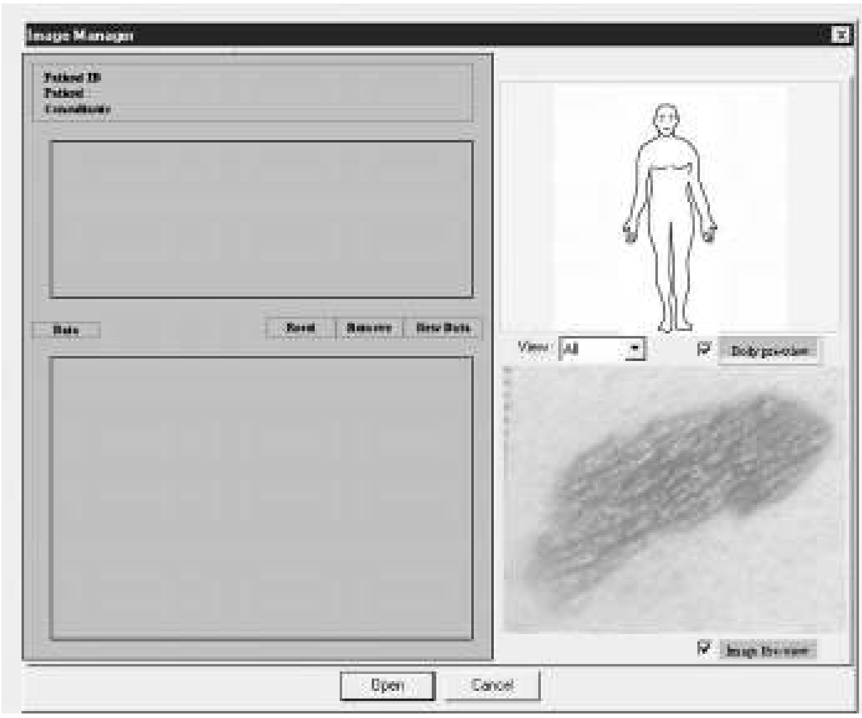

2 Working with images

The interface shown in Figure 2 is used to select from the available functions the required operation to be carried out on an image. The image is displayed in a window on the interface, it can be zoomed in and out, panned, etc. An intuitive and effective tool has been designed in the form of a small button which appears in the right bottom window corner (not shown in Figure 2), clicking the button opens a small preview window displaying the whole image. Moving the cursor with mouse button pressed then drags the view to another image area quickly. Measurement on the lesions is possible if the distance of the camera from the patient is known and calibrated.

Figure 2 Client and user workspace

The user can select to view an image, its corresponding clinical information and the description panel containing detailed diagnostic information and optional clinician's comments. At present it has been designed to have comments from two consultants only, but the number can be increased if deemed necessary. Any of these views can be displayed individually or overlaid on each other (see view-type checkboxes on the left side panel of the main workplace window in Figure 2).

The selected image can be copied by a remote user to a local file, as can the diagnostic data.

2.1 Text search

The system offers a simple natural-language search tool. The tool may be used for navigating the database – to search for records with given names or descriptions, or according to their clinical properties. Queries of the following types are recognised:

- find patient ‘Laura Smith’ – searches for patient names
- find image with ‘Seborhoeic keratoses’ – searches for any images related to ‘Seborhoeic keratoses’
- find record on ‘benign mole’ – searches for any record related to ‘benign mole’.

This interface provides a very intuitive means of retrieving relevant information, see Figure 3.

Figure 3 Natural language search tool

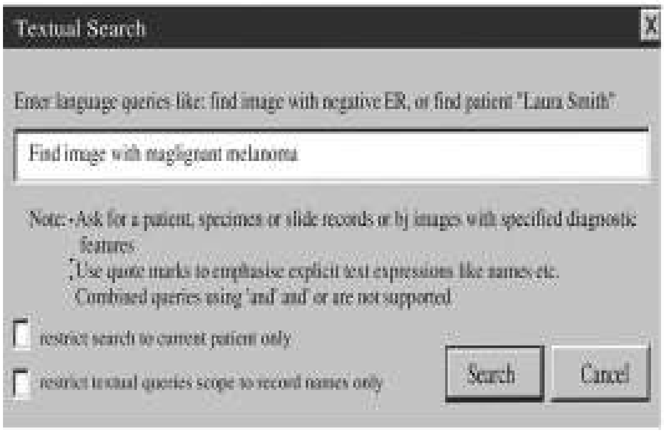

2.2 Annotation tool

Images in their raw form contain vast amounts of data within which a smaller amount of information may be relevant to particular applications, e.g. educational purposes and/or diagnosis. An annotation tool has been developed to provide the facilities for the clinicians graphically to annotate the areas of diagnostic interest or concern by drawing an ellipse around them and entering corresponding descriptive text input if required. The ellipses are in fact image masks, both the graphical and textual annotations are an overlay to the image. Figure 4 shows an annotated melanoma image, the screen is divided into four parts:

- the menu bar across the top embodies all the functions available (some of the most useful features are also available on the toolbars at the sides)
- the view toolbar on the left-hand side allows the user to select what image to view – the mask, the image, or an overlay of the mask(s) on top of the image
- the drawing toolbar on the right hand side allows the user to select the ellipse drawing tool. The ‘A’ button enables the user to attach textual annotation/information to the mask(s) if deemed necessary. There is also an eraser tool
- the main window displays user-loaded images. The user also uses the main window to draw the mask(s) to mark the regions of interests.

The front end for the annotation tool was designed to be simple and intuitive. The number of options has been kept small, so as not to overwhelm or confuse the user. The annotated images can also be used for educational purposes. Figures 5 and 6 show some example images of the form handled within the system.

Figure 4 Graphical and textural annotation

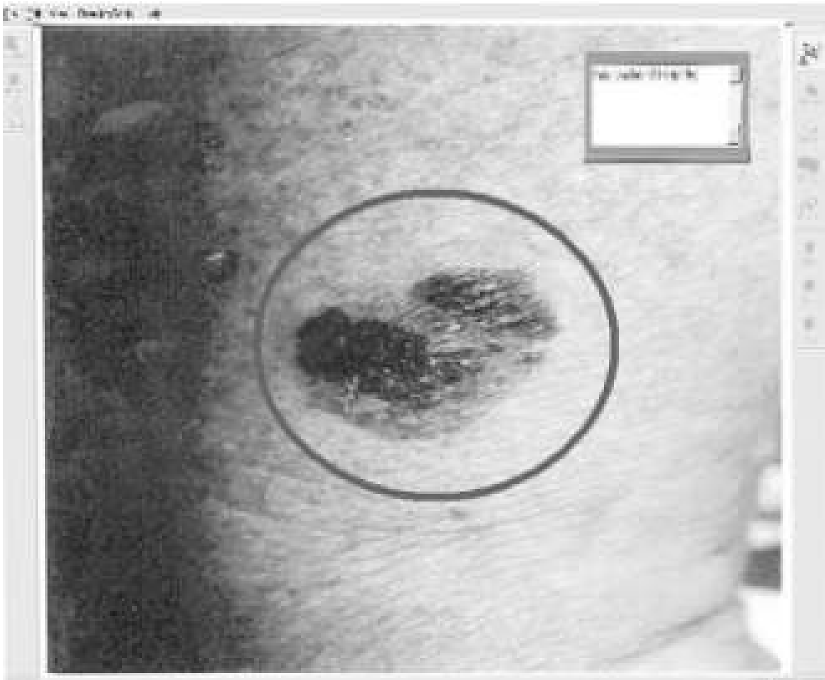

Figure 5 shows a sample annotated region of interest. In this case a Seborrheic keratosis is annotated graphically as a region of interests and can be used for educational purposes (Maize and Snider, 1995). Seborrheic keratoses are the most common benign tumours in older people. They are also known as basal-cell papillomas. They can be found in any part of the body except on the palms, soles and mucous membranes. They are superficial yellowish spots that slowly darken and become warty over the years. Initially they are about 1 cm but can grow to several centimetres or more. They never become malignant but can be troublesome to the sufferers. However, in some cases where patients have many seborrheic keratoses, they may obscure the detection of a dyplastic nevus or malignant melanoma. Sometimes it is difficult to differentiate seborrheic keratoses from lentigo maligna and pigmented actinic keratoses. Furthermore, an inflamed seborrheic keratosis can be confused with a malignant melanoma or a squamous cell carcinoma. The annotation can be used to signify areas of concerns for diagnostic and review purposes.

Figure 6 shows another example of an image; in this case of a possible melanoma for which a second opinion might be required for the confirmation of the diagnosis (Mackie, 1996). Melanoma is a malignant skin cancer, which usually occurs in adults but it may occasionally be found in children and adolescents. It is a more serious type of cancer than the more common of skin cancers. Like most cancers, melanoma is best treated when it is diagnosed early. It can spread quickly to other parts of the body through lymph system or blood. Early warning signs of melanoma are change in the size, shape or colour of a mole, oozing or bleeding from a mole, or an itchy, hard, lumpy, swollen or tender to touch mole. Melanoma can also appear on the body as new mole. Men most often get melanoma on the trunk, head or neck while women often get melanoma on the arms or legs.

Figure 5 Seborrhoeic keratoses

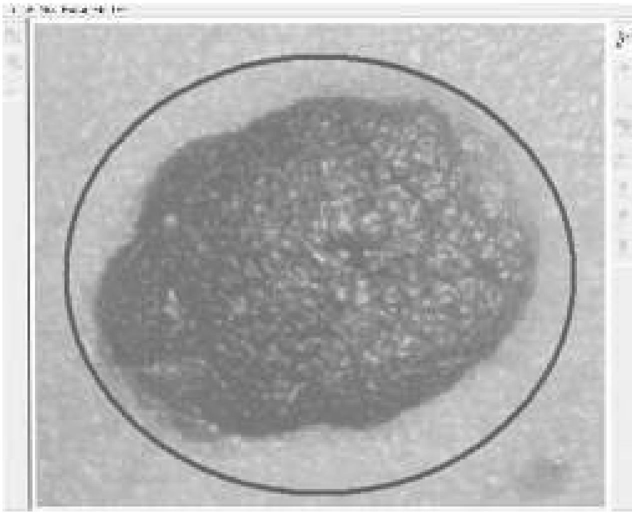

Figure 6 Melanoma image

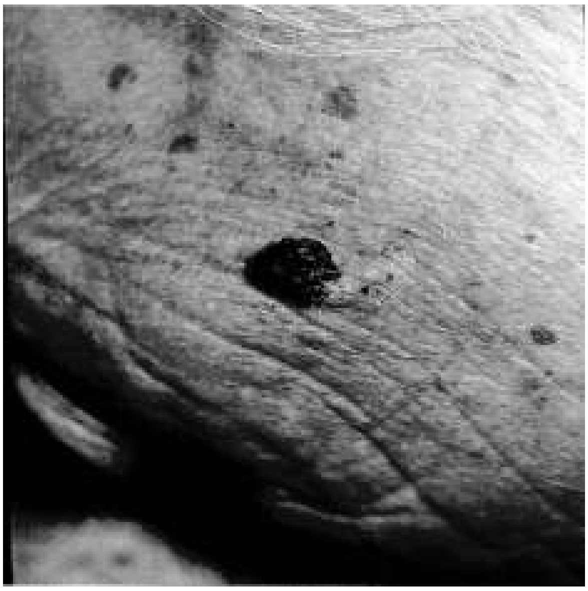

At present face-to-face consultation between a GP and the patient can result in a referral to a consultant. However, a tele-consultation, between the GP and a consultant is possible to manage and diagnose dermatological conditions without the patient going to see the consultant, which means saving patient's and consultant's time. This system provides the requisite support for management and diagnosis of dermatological lesions it allows:

- monitoring
- diagnosis of changes
- provision and management of further specialist referral
- recommendation of treatment or further specialist referral.

Due to the nationwide shortage of specialists, tele-consultancy will become more and more common to address this task.

3 Telemedicine

Telemedicine is a means of practicing medicine at a distance, viewing images on a computer display rather than a face-to-face consultation (Crowe, 1998; Eedy and Wootton, 2001; Perednia and Allen, 1995; Raab *et al.*, 1997). It requires transmitting digital images over networks or telecommunication links. Although the need for telemedicine is well established, the necessity for high-quality image capture devices and efficient transmission has, until now, inhibited the widespread and routine use of telemedicine in some applications, e.g., telepathology. However, the need to transfer images and video across the ever-expanding internet has resulted in more affordable means for capturing and delivering high quality digital images across standard telecommunications links.

The images may be transmitted as real-time video or as still images. The choice is a trade-off between cost considerations and application requirements. Real-time video-conferencing has the advantage of real-time interactivity. However, it has the following disadvantages:

- the need to get the specialists (who may be from different geographical locations in different time zones) together at the same time
- poor imagery quality (which limits its applications)
- difficult record keeping and tracking
- costly.

The alternative choice is a store-and-forward examination of still digital imagery or small video clips (i.e., capture images and records, store them and then review them at a later convenient moment) *c.f.* face-to-face consultation. This is a more practical and cost effective approach, which is commonly used. To transmit the digitised image within a reasonable time across the available bandwidth the image has to be compressed.[3] The compressed images can be sent as e-mail attachment, for example, via the NHS net to an assigned dermatology department or other specialists for tele-consultancy. The consultant dermatologist can then examine the images from the GP or primary care physician, make a diagnosis and recommend treatment. In general 3–5 images are usually adequate for diagnostic purposes. Indeed teledermatology holds great potential for revolutionising the delivery of dermatology service by allowing primary care physicians to refer patients to centres of excellence at a distance (Eedy and Wootton, 2001; Harrison, *et al.*, 1998; Lim, *et al.*, 2001; Oakley, Duffill and Reeve, 1998). However, before its routine application as a service is realised its reliability, accuracy and cost-effectiveness need to be verified by vigorous assessment. It is believed that poor image quality is one of the major concerns to the widespread use of teledermatology, colour and sharpness play a significant part in the confidence and accuracy of the diagnosis (Krupinski, 1999; Loane *et al.*, 1997). While there is a variable range of reported success, overall there are indications that teledermatology will become increasingly used and incorporated in traditional dermatology service. Studies have also been conducted regarding patient satisfaction with teledermatology. However, to date there are no conclusive findings regarding the underlying reasoning for patient satisfaction or lack of satisfaction (Loane, *et al.*, 1998; Mair, and Whitten, 2000). It is important to understand the patient's response to the use of real-time video conferencing and store-and-forward consultation. In this paper only the store-and-forward of still images is considered.

4 Image compression and decompression

4.1 Backgound

A review of the current state-of-the-art of image compression techniques and systems was carried out. It found that, at present, in order to achieve the necessary high levels of compression so called 'Lossy' compression must be used and not 'Lossless' compression. In lossy compression information is lost and there is degradation in image quality, whereas in lossless compression no information is lost and there is no degradation in image quality (Clarke, 1995; Mallat, 1998; Manduca and Said, 1996). However, the review also showed that current lossy coders, with the exception of the QinetiQ Content Based Image Compression (QCBIC)[4] (Ducksbury, 2000; Varga *et al.*, 2000; Varga, Ducksbury and Callagy, 2002a; 2002b) and the more recent JPEG2000 Standard, achieve their compression ratio by compressing indiscriminately over the entire image, which results in the loss or corruption of clinically important information. None of the currently available lossy compression techniques addresses the issues of what information should be preserved and what data can be discarded, i.e., their objective is to compress as much as possible and they are not concerned with the usability of the decompressed images. In many applications the fact that there is significant corruption caused by the compression is often tolerated, particularly as there is no alternative: however, when detailed information is required this corruption makes clinical diagnosis difficult and at times impossible.

4.2 Compression overview

Still image compression techniques exploit the massive redundancies in images as well as the non-linearities in the human vision system. Compression is achieved by removing redundant information (2-dimensional correlations) within image segments. The redundancy is the predictability, randomness or smoothness of the data; for instance, the content of a pure noise image is unpredictable and it is therefore necessary to store every pixel for reproduction, i.e., it cannot be compressed. During the compression process the original image is mapped into fewer/smaller target representations: while in the decompression process the target representation is converted back into the original image or an image similar to the original. Broadly, image compression can be split into two main categories of techniques, namely lossy and lossless. Lossy compression, where the reconstructed image loses information with respect to the original image, can achieve a significantly higher compression ratio (10:1 to 30:1) than lossless compression (2:1 on average) where the original image is preserved. In some applications images reconstructed following a 10:1–30:1 compression are acceptable, i.e., where detailed information is not required and only the overall impression is important. However, accurate diagnosis is not guaranteed if it is based on degraded images. It must thus be realised that there is a limit on how much an image can be compressed and this limit is dependent on the application.

4.3 JPEG2000 region of interest compression

This telemedicine application uses the capabilities provided within one of the latest image compression Standards, namely JPEG2000.[5] JPEG2000 provides a very extensive suite of functions, which can be built into an application.

JPEG2000 can be used to compress uniformly across an image (so-called global compression) or selectively within an image (so-called region of interest compression). In region of interest (ROI) compression different degrees of compression can be applied selectively, thus regions not of interest can be compressed heavily, causing quality degradation or corruption; while regions of interest are maintained. Overall a high degree of compression can be achieved preserving the important information at the relative expense of the non/less important parts of the image.

Figure 7 shows a schematic of the ROI compression system and evaluation framework. The compression system is composed of three components:

1 Annotation to create a ROI 'mask' or 'masks' to direct the compression

2 Compression

3 De-compression

Figure 7 Content-based image compression framework

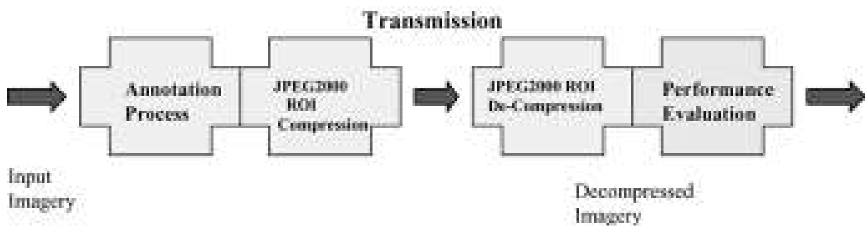

4.4 Performance evaluation

Various images with different characteristics were used in the evaluation work to establish the robustness of the ROI compression approach. Figure 8 shows the global compression of the image in Figure 6 at 50:1 and 100:1. It can be seen that there is a general smoothing effect across whole image which includes the boundary of the lesion, i.e., it looses its sharpness. The clarity of the lesion outline plays an important role in the diagnostic process. Figure 9 shows the ROI compression at the same compression ratios, it can be seen that the sharpness of the lesion outline is well preserved while the background data is smoothed. Figures 10 and 11 provide a further example. Figure 10 shows another melanoma image, globally compressed images at 50:1 and 100:1, the general smoothing effect is again evident. Figure 11 shows the ROI compressed version at 50:1 and 100:1. The benefits of using the ROI compression approach as compared with the global approach can be observed clearly.

Figure 8 Original raw image and global compression at 50:1 and 100:1 compression ratio

Figure 9 ROI compression at 50:1 and 100:1 compression ratio

The typical requirement for diagnosis is five images; each raw image is typically 5–10MB and if these were to be sent using a typical broadband connection, for example 256kbit/sec ADSL, they will take about 13–26 minutes to transmit (larger images with higher resolution could take up to 2 hours to transmit). If, however, the images are compressed by a ratio of 50:1 or 100:1, as described above, then it will take about 14–36 seconds or 7–18 seconds to transmit all the images respectively.

5 Conclusion

The database system enables the remote and interactive management (annotation, measurement, indexation, storage and retrieval) and diagnosis of patient's medical images. It provides the capability of generating patient reports and easy access to patient's image data. The system also offers a natural-language processing search tool for navigating the database. The system is useful for diagnosis, telemedicine and teaching.

The JPEG2000 ROI image compression approach used in the telemedicine component utilises the selective preservation of diagnostic information. This offers a higher degree of usable compression compared with applying one of the existing Standard compression techniques uniformly across the whole image. It is apparent from observations that not all of the data present in an image are of diagnostic value. The ROI compression approach takes advantage of this fact and aims at retaining salient diagnostic features to render an accurate diagnosis. The ROI compression technique can be applied effectively in teledermatology for second opinion consultation, where images are transmitted as e-mail attachments. Initial experiments have shown that a compression ratio of between 50–100:1 is achievable without degradation of the diagnostic content of the images. The precise gain in performance depends on the application, the diagnostic content of the image as well as the level of expertise. When compared with lossless compression it can provide the same degree of information with a smaller amount of data. When compared with lossy compression it can provide more information for a given amount of data.

Figure 10 Raw melanoma Image, global compression at 50:1 and 100:1

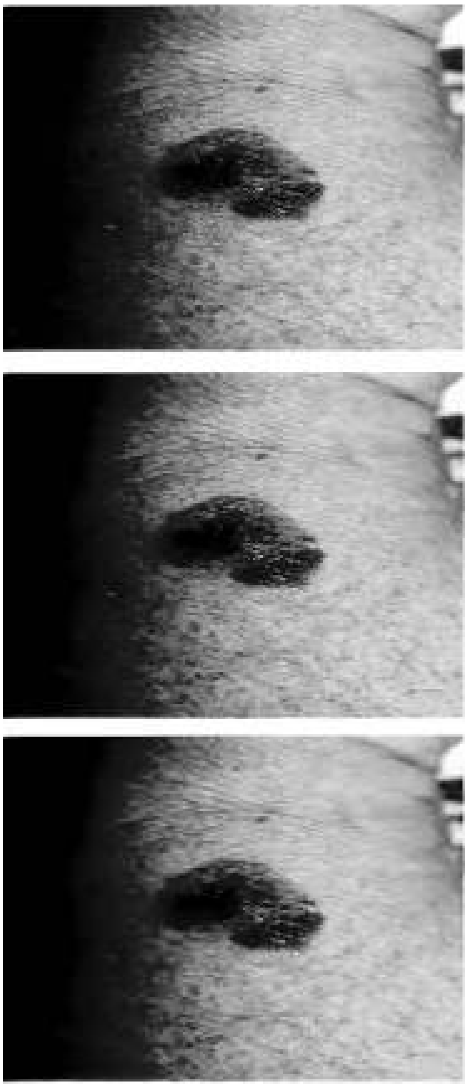

Figure 11 ROI compression at 50:1 and 100:1 compression ratio

It can be seen that at 100:1 compression ratio the areas of clinical interest are preserved while the non-annotated areas are compressed to reduce the size and transmission time. However, even with this degradation the background areas still provide useful contextual information to support diagnosis. The precise gain in compression performance depends on the application, the content of the image as well as the level of expertise.

Acknowledgement

Part of this work was undertaken as part of a MoD sponsored Corporate Research Programmes.

References

Clarke, R.J. (1995) *Digital Compression of Still Images and Video*, Academic Press.

Crowe, B.L. (1998) 'Cost-effectiveness analysis of telemedicine', *J. Telemed Telecare*, Vol. 4, No. 1, pp.14–17.

Ducksbury, P.G. (2000) 'Target detection and intelligent image compression', *SPIE Aerosense 2000, Proceedings of SPIE*, Orlando, 24–28 April, Vol. 4050.

Eedy, D.J. and Wootton, R. (2001) 'Teledermatology: a review', *Br.J. Dermatology*, April, Vol. 144, No. 4, pp.696–707.

Harrison, P.V., Kirby, B., Dickinson Y. and Schofield, R. (1998) 'Teledermatology – high technology or not?', *J. Telemed Telecare*, Vol. 4, No. 1, pp.31–32.

Krupinski, E. (1999) 'Diagnostic accuracy and image quality using a digital camera for teledermatology', *Telemedicine Journal*, Vol. 5, No. 3, pp.257–263.

Lim, A.C., Egerton, I.B., See, A. and Shumack, S.P. (2001) 'Accuracy and reliability of store-and-forward teledermatology: primary results from the teledermatology project', *Australia J. Dermatology*, Nov., Vol. 42, No. 4, pp.246–251.

Loane M.A., Bloomer, S.E., Coprbett, R., Eedy, D.J., Gore, H.E., Mathews, C., Steele, K. and Wootton, R. (1998) 'Patient satisfaction with realtime teledermatology in northern Ireland', *J. Telemed Telecare*, Vol. 4, No. 1, pp.41–46.

Loane, M.A., Gore, H.E., Corbett, R., Steele, K., Mathews, C., Bloomer, S.E., Eedy, D.J., Telford, R.W. and Wooton, R. (1997) 'Effect of camera performance on diagnostic accuracy, preliminary results from northern ireland arms of the UK multicentre teledermatology trial', *J. Telemed Telecare*, Vol. 3, No. 2, pp.83–88.

Mackie R.M. (1996) 'An illustrated guide to the aetiology, clinical features, pathology and management of benign and malignant cutaneous tumours', in Martin Dunitz, *Skin Cancer: Second Edition*.

Mair, F. and Whitten, P. (2000) 'Systematic review of studies of patient satisfaction with telemedicine', *BMJ*, Vol. 320, No. 7248, pp.1517–1520.

Maize, J.C. and Snider, R.L. (1995) 'Nonmelanoma skin cancers in associations with seborrheic keratosies, clinicopathologic correlations', *Dermatol Surg*, November, Vol. 21, No. 11, pp.960–1002.

Mallat, S. (1998) *A Wavelet Tour of Signal Processing*, Academic Press, 1998.

Manduca, A. and Said, A. (1996) 'Wavelet compression of medical images with set partitioning in hierarchical trees', *SPIE Symposium on Medical Imaging*, March, Cambridge, MA.

Oakley, A.M., Duffill, M.B. and Reeve, P. (1998) 'Practising dermatology via telemedicine', *N.Z. Med J.*, Vol. 111, No. 1071, pp.296–309.

Pak, H. and Maj., S. (1999) *Basic Guide to Dermatologic Photography, Dermatology*, Walter Reed Army Medical Center, Washington, DC 20307.

Perednia, D.A. and Allen A. (1995) 'Telemedicine technology and clinical applications', *JAMA*, Vol. 273, pp.483–488.

Raab, S.S., Robinson, R.A., Subder, T.E., McDaniel, H.L., Sigman, J.D., Leigh, C.J. and Thomas, P.A. (1997) 'Telepathologic review: utility, diagnosis accuracy, and inter-observer variability on a difficult case consultation service', *Modern Pathology*, June, Vol. 10, No. 6, pp.630–705.

Varga M.J. Ducksbury P.G. and Callagy, G. (2002) 'Performance visualisation for image compression in telepathology', *SPIE Medical Imaging Conference*, February, San Diego.

Varga, M.J., Ducksbury, P.G. and Callagy, G. (2002) 'Application of content-based image compression to telepathology', *SPIE Medical Imaging Conference*, February, San Diego.

Varga, M.J., Ducksbury, P.G., Leong, F.J. and McGee, J.M. (2000) 'The application of intelligent compression to telepathology', *2nd National Corrections Telemedicine Conference*, 19–21 November, Tucson, Arizona.

Notes

1 A variety of digital cameras that are suitable for teledermatology are available from Canon, Fuji, Nikon, Olympus, Minola etc.

2 Data has been made available to the work by Stuart Andrews from his PhD's dataset at Lancaster University.

3 Image compression is the process of reducing the amount of data (number of bytes) used to represent the image.

4 The QCBIC technique is a simpler approach than the JPEG2000 standard. It pre-dates JPEG2000 and is easier to apply however, for the purposes of this paper JPEG2000 is considered to illustrate the generality of the findings.

5 JPEG2000, http://www.jpeg.org/jpeg2000