RECONSTRUCTING MISSING REGIONS IN COLOUR IMAGES USING MULTICHANNEL MEDIAN MODELS

Steven Armstrong, Anil C. Kokaram and Peter J.W. Rayner
Signal Processing and Communications Laboratory,
Cambridge University Engineering Department,
Cambridge CB2 1PZ, England
Tel: +44 (0)1223 332767; fax: +44 (0)1223 332662
e-mail: (sa210, ack, pjwr)@eng.cam.ac.uk

ABSTRACT

This paper presents a method for reconstructing missing regions in colour images. A multichannel median model is proposed as the underlying image model and a statistical framework is employed to generate sampled realisations of the missing data. The nature of the model leads to a posterior expression for the missing data that does not involve an easy to manipulate multivariate probability distribution. Therefore, the problem of sampling is solved using a numerical approach. Results are included which show that this approach leads to excellent reconstructions.

1 INTRODUCTION

The reconstruction of missing blocks of image data in grey-scale images has been addressed in a number of papers [5],[4] and [9]. These responded to the need to reconstruct severe degradation that a combination of the unstable film stock of early films and the process of running a film through a projector causes. The restoration techniques developed are now being applied to these movies in digital film archives.

The loss of blocks of data is also a problem in the world of colour. Projection can damage a colour film and errors that result in missing blocks of data can arise during transmission and recording of digital and analogue image data. Furthermore, in implementations of MPEG-II that incorporate error concealment, small blocks of missing data can be observed when random errors are introduced into the bit stream. The need for a method of reconstructing missing image data in colour images is therefore evident.

Three further problems that are considered are the choice of a colour space and excitation process and the need for a model selection process to choose the appropriate support for the vector median model.

2 THE VECTOR MEDIAN MODEL

The image model that is used in this paper arises from the vector median filter, which is one of several multichannel filters used successfully for noise reduction in colour images [1], [10] and [6]. The vector median, $\mathbf{x}_{med}$, of a set of n-dimensional vectors, $\mathbf{X} = \{\mathbf{x}_q : q = 1, \ldots, N\}$, is defined as

$$\mathbf{x}_{med} = \arg \min_{\mathbf{x}_j \in \mathbf{X}} \sum_q \|\mathbf{x}_q - \mathbf{x}_j\| \qquad (1)$$

where $\|\cdot\|$ is an L_p norm – the L_2 norm was used to obtain the results in this paper. When n equals 1, Equation 1 is equivalent to the median filter.

It is proposed that a pixel in a colour image be modelled as a vector median *prediction* of the pixel $\mathbf{x}_k$, based on the set of pixels $\mathbf{X}'$, in some neighbourhood of the pixel. This is defined as

$$\mathbf{x}_k = \arg \min_{\mathbf{x}_j \in \mathbf{X}'} \sum_{q, q \neq k} \|\mathbf{x}_q - \mathbf{x}_j\| + \mathbf{e}_k \qquad (2)$$

where the error in the prediction, $\mathbf{e}_k$, is modelled by a 3-dimensional Laplacian excitation process.

2.1 The Laplacian Excitation Model

The n-dimensional Laplacian distribution used has the form

$$L_n(0, \mathbf{C}) = \frac{|\,\mathbf{C}^{-1}\,|}{2^n} \exp(\|\mathbf{e}^T \mathbf{C}^{-1}\|_1) \qquad (3)$$

where $\mathbf{C}$ is the covariance matrix, $|\,\mathbf{C^{-1}}\,|$ is the determinant of the inverse covariance matrix and $\|\cdot\|_1$ is the L_1 norm, which sums the absolute values of the elements of its vector or matrix argument.

A Laplacian model is chosen, as opposed to the standard Gaussian, because the errors observed from typical applications of vector median filters to images are distributed with the sharper peak and heavier tail of the Laplacian. In the single channel case the median is the maximum likelihood estimate of location for the Laplacian[8].

3 SAMPLING FROM THE POSTERIOR DISTRIBUTIONS

The process of reconstructing the missing data by ‘sampling’ requires the determination of the posterior density for the missing data. It follows from Bayes rule that this density is proportional to the likelihood of the data multiplied by the prior for the covariance, as follows

$$p(\mathbf{z} \mid \mathbf{y}, \mathbf{C}) \propto p(\mathbf{z}, \mathbf{y} \mid \mathbf{C}) p(\mathbf{C}) \quad (4)$$

where $\mathbf{z}$ is the set of missing pixels and $\mathbf{y}$ is the set of known pixels in the image and $\mathbf{C}$ is the covariance of the excitation process.

By taking a uniform prior probability for $\mathbf{C}$, the likelihood for the missing data in the image becomes proportional to the posterior density for the missing data. The likelihood, assuming that the residuals are i.i.d., is given by

$$\begin{aligned} p(\mathbf{z}, \mathbf{y} | \mathbf{C}) &= \prod_{i=1}^{N} p(\mathbf{e}_i) && (5) \\ &= \frac{1}{(2^3 \mid \mathbf{C}^{-1} \mid)^N} \times && (6) \\ & \quad \exp(-\|(\mathbf{e_1}, \mathbf{e_2}, \ldots, \mathbf{e_N})^T \mathbf{C}^{-1} \|_1) \end{aligned}$$

where $\{\mathbf{e_1}, \mathbf{e_2}, \ldots, \mathbf{e_N}\}$ are the 3×1 error vectors corresponding to the N missing data points.

3.1 THE GIBBS SAMPLER

A *Gibbs Sampler* is used to draw samples for the missing data [3]. This method is implemented such that it decomposes the problem into a series of univariate draws from the required marginal distribution. From an initial guess for the missing data and for the variance (denoted $\mathbf{z}^0$ and $\mathbf{C}^0$), the sampler proceeds iteratively according to

$$\left.\begin{aligned} z^1_{iR} &\sim p(z^0_{iR} \mid \mathbf{y}, \mathbf{C}^0, \mathbf{z}^0_{-i}, z^0_{iG}, z^0_{iB}) \\ z^1_{iG} &\sim p(z^0_{iG} \mid \mathbf{y}, \mathbf{C}^0, \mathbf{z}^0_{-i}, z^1_{iR}, z^0_{iB}) \\ z^1_{iB} &\sim p(z^0_{iB} \mid \mathbf{y}, \mathbf{C}^0, \mathbf{z}^0_{-i}, z^1_{iR}, z^1_{iG}) \end{aligned}\right\} \forall \mathbf{z}_i \in \mathbf{z} \quad (7)$$

$$\mathbf{C}^1 \sim p(\mathbf{C}^0 \mid \mathbf{y}, \mathbf{z}^1) \quad (8)$$

where $\mathbf{z}_{-i}$ is the set of missing pixels, excluding pixel $\mathbf{z}_i$.

The order in which the missing pixels are sampled for, is such that, if possible, the missing pixel to be sampled, does not appear in the support of the previously sampled pixel. This is known to lead to faster convergence of the Gibb's sampler [2].

The posterior distributions for the missing data are determined according to the following method:

1. Choose one missing pixel in the image and set the value of the red component to zero.

2. taking the set of those residuals in which the missing pixel in question plays a part in the calculation of the prediction, calculate the probability in equation 6. Store this result.

3. Increment the value of the red component of the missing pixel by 0.01.

4. Repeat the previous two steps until the red component is equal to 1 (assuming that the RGB values have been scaled to lie between 0 and 1).

This method is used to form posterior distributions for the R, G and B components of each missing pixel. Samples from these distributions are then drawn using the 'transformation method'.

3.2 Sampling for the Variance

By assuming that the covariance matrix is diagonal - a reasonable asumption in practice - the variance of the excitation process of each channel can be drawn separately from a Gamma distribution, from which samples can be drawn using the rejection method [7]. Taking one channel and denoting its variance as, c, the derivation of the distribution from which to sample is as follows

$$\begin{aligned} p(c|\mathbf{e}) &\propto p(\mathbf{e}|c)p(c) && (9) \\ &= c^{-N} \exp(-\frac{\|\mathbf{e}\|_1}{c}) && (10) \end{aligned}$$

where $p(c)$ is the prior for the variance. A non-informative uniform prior for the variance is used here.

This distribution can be transformed into a Gamma distribution with the following substitution

$$y = \frac{\|\mathbf{e}\|_1}{c} \quad (11)$$

which gives

$$\begin{aligned} p(y|\mathbf{e}) &= p(c|\mathbf{e})|dc/dy| && (12) \\ &= p(c|\mathbf{e})\frac{c}{y} && (13) \\ &= \frac{c^{-N+1}\exp(-y)}{y} && (14) \\ &\propto \frac{y^{N-1}exp(-y)}{y} && (15) \\ &= y^{N-2}exp(-y) && (16) \end{aligned}$$

This means that samples are drawn from

$$y \sim \Gamma(N-1) \quad (17)$$

and using equation 11, samples for c can be derived.

3.3 Subset Selection

An important part of the reconstruction method is a model selection stage which precedes the sampling. This proposes a number of different mask shapes and chooses the one that fits the image best.

This choice is arrived at by searching for the mask that gives the MAP (maximum a posteriori probability) estimate. These would be adversely affected by the missing data, so they are calculated using the known data that surrounds the missing block.

3.4 Choice of Colour Space

The choice of the RGB colour space over the YUV space in the development of the preceding sections was made to simplify the explanation of the method. In fact, there are two reasons why the YUV space might be preferable to the RGB space. The distance between YUV colour vectors corresponds better to the difference in colour precieved by humans. In addition, the intensity or Y channel of the YUV space is decorrelated from the U and V channels. This ties in better with the use of a diagonal covariance matrix, which implies that all the channels are decorrelated.

Nevertheless, the results presented in the next section are achieved using the RGB space, because it has one advantage over the YUV space. Any R, G or B value that lies in the specified range will lead to a valid colour. This is not the case for the YUV space where some sampled values do not correspond to actual colours. The problem can be overcome by clipping the resulting RBG values so that they fit in the required range.

4 RESULTS

The results presented here feature the 'clown' image (Figure 1), from which blocks of data are removed and replaced with random valued colour pixels. These blocks are then reconstructed using a number of methods. The first of these is the technique described in the paper. The second also follows the framework described in the paper, but the vector median becomes the single channel median and each colour channel is reconstructed separately. The final technique displaces a vector median predictor one pixel at a time, from left to right, top to bottom of the missing area, replacing each missing pixel with the vector median prediction of that pixel. Only one pass of the predictor is made.

The sampling and approaches can be seen to give very good reconstructions – far better than the recursive application of the vector median predictor. The false colouring that occurs with 'marginal' methods (those that treat each channel separately) can be seen clearly in Figure 3. The multidimensional vector median sampling approach suffers much less from this, although the sampling approach does lead to the occasional random pixel being incorrectly coloured.

Generally the vector median reconstructions appear much more pleasing. In areas of almost constant colour, the marginal method performs almost as well as the vector median sampling method. This phenomenum is elucidated by inspecting each of the RGB channels in turn. They turn out to be regions of fairly constant value, so the lack of information interchange between the three levels is therefore not essential to the success of the reconstruction.

The vector median sampling case is the slowest, taking 180 seconds with a nine point mask to do the 30 iterations usually needed to get to a converged distribution for each 6×6 block. For some of the edge features twice as many iterations were required. The longer time needed to calculate a vector median in comparison to a median means that this is slower than the marginal median method.

This compares with a marginal AR sampling approach which takes around 5 seconds to do each 6×6 block for each channel with a nine point mask. This is 100 times faster than the full vector median technique and highlights the advantage of being able to perform multivariate draws for the missing data from a known distribution. However, this approach also leads to false colouring in the same way as the marginal median sampling approach.

The vector median recursive approach takes a fraction of a second to do the all fifteen blocks. Although this is by far the fastest method and is commonly used to solve the missing data problem, it leads to very poor results.

All these times are quoted for programs running in C on a Pentium PC 90 MHz machine.

5 CONCLUSION

This paper has presented a method for reconstructing missing data in colour images. The results demonstrate that the method reconstructs features and colours very well and drastically improves on the recursive use of vector median predictors. The use of a multidimensional model reduces the occurrence of false colouring that is a problem with reconstruction techniques that treat each channel separately.

Current work is directed into developing faster methods of sampling and into studying the effect of using fast algorithms for the calculation of the vector median. Drawing samples for the full covariance matrix is also being investigated. The technique has also been applied successfully to reconstructing sequences of damaged colour images.

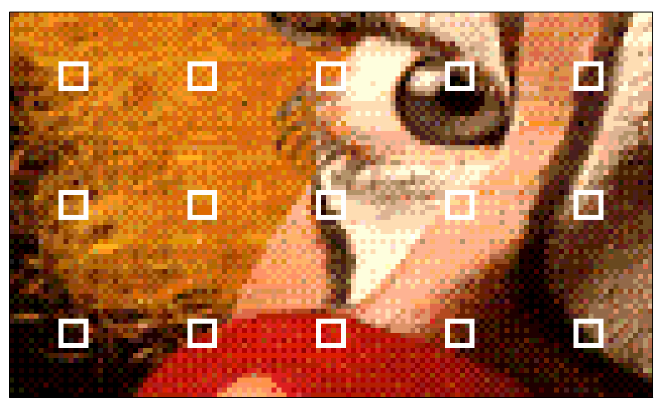

Figure 1: Original image and locations of missing blocks of data (the edges of the squares are part of the missing areas).

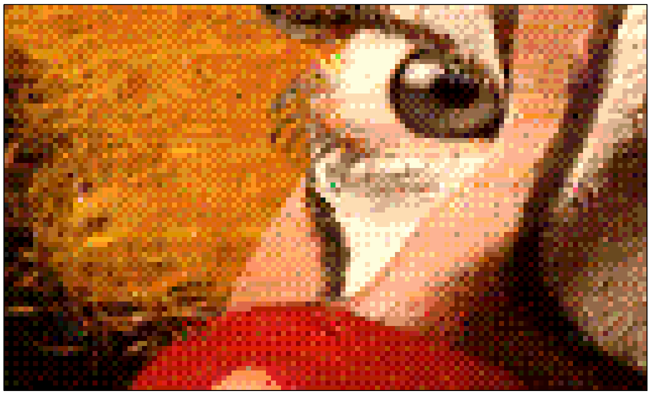

Figure 2: Reconstruction using vector median based sampling.

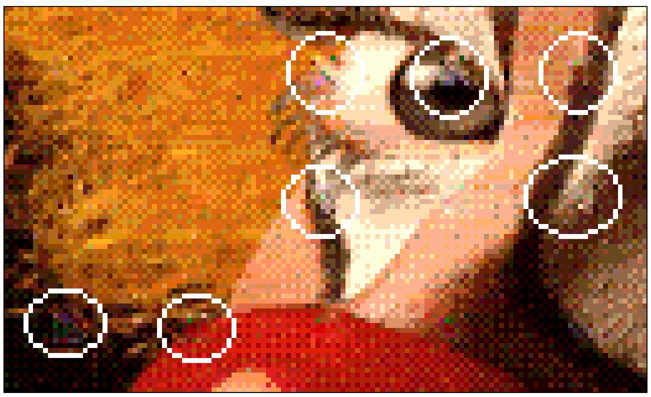

Figure 3: Reconstruction using median based sampling on each colour channel separately. Reconstructions with noticeable false colouring are ringed.

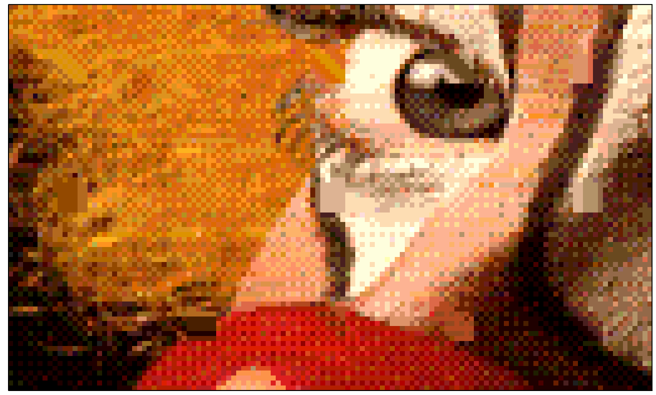

Figure 4: Reconstruction using the recursive application of vector median predictors.

References

[1] J. Astola, P. Haavisto, and Y. Nuevo. Vector median filters. In *Proceedings of the IEEE*, volume 78, pages 678–689, April 1990.

[2] J. Besag. On the statistical analysis of dirty pictures. *Journal of the Royal Statistical Society*, B(48):259–302, 1986.

[3] S. Geman and D. Geman. Stochastic relaxation, gibbs distributions, and the Bayesian restoration of images. *IEEE Transactions on Pattern Analysis and Machine Intelligence*, PAMI-6:721–741, 1984.

[4] Simon J. Godsill and Anil C. Kokaram. Joint interpolation, motion & parameter estimation for image sequences with missing data. In *Proceedings of the Eighth European Signal and Image Processing Conference (EUSIPCO)*, volume I, pages 1–4, September 1996.

[5] Anil C. Kokaram, Robin D. Morris, William J. Fitzgerald, and Peter J.W. Rayner. Interpolation of missing data in image sequences. *IEEE Transactions on Image Processing*, 4(11):1509–1519, November 1995.

[6] K.N. Plataniotis, D. Androutsos, V. Sri, and A.N. Venetsanapoulos. A nearest neighbor multichannel filter. *Electronic Letters*, pages 1910–1911, 1995.

[7] William H. Press, Saul William T. Vetterling, and Brian P. Flannery. *Numerical Recipes in C: The Art of Scientific Programming*. Cambridge University Press, 1992.

[8] D.M. Raymond and M.M. Fahmy. On the multivariate distributions of order statistics of discrete-state processes. In *IEEE Conference on Acoustics, Speech and Signal Processing*, pages 1717–1720, 1991.

[9] Paul Richardson and David Suter. Restoration of historic film for digital compression: A case study. In *IEEE International Conference on Image Processing(ICIP)*, volume 2, pages 49–52, October 1995.

[10] P.E. Trahanias and A.N. Venetsanapoulos. Vector directional filters: A new class of multichannel image processing filters. *IEEE Transactions on Image Processing*, 2:528–534, April 1993.