Chroma Reconstruction from Inaccurate Measurements

Alexander Balinsky
Cardiff University, UK
BalinskyA@cardiff.ac.uk

Nassir Mohammad
Cardiff University & HP Labs, UK
MohammadN3@cardiff.ac.uk

ABSTRACT

Non-linear filter responses of natural colour images have been shown to display non-Gaussian heavy tailed distributions which we call sparse. These filters operate in the YUV colour space on the chroma channel U (and V) using weighting functions obtained from the gray image Y. In this paper we utilise this knowledge for denoising the chroma channels of a colour image from inaccurate measurements. In our model the U (and V) elements are affected by noise, with a good version of the gray image Y obtainable through existing methods. We show that accurate reconstruction of the chroma components can be accomplished by solving an L1 constrained optimisation problem, where the sparse filter response on natural images is used as a regularization term. This scheme gives comparable results to leading commercial and state of the art denoising algorithms, and exceeds for chroma noise that does not correlate with the luminance structure.

Keywords: Natural images, filter response, sparse distributions, denoising, L^1 optimisation.

1 INTRODUCTION

Denoising is a fundamental problem in image processing due to the fact that images, no matter their content, usually contain some degree of noise. This is often regarded as a form of image degradation and the goal of denoising algorithms are to form an estimate x' of the the original image x given the observed noisy version x^*, modeled as

$$x^* = x + n, \tag{1}$$

where n is the matrix of the random noise pattern.

The principal causes of noise in digital images arise during image acquisition (digitization) and/or transmission. This can be caused by several factors such as low light levels, sensor temperature, electrical interference, malfunctioning pixels and interference in the channels used for transmission. The distribution of noise can be several, such as white, impulse or multiplicative, each giving its own characteristic form of degradation.

Various algorithms have been introduced with success over the past few decades for denoising images. The proposals, in their original form, have sparked an abundant literature resulting in many improvements in quality and speed. These algorithms can be categorized into several groups including Wavelets, Bilateral filtering, Anisotropic diffusion, Total variation and Non-local methods. Readers are advised to see [BUA05] and [MAI08] for comprehensive reviews and comparisons of the best available versions together with powerful novel approaches.

Some recent algorithms to mention include [LIU08] where the authors propose a unified framework for two tasks: automatic estimation and removal of colour noise from a single image using piecewise smooth image models. Their segmentation-based denoising algorithm is claimed to outperform current methods. This paper also contains an interesting introduction that discusses the current state of the art methods for image denoising. Another recent algorithm which claims to lead to excellent results is C-BM3D [DAB07]. In this scheme the authors propose an effective colour image denoising method that exploits filtering in a highly sparse local 3D transform domain in each channel of a luminance-chrominance colour space. For each image block in each channel, a 3D array is formed by stacking together blocks similar to it. The high similarity between grouped blocks in each 3D array enables a highly sparse representation of the true signal in a 3D transform domain, thus a subsequent shrinkage of the transform spectra results in effective noise attenuation.

The importance of denosing in image processing has also led to many commercial and freely available software. These include Neat Image, Noise Ninja, DenoiseMyImage, Photoshop, Topaz Denoise, Gimp and many more. The programs often incorporate a host of image enhancement tools to collectively remove typical forms of image degradation. A full evaluation of so many programs is difficult, especially since each has parameters which a user can change for subjective suitability. However, from general usage and reading it has been found that Noise Ninja and Neat Image are among the best used noise reduction programs. DenoiseMyImage is also a current alternative that uses a modified form of the state of art non-local means method. Readers may view [ALM] for a comprehensive *user* comparison of current software.

Denoising algorithms are usually fed a noisy *RGB* image corrupted in each channel. Most methods have been formulated as a channel by channel or vectorial model. In the former case the *RGB* values are mapped to a colour space such as YUV or *Lab* or any other suitable space to separate the luminance and chroma, with the denoising algorithm *usually* applied to each band. Since the luminance channel contains the main structural information and chroma noise is more objectionable to human vision (as opposed to the film grain appearance of luminance noise), separation allows more intensive denoising of the chroma channels without too much loss of detail. These models take into account the human perception of colour and allow us to handle the particular characteristics of the noise affecting each component. Methods based on their luminance-chromatic decomposition are well known for their excellent results with [DAB07] being a recent example. Furthermore, in the process of transmission, the reduction of bandwidth for the chroma allows errors and artifacts to be more easily compensated for than using a typical *RGB* model.

In this paper we propose a novel algorithm for removing noise from real images and also white and impulse noise from the chroma channels of an image in the colour space YUV, where a *good* version of the Y component is obtainable. (Due to the similarity of the colour components, from here on we interchangeably mention either the U or V channel, where analysis of the other is obtained by substitution). Algorithms such as those in [DAB07], [FOI07] and [BOR05] have successfully exploited the information in the luminance channel for effectively filtering the chroma components. In line with this philosophy our approach utilises the non-linear filter response distributions observed in [BAL09] as a regularization term (a prior, in Bayesian analysis) to penalize solutions that don't give a *desired* sparse solution when filtering.

The rest of the paper is laid out as follows: section 2 describes the motivating details behind our regularization function. Section 3 outlines our denoising procedure while section 4 gives results for denoising images. Section 5 summarizes the paper and directions for future work.

2 REGULARIZATION USING THE SPARSE DISTRIBUTION OF THE FILTER RESPONSE

Our approach in denoising the chroma components involves introducing a regularization term which incorporates knowledge of the statistics of natural images. More specifically we consider the recent non-linear filter response distributions of natural images observed in [BAL09]. In that paper the authors show that colour images, when filtered by the following:

$$F(U)(\mathbf{r}) = U(\mathbf{r}) - \sum_{\mathbf{s}\in N(\mathbf{r})} w(Y)_{\mathbf{rs}} U(\mathbf{s}), \tag{2}$$

display non-Gaussian heavy tailed distributions, i.e. sparse. Here $\mathbf{r}$ represents a two dimensional point, $N(\mathbf{r})$ a neighborhood (e.g. $3x3$ window) of points around $\mathbf{r}$, and $w(Y)\mathbf{rs}$ a weighting function. The proposed filter thus takes a point $\mathbf{r}$ in U (or in V) and subtracts a weighted average of chroma values in the neighborhood of $\mathbf{r}$. The $w(Y)\mathbf{rs}$ is a weighting function that sums to one over $\mathbf{s}$, large when $Y(\mathbf{r})$ is similar to $Y(\mathbf{s})$, and small when the two intensities are different. (See [BAL09] for further details).

The response of the filter can be modeled by a generalized Gaussian distribution (GGD)

$$f(x) = \frac{1}{Z} e^{-|x/c|^{\alpha}}, \tag{3}$$

where Z is a normalising constant so that the integral of $f(x)$ is 1, c the scale parameter and α the shape parameter. It is found for natural images that $\alpha < 1$ which results in a non-convex function. However, due to the recent success of L^1 optimization in recovering approximately sparse signals [CAN06], we convexify our model i.e. take $\alpha = 1$, and use this as a regularizer in (4).

3 CHROMA DENOISING PROCEDURE

We consider real noisy *RGB* images that have been corrupted by unknown noise which are then transformed to the YUV colour space. Due to the properties of the underlying natural colour images, such as high correlation between R, G, and B channels, we note that Y has higher SNR than U and V and that it contains most of the valuable information such as edges, shades, objects, texture patterns, etc. The U and V contain mostly low-frequency information with iso-luminant regions, i.e. variation in only U and V, being unlikely. Thus removing chroma noise through knowledge of gray information is plausible. We chose to use Neat Image or DenoiseMyImage when appropriate to denoise the Y channel when needed. We additionally used them as a benchmark for testing our algorithm. Furthermore, our algorithm is also tested against images in the YUV space suffering from impulse noise only in the chroma channels.

Thus, given the noisy chroma component U^* and a denoised gray image Y, our task is to recover a good approximation U' of the original element U. This model results in the following optimisation scheme,

$$argmin_{U'} \ ||F \cdot U'||_1 + \lambda ||U' - U^*||_d. \tag{4}$$

Given an $n \times m$ image, (we abuse the notation a little and have) F here is an $nm \times nm$ matrix whose rows correspond to filtering a single pixel where U' and U^* are $nm \times 1$ column first rasterized vectors. U' is the estimate we seek of U, while U^* is the noisy observation of U.

The first term is our penalizing function which takes small values for desirable solutions and the second is the *fidelity* term. The parameter d is taken to be either 2 or 1 reflecting the norms proposed in the measurement of the distance between the two vectors. In words, this optimisation scheme searches for the estimate image U' with the sparsest filter response and with the second term encouraging the solution to be close to a noisy chroma measurement U^*.

For an image assumed to be corrupted by Gaussian noise our reconstruction process involves solving (4) with $d = 2$, where the fidelity term encourages solutions to be close to the noisy version in the L^2 sense. When the noise is taken to be impulsive and affecting the image at random points by taking extrema values, we solve (4) with $d = 1$. Modifying the fidelity term to $d = 1$ (i.e. L^1 norm) has been studied with success within the Total Variation framework, as reviewed in [CHA05].

An important parameter in our algorithm is the value of λ which controls the relative weight of the difference between the noisy channel and the solution. Too small a value and the optimisation results in an overly smoothed output, while too high a value results in a solution that is too close to its noisy version. We found experimentally that $\lambda \in (0,5]$ gave the best results, with half-integer increments for optimality.

4 RESULTS

Our optimisation problem was solved using CVX [GRA] which is a convex programming package implemented in Matlab. The images that we used are of sizes in the region of 200×200 pixels, which took on average a couple of minutes to denoise. However, our aim here is not to pose a fast algorithm but only to show the applicability of such a scheme for denoising chroma channels. The algorithm is parameterised by the value of λ whose value is given in the text accompanying the figures.

Fig. 1(a) shows an example *RGB* image which is made severely noisy by adding Gaussian noise of mean zero and variance 0.01 to all the channels as shown in (b). (c) shows the denoised image obtained using Neat Image and (d) the result obtained using DenoiseMyImage. Neat Image was used at maximum setting while DenoiseMyImage was used at an adjusted medium level to obtain the best results. Neat image still left considerable noise like artifacts in the image, while DenoiseMyImage gave a less noisy but much smoother output. The result using our algorithm is shown in (e) where we used DenoiseMyImage to denoise the gray component. Visually comparing the results shows that our algorithm gives an intermediate result which is better than using NeatImage, while the colours are much more vibrant and appear sharper than when using DenoiseMyImage. This is also further justified by the peak signal to noise ratios (PSNR) which quantify the results, and shows our algorithm having a higher but similar value.

The next examples focus on real world images where the type of noise affecting the image is unknown. We begin with Fig. 2(a) which shows an image that is severely affected by colour noise. This is typical of an image taken in low light conditions with high ISO settings. (b) shows the image having been denoised using Neat Image. This program requires a suitable region to be selected for noise estimation, after which luminance and chrominance noise reduction can be individually adjusted. We required 100% noise reduction on all components due to the high amount of noise present in the image. (c) shows our algorithm where the luminance channel was denoised using Neat Image and the filter matrix F constructed from it for reconstructing the chroma channels. (d) shows the result of using DenoiseMyImage. We observe that our algorithm gives similar noise reduction compared to the existing methods, although on close inspection our result gives less colour aberrations.

Fig. 3(a) has been taken from some examples given on the Neat Image website. This is a crop of a television frame captured with a computer TV card. The image has strong colour banding visible across all the image caused by the electric interference in the computer circuitry. Similar banding is sometimes observed in digital camera images (caused by interference too). The banding degradation does not affect the luminance, however all channels still show grain like noise. (b) shows the best Neat Image result obtainable by denoising the chroma and luminance at 100%. However, the banding is still evident in the result. (c) is the result of our algorithm which clearly removes the noise. (d) is the best result obtainable using DenoiseMyImage which is still unable to remove the banding noise.

Our algorithm is able to remove this type of noise by filtering only the chroma channels and using Neat Image for clearing the fine grain luminance noise. The result is free of the colour banding and (f) shows that the V channel does not display any of this degradation against the V channel when using Neat Image (e). We are able to attain this result as we are filtering the chroma channels through taking account of the underlying gray level structure. Since the colour banding is not appearing in the luminance, minimisation of the filter response favours areas of homogeneous colours while the fidelity term bounds the colours to being close to the original.

The final two examples illustrate the flexibility of the model in handling chroma noise taking a different distribution. Fig. 4 shows an example of a clean image (a) which is transformed to the YUV colour space and impulse noise of density 0.05 added to the U and V channels only. Our algorithm, with the fidelity term measuring L^1 norm, is able to denoise such that the recombined *RGB* image shown in (b) is visually identical to the original. The detailed look at the chroma components reveals no sign of the impulse noise, while the PSNR is of a good value.

Fig. 5 shows another example of an image that has been corrupted by impulse noise and reconstructed. (a) shows the original image, (b) the *RGB* image with noise having been added to only the chroma channels and (c) shows our reconstructed image. The results illustrate again that noise has been successfully removed to a very high standard with good PSNR values, and this is further justified by looking at the chroma channels which have had their impulse noise removed. Neat Image and DenoiseMyImage are unable to effectively denoise the images affected by impulse noise. Instead we obtain a 'washed out' look with the impulse points still remaining. An example is shown by (d).

5 CONCLUSION

We have illustrated how knowledge of the statistics of natural images can be incorporated into an effective denoising scheme. Our objective was to propose a novel algorithm for removing chroma noise from digital images by operating in a luminance-chrominance colour space. We utilised the sparse filter response distribution of the filter studied in [BAL09] as a regularization term, and introduced a quadratic fidelity term to ensure the solution remained close to the original. This model allowed us to denoise real images with results comparable to current alternatives. The flexibility of the model was also shown by its ability to handle chroma impulse noise very effectively, giving results that are virtually identical to the original image. This was accomplished by altering the fidelity term to measure L^1 norm and shows concentration on gray level denoising gives sufficient information for colour channel reconstruction.

In future it would be most useful to robustly test this approach across diverse datasets of images and also in other colour spaces where we may observe increased performance. We are also looking at algorithms for solving the optimisation scheme much more quickly and looking at applying the approach to denoising hyperspectral images.

ACKNOWLEDGMENTS

This work was supported in part by grants from EPSRC and Hewlett Packard Labs awarded through the Smith Institute Knowledge Transfer Network.

Figure 1: Denoising example. (a) shows the original image, (b) the image with Gaussian noise added to all *RGB* channels. (c) is the result using Neat Image at maximum filtering. (d) shows the denoising result using DenoiseMyImage. (e) is the result obtained using our algorithm. PSNR: (c) 26.69, (d) 26.35, (e) 27.20 ($\lambda = 5$)

REFERENCES

[BAL09] A. Balinsky and N. Mohammad, "Non-linear filter response distributions of natural colour images." *LNCS* 5646, pp. 101-108. *Springer-Verlag Berlin Heidelberg* 2009.

[BUA05] A. Buades, B. Coll, and J. M. Morel, "A review of image denoising algorithms, with a new one, Multiscale Modeling & Simulation, vol. 4, no. 2, pp. 490530, 2005.

[MAI08] J. Mairal, M. Elad, and G. Sapiro, "Sparse representation for color image restoration", IEEE Transactions of image processing, vol. 17, No. 1, January 2008.

[LIU08] C. Liu, R. Szeliski, S.B. Kang, C.L. Zitnick, W.T. Freeman, "Automatic Estimation and Removal of Noise from a Single Image", IEEE Transactions on pattern analysis and machine intelligence, vol. 30, NO. 2, Feb 2008

[DAB07] K. Dabov, A. Foi, V. Katkovnik, and K. Egiazarian, "Color image denoising via sparse 3d collaborative filtering with grouping constraint in luminance-chrominance space", ICIP 2007. (Matlab code available at www.cs.tut.fi/ foi/GCF-

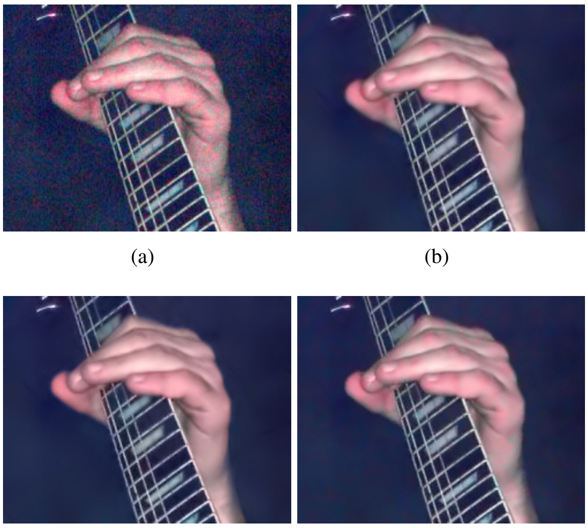

(a) (b) (c) (d)

Figure 2: Real image denoising example. (a) is an image that has been affected by severe chroma noise resulting in the appearance of 'blotches' of colour. (b) shows the denoised image obtained using Neat Image and (c) is obtained using our algorithm. (d) is the result obtained using DenoiseMyImage. We observe that all the reconstructions are visually similar, although on close inspection our result gives less colour aberrations. ($\lambda = 0.5$)

BM3D).

[ALM] M. Almond, http://www.michaelalmond.com/Articles/noise.htm

[FOI07] A. Foi, V. Katkovnik, and V. Egiazarian, "Pointwise Shape-Adaptive DCT for High-Quality Denoising and Deblocking of Grayscale and Color Images", IEEE Trans. Image Process., vol. 16, no. 5, May 2007.

[BOR05] D Borkowski, "Chromaticity Denoising using Solution to the Skorokhod Problem", Image Processing Based on Partial Differential Equations, Proceedings of the International Conference on PDE-Based Image Processing and Related Inverse Problems, CMA, Oslo, August 812, 2005.

[CAN06] E. Candes, "Compressive Sampling," *Int. Congress of Mathematics,* 3, pp. 1433-1452,Madrid, Spain, 2006.

[CHA05] T. F. Chan and S. Esedoglu, "Aspects of total variation regularized L^1 function approximation", SIAM J. Appl. Math., 65 (2005), pp. 18171837.

[GRA] M. Grant and S. Boyd, http://cvxr.com/cvx/

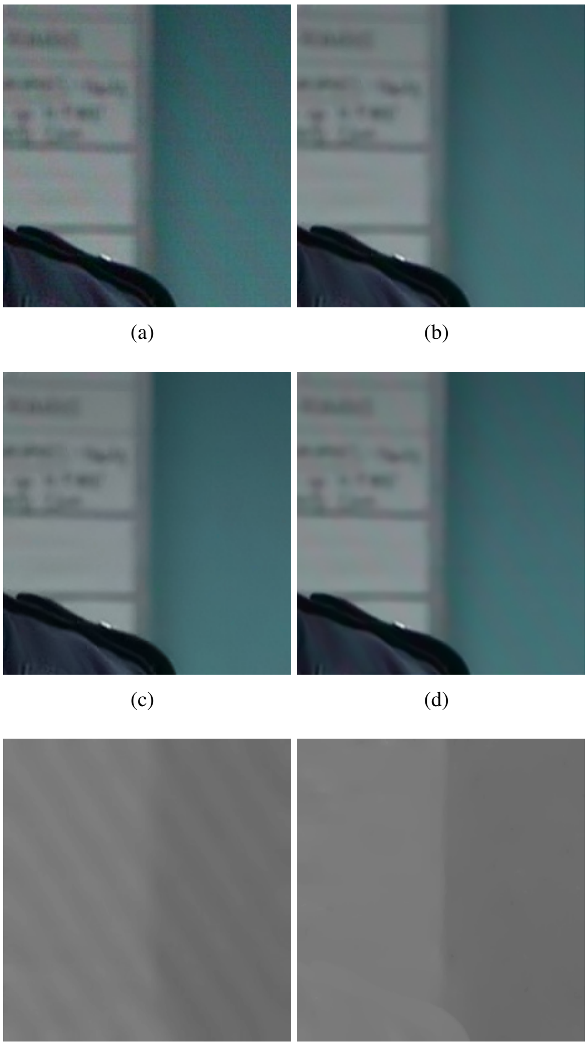

(a) (b) (c) (d) (e) (f)

Figure 3: Real image denoising example. (a) shows an example image affected by chroma noise that appears as bands in the colour channels. (b) is the result obtained using Neat Image which still leaves evident colour banding. (c) is our result which is able to remove the noise leaving a clean image as the colour banding does not correlate with the luminance structure. (d) is the best result obtained using DenoiseMyImage. (e) shows the banding still remaining in the V channel of the image when using Neat Image, while (f) clearly shows that the banding structure has been removed in our reconstructed V channel. ($\lambda = 0.1$)

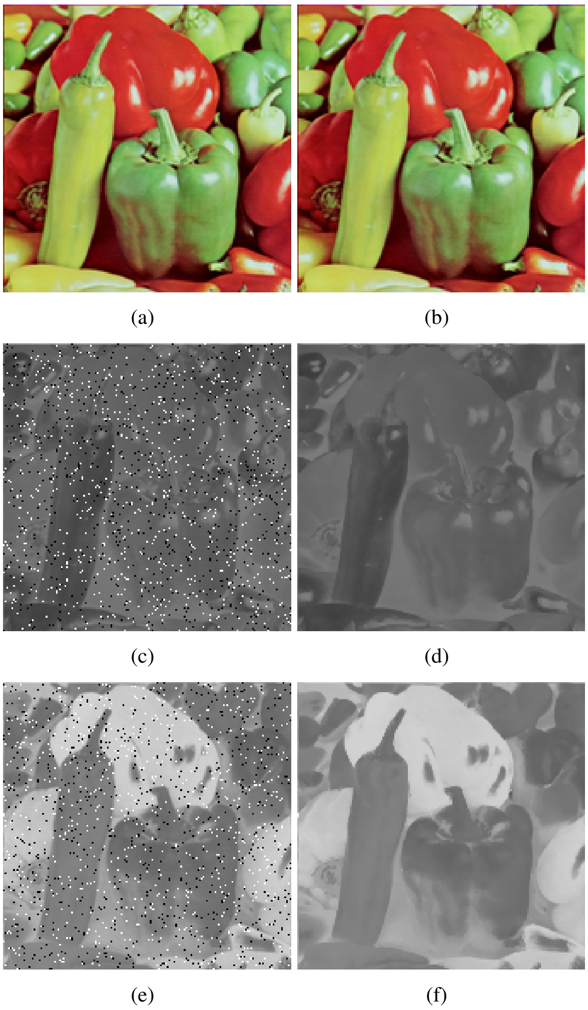

Figure 4: Impulse noise removal example. (a) shows the original image and (c) and (e) illustrate the colour channels with impulse noise added. (b) is the reconstructed image which does not display the impulse noise and is visually identical to the original. (d) and (f) shows the denoised chroma channels which have had their noise successfully removed. PSNR: (b) 37.68. ($\lambda = 0.5$)

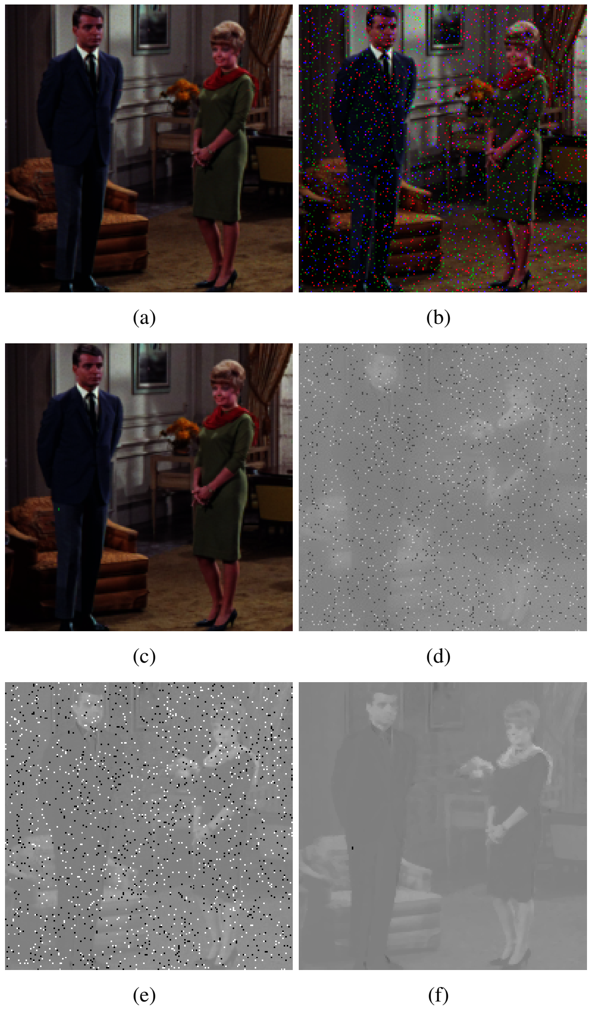

Figure 5: Impulse noise removal example. (a) shows an original colour image and (b) a noisy version that has had impulse noise added to the chroma channels in the YUV space. (c) is our reconstructed image which is virtually identical to the original. (d) is a typical result obtained using Neat Image or DenoiseMyImage. The impulse noise affecting the chroma is illustrated by (e) while the success of our algorithm for impulse removal is shown by (f). PSNR: (c) 42.20. ($\lambda = 0.5$)