Paper

Gaze-J2K: gaze-influenced image coding using eye trackers and JPEG 2000

Anthony Nguyen, Vinod Chandran, and Sridha Sridharan

Abstract— The use of visual content in applications of the digital computer has increased dramatically with the advent of the Internet and world wide web. Image coding standards such as JPEG 2000 have been developed to provide scalable and progressive compression of imagery. Advances in image and video analysis are also making human-computer interaction multi-modal rather than through the use of a keyboard or mouse. An eye tracker is an example input device that can be used by an application that displays visual content to adapt to the viewer. Many features are required of the format to facilitate this adaptation, and some are already part of image coding standards such as JPEG 2000. This paper presents a system incorporating the use of eye tracking and JPEG 2000, called Gaze-J2K, to allow a customised encoding of an image by using a user's gaze pattern. The gaze pattern is used to automatically determine and assign importance to fixated regions in an image, and subsequently constrain the encoding of the image to these regions.

Keywords— eye tracking, image compression, importance map, JPEG 2000, region of interest.

1. Introduction

The use of visual content in applications of the digital computer, such as the Internet and world wide web, has increased dramatically in recent years. Image compression standards such as JPEG 2000 [1, 2] have been developed to provide scalable and progressive compression, and thus images can be displayed with varying resolution and quality depending on the bandwidth, memory and time available. Applications such as electronic commerce have become a reality, allowing merchandise in e-commerce to be displayed as images.

Advances in image and video analysis are also making human-computer interaction multi-modal, rather than through the use of a keyboard or mouse. New sensors and input devices such as the eye tracker have emerged. An eye tracker can locate on the monitor screen where a user is looking. Recent advancements in eye tracking technology, specifically the availability of cheaper, faster, accurate and user-friendly trackers, have inspired new research into eye movements and gaze patterns. Eye trackers are no longer intrusive or require cumbersome headgear to be worn.

Eye tracking can be used by an application that requires the display and adaptation of visual content to the viewer, provided the format in which the image is represented (coded) and reconstructed (decoded and displayed), and the environment (operating system extensions) in which the format is utilised can support such adaptation. Many features are required of the format to facilitate this and some are already a part of image coding standards such as JPEG 2000. Images can be encoded and decoded in JPEG 2000 with scalable resolution and quality in a progressively increasing manner, and regions of interest (ROI) can be selected in images and used to encode/decode the image in a non-uniform manner. This paper presents a system, called Gaze-J2K, which uses eye trackers and JPEG 2000 to allow an image author (i.e., user at the encoder) to use their gaze to automatically determine and assign importance to fixated regions in an image, and subsequently constrain the encoding of an image to these regions. This allows the user to receive the image as desired by the image author. This is very useful for the Internet and world wide web for applications such as merchandising, where faster interpretation of image contents would imply the faster rejection of unwanted images and hence improve user productivity.

2. Gaze-J2K system overview

The Gaze-J2K system incorporates a multi-modal interaction device provided by an eye tracker to allow an image author to influence and direct the encoding of an image to particular objects or ROIs in an image using JPEG 2000. The system comprises of three stages of operation as shown in Fig. 1, namely, gaze point collection, gaze point analysis and ROI prioritised JPEG 2000 image coding. Here the gaze point collection stage records information on the location and sequence of regions in an image followed by the image author. The structure of the spatial and temporal characteristics of the gaze pattern can then be used as parameters and analysed to generate ROIs and its measure of importance. These ROIs and importance scores define a ROI "importance" map, which can be input to a JPEG 2000 ROI encoder to prioritise the image code-stream according to the importance map specification. A client at the decoding end can receive the image progressively with the default ROIs, which were considered important by the image author, reconstructing faster than other regions in the image. Each stage of operation is described in further detail in the following sections.

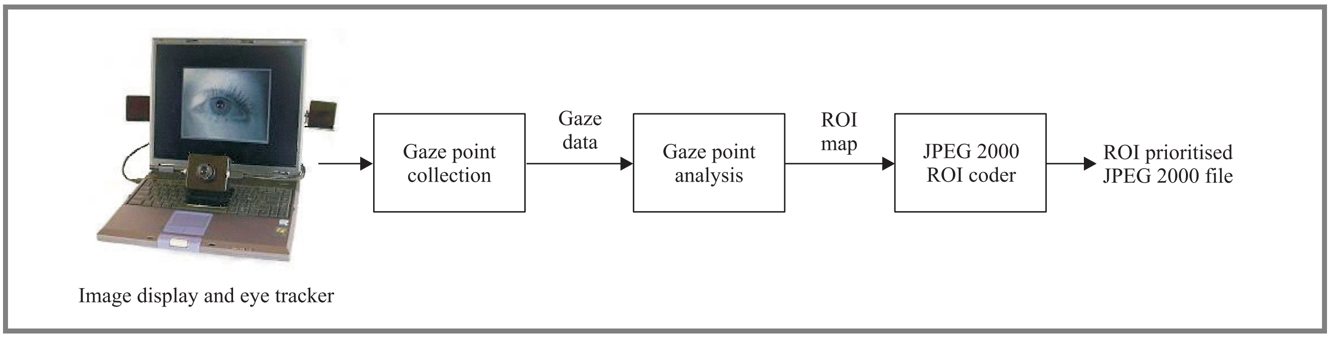

Fig. 1. Gaze-J2K system block diagram.

3. Gaze point collection

The goal of gaze or eye tracking is to determine the gaze point on the field of view where a user is looking. The device used to record eye movements was an EyeTech video-based corneal reflection eye tracker [3]. Infrared lights were mounted on both sides of a computer monitor to illuminate the eye and provide reference points for the eye tracker. The method of operation relies on focusing and tracking the infrared reflections from a user's

Number of recorded samples

1 second of gaze data

150

0	551	473	-10375	0
1	537	464	70	10445
2	514	297	100	30
3	514	301	150	50
4	509	301	220	70
5	504	280	290	70
6	505	280	350	60
7	505	286	421	71
8	510	288	491	70
9	511	288	551	60
10	511	366	621	70
11	502	384	691	70
12	503	384	751	60
13	503	379	821	70
14	468	383	891	70
15	469	383	951	60
16	480	383	1021	70
17	480	383	1091	70
18	473	372	1152	61
19	492	342	1222	70
20	579	342	1292	70

Fig. 2. Example eye tracker gaze data output (extract).

eye using an image sensing camera mounted in front of the monitor screen. By analysing the position of the infrared light reflections and the center of the pupil contained in the image captured, the gaze point can be determined. The gaze-tracker can operate at 15 to 30 frames per second (fps) and records the position and time of gaze. The system was setup so that gaze data collection can be conducted on an image and screen resolution of 1024×768 pixels.

An example extract of a recorded gaze data output by the eye tracker at 15 fps is shown in Fig. 2. The first line of the recorded gaze data shows how many samples were recorded. Each line thereafter contains information for one gaze point, which details the sample number, x-position (pixels), y-position (pixels), time from start (ms), and time from last sample (ms). The position and time information provided by the gaze data provides a couple of parameters that can be used to analyse the gaze pattern and determine, if any, ROIs fixated by the viewer.

A collection of gaze data were obtained from 13 subjects, of which 11 were naive to the purpose of the study. Six colour images (boat, cow, horse, paddock, rockclimb, and yacht), each with at least one primary object of interest clustered in a scenic background as shown in Fig. 3 (in this edition black-white), were displayed on the computer monitor, and each subject was told to locate and examine the objects in the image. For each image, the task was repeated three times for a duration of 15 seconds each. The eye tracker sample rate was set to 20 fps.

A few problems can occur at this stage of the Gaze-J2K process due to the eye tracker failing to track the infrared reflections or the pupil of the eye. As a result, the number of gaze points recorded by the eye tracker can be considerably less than normal. Drifts in the location of the gaze points cause by the tracking of the infrared reflections can also correspondingly produce an offset relative to the actual location fixated. If this offset is large enough, then subsequent stages of Gaze-J2K will generate a ROI map that will not correspond to the ROI. Minor drifts, however, will not adversely affect the results.

A culling process was performed to discard gaze data sets that were found to be unsuitable for further processing. A total of 229 out of the 234 gaze patterns were retained for the testing of subsequent stages of Gaze-J2K. These discarded gaze data sets were mainly due to the eye tracker failing to locate the viewer's location of fixation for the vast majority of the viewing duration.

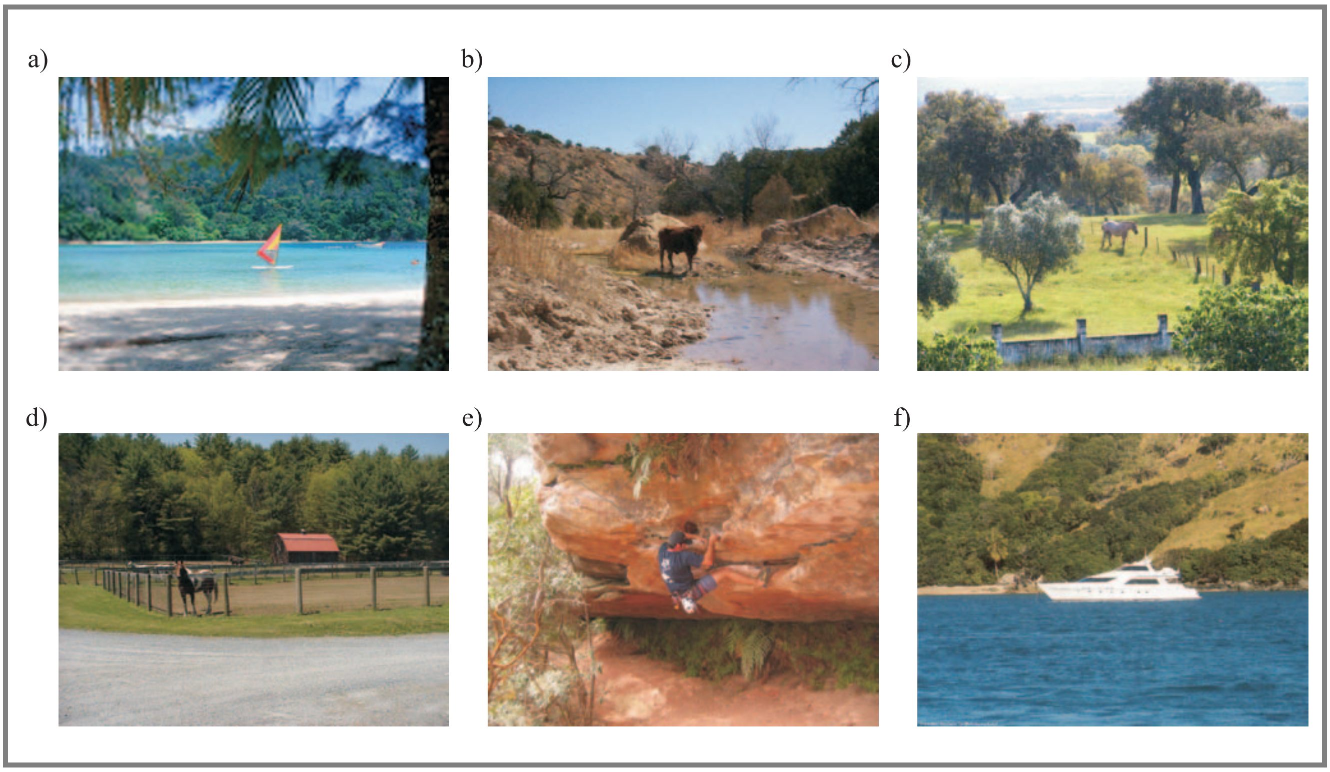

Fig. 3. Gaze-J2K colour test images (1024 × 768, 24 bpp): (a) "beach"; (b) "cow"; (c) "horse"; (d) "paddock"; (e) "rockclimb"; (f) "yacht".

4. Gaze point analysis

The purpose of gaze point analysis is to analyse a viewer's gaze pattern to determine ROIs fixated by the viewer. This procedure reduces the spatial characteristics of the gaze pattern to a limited subset of clusters that would represent ROI candidates. The choice of clustering technique is influenced by a number of factors such as whether the probability densities of the data are known or can be modelled, and the size of the data set. Since the number of gaze location points are limited and its spatial distribution is unknown, an unsupervised clustering technique, such as K-means, was used. In addition to the clustering procedure, a means to determine the importance of the ROI candidates was also investigated. The following subsections detail the development of the ROI clustering and ROI mapping stages.

4.1. ROI clustering

The ROI clustering involves the partitioning of gaze points into mutually exclusive clusters such that the loci of the points belonging to the clusters represent ROI candidates for the particular gaze pattern. Here, a K-means clustering method is used to assign data to one of K clusters using the distance from the means of these clusters. A data vector is assigned to the nearest cluster mean. After all data vectors are classified, the means are updated using the sample means of the data vectors assigned to that cluster. The process is iterated until convergence (i.e., the means do not change significantly when compared against a precision threshold). The result is a set of clusters that are as compact and well-separated as possible. The K initial values for the cluster means were chosen randomly from the data set. These initial values can cause K-means to converge to a local minima, where the total sum of distances are a minimum, from which a better solution may exist. To avoid this, K-means was repeated a number of times and if different local minimums exists then the case with the lowest total sum of distances, over all repetitions, was returned.

The value K can be arbitrarily chosen based on examination of the gaze tracking data, or simply by increasing the number of clusters to see if K-means can find a better grouping of the data. One method to automatically determine K is to determine how well-separated the resulting clusters are and choose a K which gives maximum separation. A silhouette score can be used to measure how close each point in one cluster is to points in neighbouring clusters. This measure ranges from $+1$, indicating that the points are very distant from neighbouring clusters, through to 0, indicating points that are not distinctly in one cluster or another, to -1, indicating points that are probably assigned to the wrong cluster. The average of the silhouette scores for each K can be used as a quantitative measure to compare different K's. In this paper, K-means was repeatedly performed by increasing K by 1 at each stage until the silhouette score for the grouping of data for $K+1$ is less than that for K. In such a case, K would give a maximum silhouette score

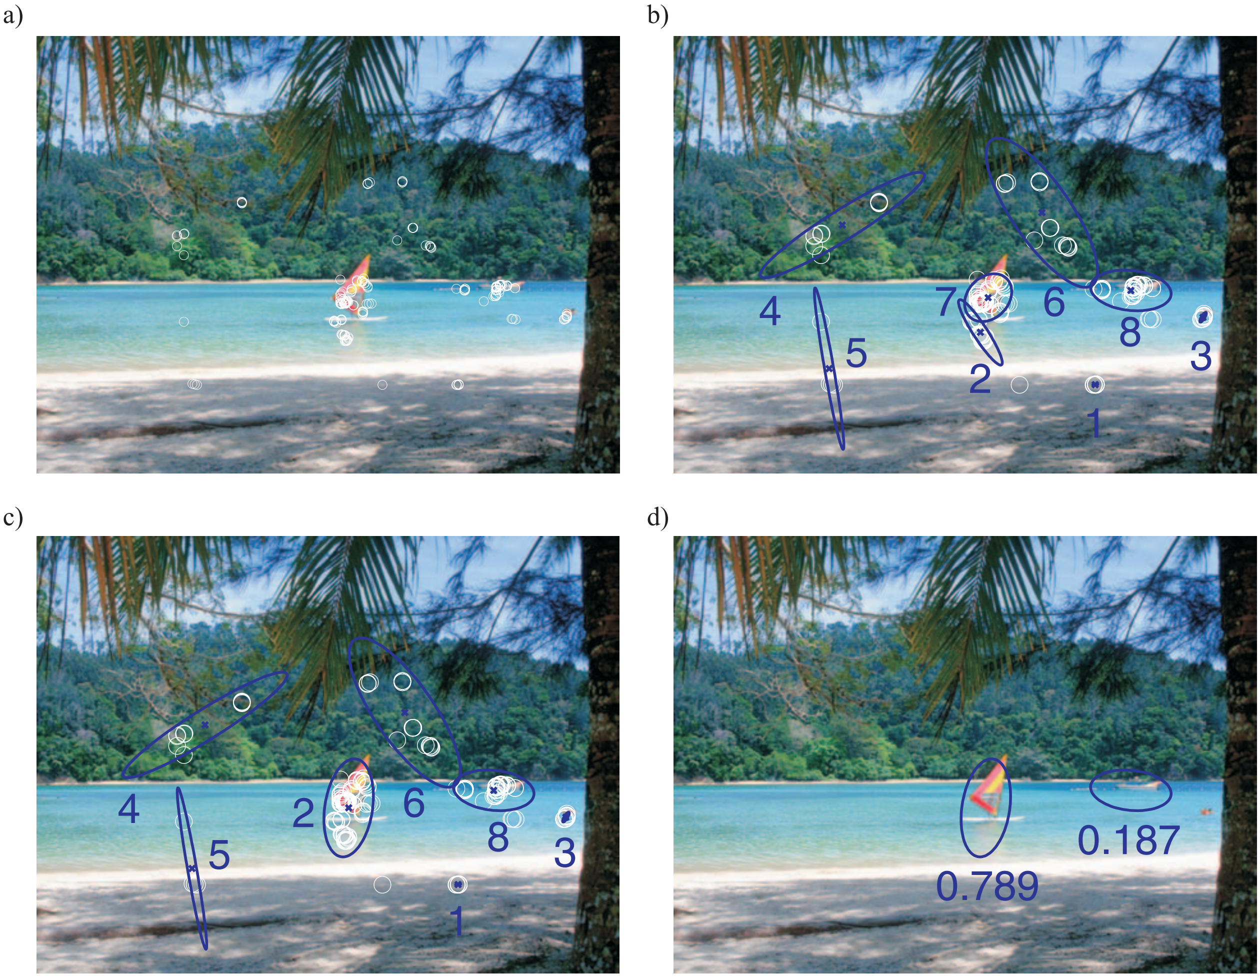

e)

ROI cluster	1	2	3	4	5	6	7	8
Importance score *	0.001	**0.789**	0.002	0.004	0.000	0.017	NA	**0.187**
* Values in bold are above the ROI threshold of 0.15 and are retained for ROI prioritisation.								

Fig. 4. An illustrated example of the ROI clustering and mapping process: (a) gaze locations (circles) recorded by an eye tracker; (b) output of K-means clustering process showing cluster loci of ROI candidates (ellipses); (c) set of ROI cluster candidates after merging of close clusters; (d) final ROI importance map output after culling of "unimportant" clusters; (e) importance scores using the *visit weighted (cluster count)*^{2} metric for the clusters shown in Fig. 4c.

and the data vectors and mean of the clusters that correspond to K would represent the "best" grouping of the data.

In some cases, several clusters may be close together and a procedure is required to merge the two clusters. The rule used to merge was if any two cluster means fall within a distance threshold of 10% of the average of the image dimensions, then the two clusters would merge into one. The merging process would often merge multiple clusters belonging to the same object into a single cluster. The number of clusters, K, and the cluster mean are updated as the clusters are merged.

The cluster means and covariances of the data vectors that were assigned to the clusters were used to generate ellipses to represent the loci of ROI candidates. The major and minor radial components of the ellipses were chosen to be 2.58 standard deviations in each direction. In such a case, if the cluster's spatial distribution was Gaussian, then this will represent approximately 99% of data points belonging to the cluster. The total ROI size (area bounded by all the ellipses) was also restricted to less than 25% of the image space. This is to ensure that during the encoding process, the reconstructed quality of ROIs more than compensates for the overhead in encoding the ROI [4]. If the ROI area did not satisfy this condition, then the process was repeated using $K = K + 1$.

An illustrated example of the ROI clustering process using the "beach" image is shown in Fig. 4. The "beach" image contains a number of objects of interest that a viewer may gaze upon, namely the windsurfer and perhaps the boat and people swimming (towards the right hand side of the image). For a particular viewer, the location of gaze points recorded by the eye tracker are shown in Fig. 4a. The plot does not contain any information about the sequence these points were viewed in. Figure 4b shows the grouping of data output by the clustering process described above. In the case shown, the number of clusters that resulted in a maximum silhouette score was $K = 8$. Finally, as shown in Fig. 4c, the set of clusters was refined by merging clusters whose geometric means were close together. In the example shown, cluster 2 was merged with cluster 7 and the new cluster mean and loci were recomputed.

The quality of the ROI clustering results is very dependent on the gaze data being clustered. If the ROI in an image is large, such as those in the "rockclimb" and "yacht" image, then multiple clusters may result for the single object. There were also cases where viewers only fixated on a particular region of a ROI, such as the head of the cow, which meant that the clustering procedure will not produce a ROI loci which would encompass the whole object. Other problems that may exist is that some gaze points not belonging to the ROI may be included in the ROI cluster simply because the gaze point was closer to the ROI cluster than any other clusters. These problems can be overcome by improving the clustering algorithm to reduce the sensitivity of "outlier" gaze points and/or by having viewers get more experience with the eye tracker hardware and be more aware of the purpose of the task required for the application at hand. Since most viewers were naive to the purpose of the gaze tracking experiment, a diversity in range of gaze patterns resulted.

4.2. ROI mapping

Given that the ROI clustering procedure outputs K candidate ROIs, a ROI mapping procedure is required to assign an importance measure or score to each ROI. This importance score can be interpreted as the degree of importance of the ROI relative to other regions in an image. Regions with a high importance score represents regions of high importance, which should be retained and prioritised by the encoder with higher priority than regions with a low importance score, which are to be removed and prioritised along with the image background.

It is conjectured that the fixation-saccade sequence provided by the gaze patterns would reveal underlying visual attentional processes that can be used to develop an importance metric. Several factors have been previously considered such as cluster count, distance, variance, area, and revisit count, and were combined using an entropy weighting procedure to weight each factor accordingly [5]. The ROI mapping procedure was found to be sound for the gaze data and particular test image. Here, such a complex metric may not be appropriate to model an importance metric that would suit all viewers and gaze patterns.

In this paper, a subset of factors which are intuitive from the gaze pattern sequences are studied. This includes the count or duration of gaze points belonging to a cluster and it's sequential behaviour in terms of the number of visits and revisits to a given cluster. These factors provide several possible derivations for an importance metric. The importance measures investigated in this paper are *cluster count*, $(\textit{cluster count})^2$ and *visit weighted* $(\textit{cluster count})^2$.

The *cluster count* measures the number of gaze points that belong to a given cluster. This factor is analogous to the duration of gaze within the cluster, since uniform gaze sampling was recorded. This factor represents the total time spent viewing/gazing at that region. The cluster's count was mapped to a range 0 to 1 by dividing the measures by its sum. The mapped range indicates a cluster's relative importance for the given metric.

To determine whether or not a cluster is classified as a ROI or not, a threshold of 0.15 was applied to the measures. A cluster importance measure greater than 0.15 would be retained and considered as an ROI, else the clusters would be considered as part of the background for ROI coding purposes. The threshold was chosen such that the number of clusters classified as a ROI, in general, would be slightly larger than the number of "actual" ROIs. This would minimise the number of ROIs that would be misclassified, while taking into account the fact that each ROI may contain a number of clusters.

The performance of the metric was evaluated in terms of ROI misses and ROI false alarms. ROI misses are those cases where the ROI mapping algorithm did not pick up a primary object of interest in an image as an ROI, while a ROI false alarm is the case where the algorithm considered a cluster as a ROI when it contains no object of interest. Because of the varied range in objects fixated by a viewer and the number of objects that may exist in an image, an "intuitive" rule-based ROI definition was formulated to define the ROI. The following rule-based definitions were used:

- "Beach" image – ROI must contain the windsurfer. The boat and people swimming on the right of the image were not considered as ROI false alarms.
- "Cow" image – ROI must contain the cow.
- "Horse" image – ROI must contain the horse.
- "Paddock" image – ROI must contain the horse in the foreground. The horses in the background and the barn are not ROI false alarms.
- "Rockclimb" image – ROI must contain the upper body of the rock climber. The rock climber's lower body were not considered to be a ROI false alarm.
- "Yacht" image – ROI must contain at least the centre of the yacht. Other parts of the yacht were not ROI false alarms.

With the above ROI definition, the *cluster count* metric contained 20 ROI misses and 48 ROI false alarms. The square of the number of cluster gaze points, (*cluster count*)^2, was found to be a more useful metric as it emphasises regions with a high cluster count and penalises those with a small cluster count. This resulted in a much improved ROI performance with 13 ROI misses and 41 ROI false alarms.

Furthermore, it was hypothesised that by making use of the number of visits (or revisits) to a given cluster during the course of viewing would provide additional information to determine the cluster's importance. The more visits to a given cluster, the more important the cluster should be. The *visit weighted (cluster count)*^{2} metric is given by the square of the number of cluster gaze points multiplied by n, where n is the number of visits to the cluster under consideration. To stop those clusters with a high number of visits from having a dominant effect on the cluster importance, the visit weight was capped at 3. This value was chosen since clusters with a number of visits greater than 3 were statistically unreliable with on average less than 1 cluster per gaze pattern with a number of visit greater than 3.

Table 1 shows the performance in terms of ROI misses and ROI false alarms for the three ROI mapping methods. It can be seen that *visit weighted* (*cluster count*)^2 improves the ROI performance even further by reducing the ROI misses to 7, while only marginally increasing the number of ROI false alarms.

Table 1
The ROI misses and ROI false alarm results for three ROI mapping metrics *

Importance metric	Number of ROI misses	Number of ROI false alarms
Cluster count	20 (8.7%)	48 (21.0%)
(Cluster count)^2	13 (5.7%)	41 (17.9%)
Visit weighted (cluster count)^2	7 (3.1%)	45 (19.7%)

* Values in parentheses are percentages of the total number of gaze data.

It should be noted that the majority of ROI misses and ROI false alarms are contributed by only a few viewers. Table 2 provides an indication of the distribution of ROI misses and ROI false alarms across the viewers for the *visit weighted* (*cluster count*)^2 metric. Notice that only two viewers contributed to the ROI misses, while a varying amount of ROI false alarms were contributed by different viewers. The large number of ROI false alarms indicates that viewers have their own viewing preferences and fixated on other regions in addition to the defined ROIs. As suggested earlier, the ROI performances can be improved if the viewers had more experience with the eye tracker and were more informed of the purpose of the task that was required.

The importance scores using the *visit weighted* (*cluster count*)^2 metric for the illustrated example shown in Fig. 4c is tabulated in Fig. 4d. The cluster importance scores that are in bold font represents those clusters retained as an ROI (i.e., importance score > 0.15). Note that only the windsurfer and the boat are retained as ROIs and represents the objects of interest found important for the particular viewer. The degree of importance is represented by the value of the importance score.

The duration and order of the cluster visits were also considered but the results did not provide an improved ROI performance over the *visit weighted* (*cluster count*)^2 case. Viewer's gaze patterns were too varied and these measures did not apply globally across all gaze data sets. If the underlying visual attentional processes of each viewer can be known, improved performances can be gained.

5. JPEG 2000 ROI image coding

The coding/decoding of images may be influenced to enhance the image quality in ROIs. JPEG 2000 provides several ROI coding mechanisms which can prioritise predefined ROIs, such as the max-shift [1, 2] and implicit [2] ROI coding methods. The problem with these methods are that implementations, such as that in [6], treat all ROIs with the same degree of importance and thus all pre-defined ROIs will be emphasised and prioritised with the same level of priority regardless of their degree of importance. To overcome this, an importance prioritised JPEG 2000 (IMP-J2K) image coder [7, 8] was developed to extend the concept of the Implicit method to incorporate an importance map to quantitatively model multiple ROIs and variable ROI importance scores. With IMP-J2K, the ROI was emphasised by weighting the mean square error (MSE) distortion measure of a block of coefficients by the square of its importance score. The reconstruction of the ROIs are bounded by the extent of these blocks. This is advantageous for the ROIs generated from the clustering process, since the ROIs may not fully encompass the objects in the image.

The ROI cluster loci and importance measures as generated by the ROI mapping stage can be input to IMP-J2K for ROI encoding, with an additional background importance parameter of, say, 0.01. Figure 5 shows an example of reconstructed images as the ROI encoded image is progressively (or incrementally) transmitted and received by the recipient for the illustrated example in Fig. 4. Note that only cluster 2 and 8 (in Fig. 4) were prioritised and emphasised during the encoding process. The ROIs, especially the windsurfer, can be observed to be reconstructed with better quality and at a higher resolution than the rest of the image, especially during the earlier stages of transmission when only a part of the code-stream has been received. However, when the entire code-stream has been received, the lossless (or near lossless) representation of the image is possible. The gaze-influenced selective coding of parts of an image provides a user of the image to receive the image as desired by the image author.

Table 2

The ROI misses and ROI false alarms contributed by viewers for the *visit weighted* (*cluster count*)^2 ROI mapping metric

Viewer	1	2	3	4	5	6	7	8	9	10	11	12	13	Total
Number of ROI misses	0	0	0	0	0	0	0	0	0	0	2	0	5	7
Number of ROI false alarms	0	1	1	8	6	1	0	3	4	0	4	7	10	45

Fig. 5. Progressive decoding of a ROI prioritised code-stream at: (a) 0.0625; (b) 0.125; (c) 0.25; (d) 0.5; (e) 1.0; (f) 2.0 bits per pixel (bpp) for the beach image example in Fig. 4. The ROIs improve most rapidly at reduced bit-rates, while the visually lossless (or near-lossless) reconstruction of the image as a whole is possible at higher bit-rates.

6. Conclusions

This paper has presented a system, called Gaze-J2K, which uses a combination of eye tracking and JPEG 2000 image coding, to allow an image author to customise the encoding of an image for users of an application. The system collects spatial and temporal gaze information from a viewer, uses the gaze pattern to automatically locate and assign importance to a representative subset of ROIs, and subsequently encode these regions with higher priority. Experimental results show that the accuracy of determining ROIs can be as high as 97% and can be further improved for experienced viewers. The system can be used in various applications such as the Internet and world wide web, which require the display of visual content to adapt to the end user.

References

[1] "Information technology – JPEG 2000 image coding system – Part 1: Core coding system," ISO/IEC 15444-1, Aug. 2002.

[2] D. S. Taubman and M. W. Marcellin, *JPEG 2000: Image Compression Fundamentals, Standards, and Practice*. Boston: Kluwer, 2002.

[3] EyeTech Digital Systems, "Quick glance eye-gaze tracking system", 2005, http://www.eyetechds.com/

[4] A. P. Bradley and F. W. M. Stentiford, "JPEG 2000 and region of interest coding," in *Proc. Digit. Im. Comput. Techn. Appl.*, Melbourne, Australia, 2002, pp. 303–308.

[5] A. Nguyen, V. Chandran, and S. Sridharan, "Visual attention based ROI maps from gaze tracking data," in *Proc. Int. Conf. Im. Proces.*, Singapore, 2004, pp. 3495–3498.

[6] D. Taubman, "Kakadu software: a comprehensive framework for JPEG 2000", 2005, http://www.kakadusoftware.com/

[7] A. Nguyen, V. Chandran, and S. Sridharan, "Importance prioritisation in JPEG 2000 for improved interpretability," *Sig. Proces. Im. Commun.*, vol. 19, no. 10, pp. 1005–1028, 2004.

[8] A. Nguyen, "Importance prioritised image coding in JPEG 2000", Ph.D. thesis, School of Engineering Systems, Queensland University of Technology, Brisbane, Australia, Jan. 2005.

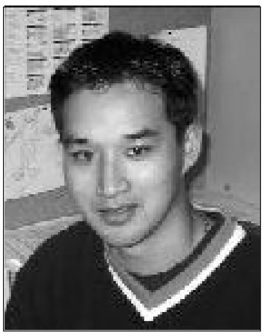

Anthony Nguyen received the B.E. (aerospace avionics) degree with first class honours in 1999, and a Ph.D. degree in the area of image processing and image compression in 2005 from Queensland University of Technology (QUT), Brisbane, Australia. He is currently a Research Fellow within the Image and Video Research Laboratory, School of Engineering Systems at QUT. He is also a tutor and a telecoms lab development coordinator for the digital communications and wireless communications units. His research interests include image processing, image compression, and pattern recognition. He is a Member of the Institute of Electrical and Electronic Engineers – IEEE (USA).

e-mail: a.nguyen@ieee.org
Image and Video Research Laboratory
School of Engineering Systems
Queensland University of Technology
GPO Box 2434, Brisbane QLD 4001, Australia

Vinod Chandran received the B.T. degree in electrical engineering from the Indian Institute of Technology, Madras, India, in 1982, the M.Sc. degree in electrical engineering from Texas Tech University, Lubbock, in 1985, and the Ph.D. degree in electrical and computer engineering and the M.Sc. degree in computer science from Washington State University, Pullman, USA, in 1990 and 1991, respectively. He is currently an Associate Professor at the Queensland University of Technology (QUT), Brisbane, Australia, in the School of Engineering Systems. His research interests include pattern recognition, higher order spectral analysis, speech processing, and image processing. He is a Senior Member of the Institute of Electrical and Electronic Engineers – IEEE (USA) and the Chairman of the IEEE Computer Society in Queensland.

e-mail: v.chandran.qut.edu.au
Image and Video Research Laboratory
School of Engineering Systems
Queensland University of Technology
GPO Box 2434, Brisbane QLD 4001, Australia

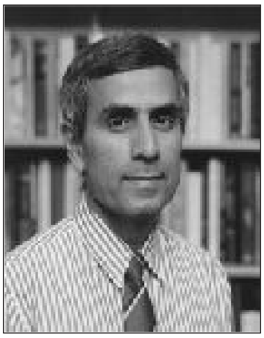

Sridha Sridharan has a B.Sc. (electrical engineering) degree and obtained a M.Sc. (communication engineering) degree from the University of Manchester Institute of Science and Technology (UMIST), UK, and a Ph.D. degree in the area of signal processing from University of New South Wales, Australia. He is currently with the Queensland University of Technology (QUT), where he is a Professor in the School of Engineering Systems. He is also the Program Leader of the Research Program in Speech, Audio and Video Technology at QUT. In 1997, he was the recipient of the award of Outstanding Academic of QUT in the area of research and scholarship. He is a Fellow of the Institution of Engineers, Australia, a Senior Member of the Institute of Electrical and Electronic Engineers – IEEE (USA) and the Chairman of the IEEE Queensland Chapter in Signal Processing and Communication.

e-mail: s.sridharan@qut.edu.au
Image and Video Research Laboratory
School of Engineering Systems
Queensland University of Technology
GPO Box 2434, Brisbane QLD 4001, Australia