Issues in multi-view autostereoscopic image compression

Druti Shah and Neil A. Dodgson

Computer Lab, University of Cambridge, New Museum Site, Pembroke St,
Cambridge CB2 3QG, UK.

Abstract

Multi-view auto-stereoscopic images and image sequences require large amounts of space for storage and large bandwidth for transmission. High bandwidth can be tolerated for certain applications where the image source and display are close together but, for long distance or broadcast, compression of information is essential. We report on the results of our two-year investigation into multi-view image compression. We present results based on four techniques: differential pulse code modulation (DPCM), disparity estimation, three-dimensional discrete cosine transform (3D-DCT), and principal component analysis (PCA).

Our work on DPCM investigated the best predictors to use for predicting a given pixel. Our results show that, for a given pixel, it is generally the nearby pixels within a view that provide better prediction than the corresponding pixel values in adjacent views.

This led to investigations into disparity estimation. We use both correlation and least-square error measures to estimate disparity. Both perform equally well. Combining this with DPCM led to a novel method of encoding, which improved the compression ratios by a significant factor.

The 3D-DCT has been shown to be a useful compression tool, with compression schemes based on ideas from the two-dimensional JPEG standard proving effective.

An alternative to 3D-DCT is PCA. This has proved to be less effective than the other compression methods investigated.

1. INTRODUCTION

In real life we gain three dimensional information from a variety of cues. Two important cues are stereo parallax: seeing a different image with each eye, and movement parallax: seeing different images when we move our heads. These cues are not provided by two-dimensional display devices. Multi-view auto-stereo displays, such as the ones developed at the University of Cambridge[1–4], provide these cues, which aid in the visualisation of complex three-dimensional structures and in remote manipulation applications.

Multi-view auto-stereo displays work by displaying not just one or two, but many different pictures of the scene. Each picture is taken from a slightly different position and is visible in a corresponding "window" in space, so that each of the viewer's eyes sees a different image of the displayed scene (stereo parallax), just as in real life, and the viewer can move his/her head to "look around" objects in the scene (movement parallax). The result is auto-stereoscopic vision with no need for any special glasses or other headgear.

To allow for multi-view auto-stereo image broadcast, transmission, and storage, a very high bandwidth for transmission and a large storage space is required. High bandwidth can be tolerated for certain applications where the image source and display are close together but, for long distance or broadcast, compression of information is essential. The authors have conducted a two-year investigation into compression of multi-view imagery. This paper presents an overall view of the results obtained in lossless and lossy image compression. Many compression techniques have been developed for two-dimensional images and for time sequences. There has also been considerable work on compression of two-view stereoscopic images but there is little prior work on the encoding of multi-view autostereoscopic images.

The investigation was carried out on four compression mechanisms: differential pulse code modulation(DPCM), disparity estimation, three-dimensional discrete cosine transform(3D-DCT), and principal component analysis(PCA).

Further author information: (Send correspondence to Neil Dodgson)
E-mail: Neil.Dodgson@cl.cam.ac.uk

The work on DPCM investigated predictors for a given pixel. Our results show that, for a given pixel, it is generally the nearby pixels *within* a view that provide better prediction than the corresponding pixel values in *adjacent* views. This led to investigations into disparity estimation to provide better correlation between pixels in adjacent views. Use of correlation and least-square error measures were made to estimate disparity. Both measures perform equally well. Combining these measures with DPCM led to a novel method of encoding, whereby similar scan-lines are collected into blocks, and each block then has its disparity estimated, rather than estimating the disparity for each scan-line independently.

The 3D-DCT has been shown to be a useful compression tool, with compression schemes based on ideas from the two-dimensional JPEG standard proving effective. The DCT can transform a matrix of pixels to a matrix of "spatial frequency information" which contain many small (or zero) entries. As the response of the human eye is frequency dependent this factor can be exploited by discarding the information the eye can not perceive without affecting the overall image quality. Since a multi-view frame is very homogeneous, an investigation into the performance of the 3D-DCT system for a number of different multi-view images using different quantisation methods has been carried out. The results indicate that multi-view images can be coded efficiently.

An alternative to DCT is PCA. This technique is one of the simplest multivariate methods designed to represent the original data set by reducing the number of variables. The transformation is performed by finding a set of uncorrelated variables whose linear combinations model the original data set, thus measuring the different "dimensions" of the data set. This allows the PCA technique to be applied to a wide range of problems. This has proved to be less effective than the other compression methods investigated.

The paper has been divided as follows. Section 2 gives a brief review of the Cambridge auto-stereo display and other multi-view auto-stereo displays. Section 3 discusses compression issues for multi-view images. Section 4 investigates the behaviour of the third dimension using entropy measures. Section 5 describes the mechanisms used for disparity estimation. Section 6 discusses the aspects of 3D-DCT and Section 7 details the Principal Component Analysis technique and results. Finally Section 8 summarises the work.

2. THE DISPLAY

Stereoscopic displays offer added realism over conventional two dimensional display devices. They are especially desirable for visualisation and remote manipulation, where the extra dimension provides much needed depth perception.

There are a variety of systems for stereoscopic viewing. "3D with glasses"[5] and auto-stereo systems are two broad classes. The former is reasonably well known and requires the viewer to wear special glasses. These displays present two different images in the same plane. The glasses select which of the two images is visible to each of the viewer's eyes. Such displays have not gained wide acceptance, partly owing to the need to wear glasses. Auto-stereo displays provide stereoscopic perception without glasses and could, therefore, prove more commercially viable.

The principle of multi-view auto-stereo displays such as those developed at Cambridge is illustrated in Figure 1. Figure 1(a) shows an observer looking at a scene. He sees a different image of the scene with each eye and different images again whenever he moves his head. He is able to view a potentially infinite number of different images of the scene. Figure 1(b) shows the same viewing space divided into a finite number of horizontal windows. In each window only one image, or view, of the scene is visible. However, the viewer's two eyes each see a different image and the images change when the viewer moves his head — albeit with jumps as the viewer moves from window to window. Thus both stereo and horizontal movement parallax cues can be provided with a small number of views. There is no fundamental restriction to horizontal movement parallax: vertical movement parallax can also be provided, but this squares the number of views. The finite number of views required in Figure 1(b) allows the replacement of the scene by an auto-stereoscopic display that outputs a different image to each window (Figure 1(c)).

There are various ways of manufacturing such displays.

- Spatially multiplexed systems: In such systems the resolution of a display device is split between the multiple views; For example, parallax barriers[6–8] and lenticular sheets[9–12] have both been used to divide the resolution of a display device between multiple views. The display is almost always a liquid crystal device, because this allows relatively simple alignment of the barrier or lenticules with the pixel structure.

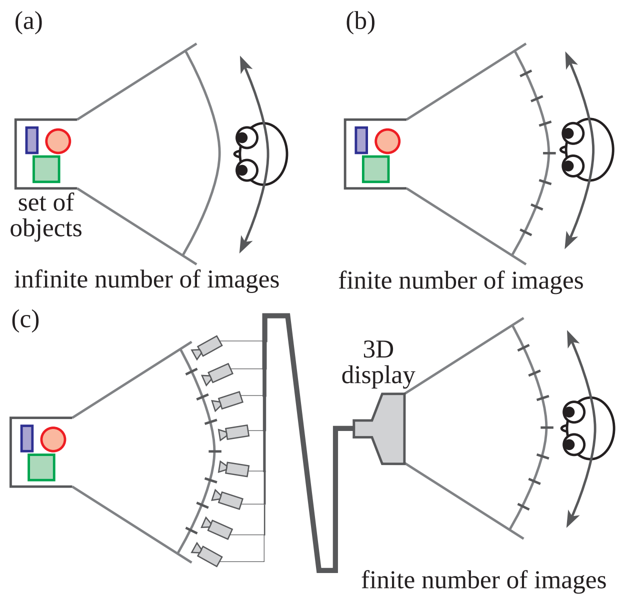

Figure 1. (a) When viewing a scene in real life, an observer sees a different image with each eye (stereo parallax). When he moves his head he sees different images (movement parallax) There are an infinite number of different images of the scene that he could see. (b) The number of different images is made finite, each visible in its own window. Each eye still sees a different image and different images are seen when the head is moved (c) An auto-stereoscopic 3D display provides a different image to each window, producing both stereo and movement parallax with a small number of views. Note that for the purpose of illustration, a radial camera system has been depicted whereas a parallel camera configuration is used for the Cambridge auto-stereo displays.

- Multi-projector systems: These devices use a single projector for each view[13] projecting their images onto a special transmissive or reflective screen, such as a double lenticular sheet. They suffer from the problems of expense: one projector per view becomes exorbitant for even a reasonable number of views; and of alignment: the projected images must be aligned precisely with one another.
- Time-sequential systems: These displays use a single display device running at a high frame rate. A secondary optical component is required to direct the images to the appropriate zones in space. The displays developed at Cambridge[1–4] are of this type.
- Hybrid systems: Combining two of the above mechanisms can produce a system with a higher number of views, at the expense of more complex technology. Combining spatial multiplexing and multi-projector has led to prototype 40-view[14] and 72-view[11] displays. A simpler 7, 13, or 21-view hybrid system has been designed by Hines.[15] The combination of time-sequential and multi-projector methods has led to the development of the latest Cambridge displays.[2,3] Kalai *et al*[16] developed a new auto-stereo display with motion parallax, which combines a parallax-barrier panaramagrams with time multiplexing.

Starks[17] presents a comprehensive survey paper on stereoscopic imaging technology. Dodgson[18] describes the categories of auto-stereo display in more detail.

Multi-view displays range from as few as four to many tens of views. While compression of two-view imagery has been investigated in some depth, the compression of multi-view images has received little attention.

3. MULTI-VIEW COMPRESSION ISSUES

3.1. Image capture

There are two principle ways of arranging the cameras for capturing multi-view imagery. The radial method has all the camera principal axes passing through a single point (e.g. Figure 1(c)) whereas the parallel method has all the camera axes parallel. For the type of multi-view imagery used by the Cambridge auto-stereoscopic displays (and most other auto-stereoscopic displays), the parallel method gives correct results. This is because the multiple views are all displayed on the same surface/screen and all the capture devices must have their principal axes perpendicular to that surface and therefore parallel to each other. In addition to the parallel axes, the images must be captured off-axis, i.e. the field of view of each camera is not located directly in front of the camera, as shown in Figure 2. A thorough analysis may be found in.[19,20]

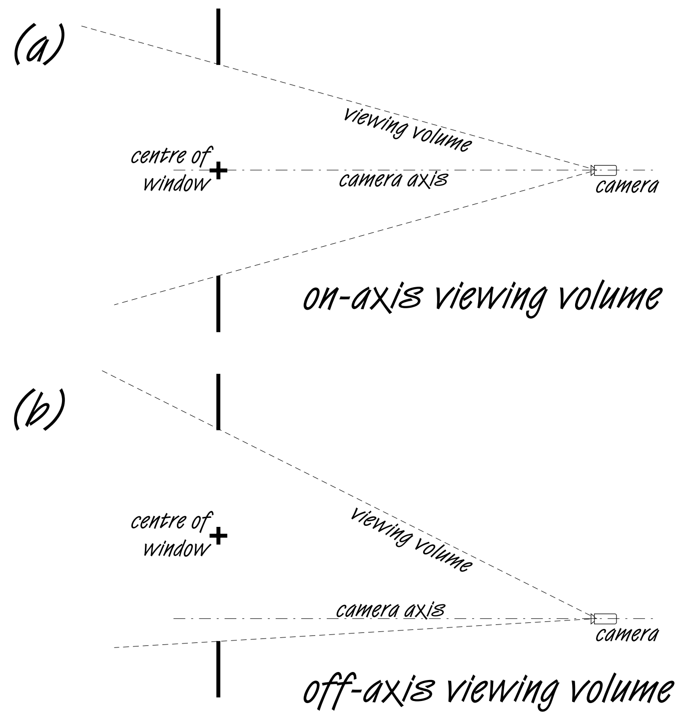

Figure 2. (a) an on-axis viewing volume, where the centre of the window (screen) is directly in front of the camera, (b) an off-axis viewing volume where the centre of the window is not directly in front of the camera.

3.2. Previous Work

Current image compression techniques are applicable to 2D static images and time sequences of images. These techniques could be applied to the individual views in a view-point sequence, however this would not take advantage of the high degree of similarity found between views, as the differences between views are more constrained that those between frames of a sequence. Most of the work to date on 3D compression has been for two-view stereoscopic images. The issues concerning stereo images such as disparity estimation, spatial correlation, MPEG compatible coding issues, the human visual system and standard 2D image and video coding techniques can all be utilized for multi-view images.

A recent overview of 3D compression and representation is given by Naemura et al.[21].

4. DCPM

Many 2D compression schemes treat a 2D image as a 1D stream of data. The order in which the 2D image data is input to the 1D compression algorithm can impinge on the compression ratio. With auto-stereo 3D imagery the extra dimension allows for many more options. At its simplest, where a 2D image could be scanned as either (x, y) or (y, x), a 3D image could be scanned up to six ways including (x, y, θ), (y, θ, x) and (θ, x, y). In addition, the behaviour of data in the x and y dimensions can be expected to be similar, while in the θ dimension the data may behave in quite a different manner.

To investigate how the behavior differs in terms of the amount of redundant information found in each frame, zero order and first order entropy measures were employed. Zero order entropy indicates the amount of compression possible if each pixel is treated independently. First order entropy indicates the compression which is possible if each pixel's value is naively predicted from the previous pixel, where "previous" could mean previous of any one of x, y or θ.

Experiments on comparing between views and then between layers, where 'layer' consists of a single row taken from all views, have shown that there is better coherence between layers than views. For example, the average zero order entropy for the layers was found to be 6.61 bpp which was lower when compared to the zero order entropy for views (7.4 bpp) indicating that the average layer contained less variability than the average view. Detailed results are given in Shah and Dodgson.[22] Therefore, naively using one view to predict the next is not going to be as good as simply predicting from one scan-line to the next within a view. There are two directions in which to proceed: one is to consider more complex predictors, as described by Penrose and Dodgson.[23] The other, is to exploit the disparity shift between views.

5. DISPARITY ESTIMATION

Differences in multi-view images of real and computer generated scenes are caused by the relative displacement of the cameras. These differences are important because they encode information that often allows a partial reconstruction of the 3D scene structure from 2D projections. Matching between the sets of points from two images gives rise to two important issues. Firstly, how to select points for matching and secondly how to determine which matches are correct. Shah and Dodgson[24] discuss these issues with respect to space search measures and hierarchical row decomposition. Since the cameras are equispaced with parallel axes, the only disparity shifts are horizontal. Thus the disparity estimation problem is simplified as the algorithm needs to find matches in corresponding scan-line of the other views.

Search space measures are measures of similarity. They work by considering a small region in image 1 which surrounds an interesting point and a search is then made in image 2 for the region of maximum similarity. The two measures compared are the correlation measure and least mean square error measure. Although the search space measures performed well, the overhead of encoding the shifts reduced the advantages of the method. This lead to considering ways of combining disparity estimation with the knowledge that there is often good correlation between layers.

A novel method of encoding was developed "hierarchical row decomposition" (HRD), whereby similar scan-lines were collected into blocks, and each block then has its disparity estimated, rather than estimating the disparity for each scan-line independently. Significant results in terms of entropy and matching were obtained when combined with the use of windowing.

6. DCT

We now turn to lossy compression methods. DCT forms a key role in several image compression standards including JPEG[25] for still picture compression, and the various MPEG[25] standards for audio-visual compression and communication.

It transforms a 2D matrix of pixels to a 2D matrix of spatial frequency information containing many small (or zero) entries. As the response of the human eye is frequency dependent, the DCT is able to separate the perceptually important information in an image from the information that cannot be perceived by the eye. The information that cannot be seen can be discarded without affecting the overall image quality. This is achieved by the transformed matrix being quantised and efficiently encoded.

An image sequence can be visualised as a volume of data and therefore conceived as a three dimensional signal. Thus extending the 2D-DCT to 3D-DCT has some inherent merits[26] such as high compression of homogeneous video and a symmetric codec, although it is reported to produce perceptually annoying artifacts in the temporal dimension.

Since multi-view images are homogeneous in nature, because of the method of capturing the images, we expect the 3D-DCT to perform well.

For a 3D-DCT scheme, the 2D based quantisation tables become inappropriate since they cannot account adequately for the variation in the third dimension. This leads to the investigation and development of a quantisation volume where volume can be viewed as a $8 \times 8 \times 8$ multi-dimensional array with dimensions x, y and θ. Figure 3 depicts a 3D volume. Furthermore, JPEG uses a particular scan order for the quantisation coefficients, to further reduce the number of bits required to represent them. An appropriate scan order for 3D needs to be established.

After experimentation with a new scanning order and different sets of quantisation volumes it was found that the 3D-DCT performed well in comparison to the standard JPEG method. Detailed results are given in Shah and Dodgson[27].

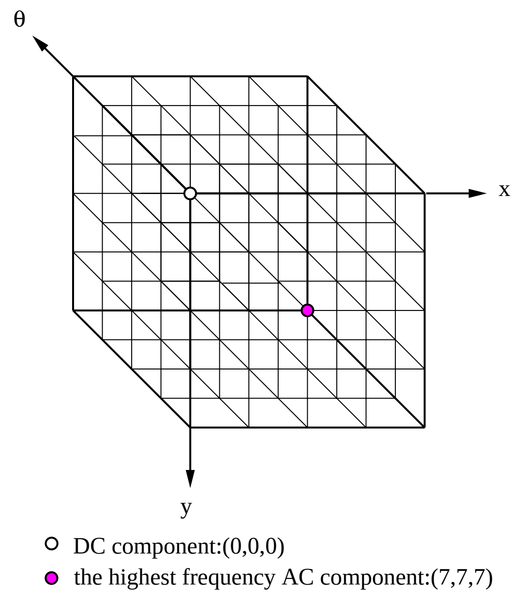

Figure 3. A three dimensional volume

7. PCA

Finally, we have adopted the technique of Principal Component Analysis (PCA) to a multi-view sequence as the method is designed to represent the original data set by reducing the number of variables. The transformation is caused by finding a set of uncorrelated variables whose linear combinations model the original data set, thus measuring the different "dimensions" of the data set. The following section describes the method used.

7.1. Mathematical Description

The PCA method is applied to image coding by forming a set of eigenvectors corresponding to a set of images of the multi-view sequence. Using the images $\mathbf{I}_1, \mathbf{I}_2, \cdots, \mathbf{I}_M$, as an example, where M is the number of frames in the training set, the average image is calculated and a set of images are derived by subtracting the average image from each image in the set. The images are converted into a single-column vector by concatenating consecutive rows of pixel values. The average image is given by

$$\bar{\mathbf{I}} = \frac{1}{M} \sum_{i=1}^{M} \mathbf{I}_i, \tag{1}$$

and each difference image is given by

$$\phi_i = \mathbf{I}_i - \bar{\mathbf{I}}, \qquad i = 1, \cdots, M \tag{2}$$

A set of eigenvectors $\mathbf{E}_l$ which maximize λ_l are sought where

$$\lambda_l = \frac{1}{M} \sum_{i=1}^{M} (\mathbf{E}_l \phi_i)^2, \tag{3}$$

subject to $\mathbf{E}'_p\mathbf{E}_i = 0$, for $p < i$,. The vectors $\mathbf{E}_l$ and scalar λ_l are, respectively, the eigenvectors (i.e. eigenpictures) and eigenvalues of the covariance matrix.

$$\mathbf{C} = \frac{1}{M}\sum_{i=1}^{M}\phi_i\phi'_i, \tag{4}$$

Since the number of images in the training set is less than the dimension of the space $M < N_c \times N_r$, where N_c is the number of pixels on a line and N_r the number of lines, there will only be $M-1$ nonzero eigenvectors. Using this training image set, an $M \times M$ matrix $\mathbf{L}$ is constructed and the eigenvectors $\mathbf{v}_l$ of $\mathbf{L}$ are obtained, where

$$L_{mn} = \phi'_m\phi_n, \tag{5}$$

The $M-1$ orthogonal eigenpictures $\mathbf{E}_l$ are given by

$$\mathbf{E}_l = \sum_{i=1}^{M} v_{li}\phi_i, l = 1, \cdots, M-1 \tag{6}$$

This set of $M-1$ eigenpictures span a basis set, which best describe the scene distribution and the first few eigenpictures carry the maximum variation. Having obtained the eigenpictures from the training ensemble, images outside this set are represented as follows: The new image $\mathbf{I}''$ is transformed into its eigenpicture components by the operation,

$$\omega_l = \mathbf{E}'_l(\mathbf{I}'' - \bar{\mathbf{I}}), l = 1, \cdots, M-1 \tag{7}$$

The pattern vector $\mathbf{\Omega}^T = [\omega_1, \omega_2, \cdots, \omega_{M-1}]$ describes the contribution of each eigenpicture in representing the new image. The new image can then be re-synthesised using the representative set of coefficients and the eigenpictures. This is given by:

$$\mathbf{I}'' = \sum_{l=1}^{M-1}\omega_l\mathbf{E}_l + \bar{\mathbf{I}} \tag{8}$$

7.2. Experimental Results

Experiments were carried out on the Granny multi-view sequence. The sequence was made up of 300 frames with 16 views. For the training set a group of 40 images out of the (300×16) images were used to produce the eigenpictures. The remaining frames were used for evaluation purposes. Thus the number of bits required to code an image was equal to the product of the number of eigenpictures and the number of bits per coefficient.

Initially all the eigenpictures were used to code the test images. It was found that variations occurring in the image sequence were effectively captured. The subjective quality of the images was not adequate since the reconstructed images showed reflective and ghosting artifacts. This was due to the sequence being computer generated which contained high frequencies(i.e sharp edges) and did not contain any natural fuzziness. This lead to eigenpictures having a high degree of edge information. Figure 4 shows an original and reconstructed image from the test set. The figure shows a small section of the Granny image sequence with each image comprising of three horizontal bands. The bands were selected as those parts of the image which contain a high degree of information.

In the next experiment we considered sub-images obtained by: 1) grouping rows of an image having a correlation coefficient larger than a specific threshold and termed it as a block and 2) concatenating corresponding blocks from each image in the training set. The above method is a simple version of the Hierarchical Row Decomposition technique.

It was found that the variations were captured efficiently and the subjective quality of the reconstructed test set was good. Figures 5 and 6 show an example of an original and decoded sub-image.

The conclusion drawn from the two experiments showed that it was possible to code new images with respect to blocks but not with respect to views. This may be due to the fact that we are dividing a very high frequency image into small sub-images where the overall high frequency is distributed.

Furthermore the lack of any real-world multi-view sequences (as opposed to computer-generated sequences), means that we could not test the efficiency of the technique for such images and have left it as an open problem.

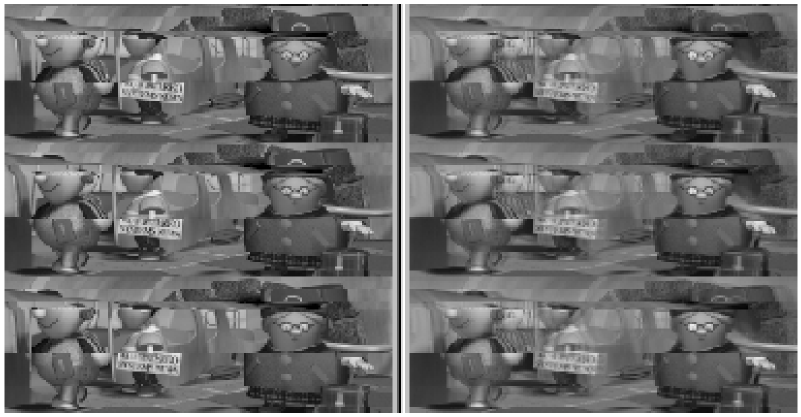

Figure 4. Selected sections of the original and reconstructed Granny sequence,three sections of an image are chosen which contain high degree of information: no of eigenpictures used = 4, bits per coefficient = 3

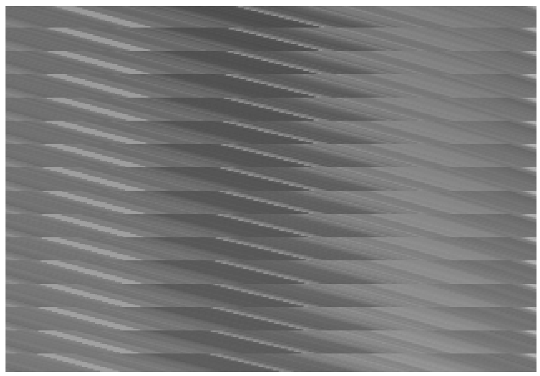

Figure 5. Original section showing layers having high degree of correlation

8. CONCLUSION

This paper reports on the results of our two-year investigation into multi-view image compression. Four different techniques have been used.

The lossless methods: DPCM and disparity estimation, have shown that it is possible to code multi-view images at an acceptable compression rate.

The DPCM has shown that, for a given pixel, it is generally the nearby pixels *within* a view that provide better prediction than the corresponding pixel values in *adjacent* views. To prove the above statement, the research was led into investigating disparity estimation using search space measures. It was observed that although the search space measures performed well, the overhead of encoding the shifts reduced the advantages of the method. This lead to considering ways of combining disparity estimation with the knowledge that there is often good correlation between rows within a view. A novel method of encoding was developed, whereby similar scan-lines were collected into blocks, and each block then has its disparity estimated, rather than estimating the disparity for each scan-line independently. Significant results in terms of entropy and matching were obtained when combined with the use of windowing.

Parallel investigations into encoding methods for lossy compression were carried out. The 3D-DCT technique which is a JPEG-like encoding method showed that multi-view images can be coded efficiently. Since a multi-view frame is very homogeneous, an investigation into the performance of the 3D-DCT system for a number of different multi-view images using different quantisation methods has been carried out. For this research a new scanning order was obtained which proved to be significantly better than a simplistic extension of the standard zig-zag order to 3D.

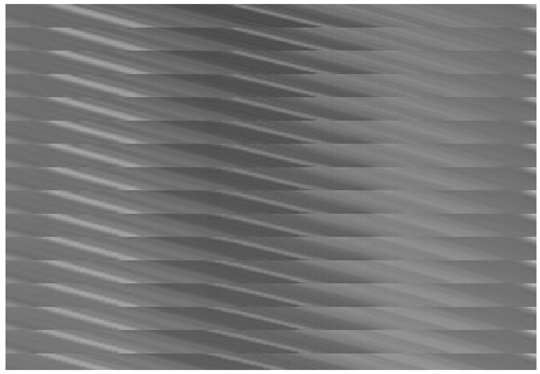

Figure 6. Reconstructed section showing layers having high degree of correlation: no of eigenpictures used = 20, bits per coefficient = 3

A set of quantisation volumes proved to be effective in reproducing good quality images as well as a similar or better compression rate when compared to the 2D JPEG coding scheme.

An alternative coding scheme used was the PCA. The conclusion drawn from the two experiments showed that it was possible to code new images with respect to blocks but not with respect to views. This may be due to the fact that we are dividing a very high frequency image into small sub-images where the overall high frequency is distributed.

Overall, we have produced results which point the way to efficient lossy and lossless compression of multi-view auto-stereo images and image sequences.

REFERENCES

1. J. R. Moore, N. A. Dodgson, A. R. L. Travis, and S. R. Lang. Time-multiplexed color autostereoscopic display. In *Proc. SPIE*, volume 2653, pages 10–19, 1996.
2. N. A. Dodgson, J. R. Moore, S. R. Lang, G. Martin, and P. Canepa. Time-sequential multi-projector autostereoscopic 3-D display. *J. Soc. for Information Display*, 8(2), 2000.
3. N. A. Dodgson, J. R. Moore, S. R. Lang, G. Martin, and P. Canepa. A 50" time-multiplexed autostereoscopic display. In *Proc. SPIE*, volume 3957, 2000.
4. G. Martin, A. Smeyne, J. R. Moore, S. R. Lang, and N. A. Dodgson. 3-D visualization without glasses - a large-screen autostereoscopic display. In *Proc. SPIE*, volume 4022, 2000.
5. D. F. McAllister. *Stereo Computer Graphics and Other True 3-D Technologies.* Princeton University Press, 1993.
6. J.B. Eichenlaub. Developments in autostereoscopic technology at dimension technologies inc. In *Proc. SPIE*, volume 1915, pages 177–186, 1993.
7. H. Isonon, M. Yasuda, D. Takemori, H. Kanayama, C. Yamada, and K. Chiba. 50-inch autostereo full-color 3-D display system. In *Proc. SPIE*, volume 1669, pages 176–185, 1992.
8. H. Morishima, H. Nose, N. Taniguchi, K. Inoguchi, and S. Matsumura. Rear cross lenticular 3-D display without eyeglasses. In *Proc. SPIE*, volume 3295, pages 193–202, 1998.
9. T. Okosi. *Three-Dimensional Imaging Techniques.* Academic Press, 1996.
10. C. van Berkel and J. A. Clarke. Characterisation and optimisation of 3D-LCD module design. In *Proc. SPIE*, volume 3012, pages 179–186, 1997.
11. K. Matsumoto and T. Honda. Research of 3-D display using the anamorphic optics. In *Proc. SPIE*, volume 3012, pages 208–219, 1997.
12. A. Arimoto, T. Ooshima, T. Tani, and Y. Kaneko. Wide viewing area glassless stereoscopic display using multiple projectors. In *Proc. SPIE*, volume 3295, pages 186–192, 1998.
13. G. R. Little, C. Gustafson, and V. E. Nikolaou. Multiperspective autostereoscopic display. In *Proc. SPIE*, volume 2219, pages 388–394, 1994.

14. J. D. Montes and P. Campoy. A new three-dimensional visualisation system based on angular image differentiation. In *Proc. SPIE*, volume 2409, pages 125–140, 1995.
15. S. P. Hines. Autostereoscopic video display with motion parallax. In *Proc. SPIE*, volume 3012, pages 208–219, 1997.
16. A. Kalai and M. Siegel. Improved rendering of parallax panoramagrams for a time-multiplexed autostereoscopic display. In *Proc. SPIE*, volume 3295, pages 835–838, 1998.
17. M. Starks. Stereoscopic imaging technology. 3DTV Corporation, 1996. http://www.3dmagic.com/.
18. N. A . Dodgson, J. R. Moore, and S. R. Lang. Multi-view autostereoscopic 3D display. In *International Broadcasting Convention*, pages 497–502, 1999.
19. N. A. Dodgson. Resampling radially captured images for perspectively correct stereoscopic display. In *Proc. SPIE*, volume 3295, pages 100–110, 1998.
20. N. A . Dodgson, J. R. Moore, and S. R. Lang. Time-multiplexed autostereoscopic camera system. In *Proc. SPIE*, volume 3012, pages 72–83, 1997.
21. T. Naemura, M. Kaneko, and H. Harashima. Compression and representation of 3-D images. *IEICE trans. Inf. & Syst.*, E82-D(3):558–567, March 1999.
22. D. Shah and N.A. Dodgson. Multi-view autostereo display and image compression. submitted to IEEE Transactions on Image Processing, 2000.
23. A. J. Penrose and N. A. Dodgson. Error resilient lossless image coding. In *ICIP*, Kobe, Japan, October 1999.
24. D. Shah and N.A. Dodgson. Compression issues in multi-view autostereoscopic displays. In *Proc. SPIE*, volume 3974, pages 994–1005, 2000.
25. V. Bhaskaran and K. Konstantinides. *Image and Video Compression Standards: Algorithms and Architectures*. Kluwer Acedemic, 1997.
26. R.K.W. Chan and M.C. Lee. 3d-dct quantization as a compression technique for video sequences. In *International Conference on Virtual Systems and MultiMedia*, pages 188–196, 1997.
27. D. Shah and N.A. Dodgson. Compression of multiview images using a 3d discrete cosine transform. submitted to IEEE Transactions on Image Processing, 2000.