Phase Preserving Tone Mapping of Non-Photographic High Dynamic Range Images

Peter Kovesi
Centre for Exploration Targeting
School of Earth and Environment
The University of Western Australia
Crawley, Western Australia 6009
Email: peter.kovesi@uwa.edu.au

***Abstract*—Non-photographic images having a high dynamic range, such as aeromagnetic images, are difficult to present in a manner that facilitates interpretation. Standard photographic high dynamic range (HDR) algorithms may be unsuitable, or inapplicable to such data. We present a method that compresses the dynamic range of an image while preserving local features. It makes no assumptions about the formation of the image, the feature types it contains, or its range of values. Thus, unlike algorithms designed for photographic images, this algorithm can be applied to a wide range of scientific images. The method is based on extracting local phase and amplitude values across the image using monogenic filters. The dynamic range of the image can then be reduced by applying a range reducing function to the amplitude values, for example taking the logarithm, and then reconstructing the image using the original phase values. An important attribute of this approach is that the local phase information is preserved, this is important for the human visual system in interpreting the image. The result is an image that retains the fidelity of its features within a greatly reduced dynamic range. An additional advantage of the method is that the range of spatial frequencies that are used to reconstruct the image can be chosen via high-pass filtering to control the scale of analysis.**

I. Introduction

Increasingly images are being captured, or created, with a high dynamic range that far exceeds what can be displayed on conventional devices. Much of the recent interest in high dynamic range images and the problem of tone mapping was stimulated by the work of Debevec and Malik [1] who first presented a practical algorithm that allowed high dynamic range radiance maps to be recovered from multiple images taken with conventional cameras. Prior to this the problem of high dynamic range images were encountered in computer graphics when realistic, physics based, illumination methods were developed for image rendering in the 1980s. In doing this the computer graphics community rediscovered the same problems that photographers encounter, for example, in trying to render an image of a room with a window opening out to a bright outdoor scene.

This paper is concerned with the rendering of high dynamic range (HDR) scientific images. The aims are potentially different from that in rendering high dynamic range photographic images. Photographic HDR maps may contain radiance values corresponding to moonlight ranging up to direct sunlight. They represent the formation of an image from light reflecting off objects. As such they are non-negative and have specific image statistics [2], [3]. Photographic HDR tone mapping algorithms have been specialised for these images and draw extensively on psychophysical research into the response of the human eye to light [4]. Ultimately their aim is to make an image 'look good'.

For non-photographic HDR scientific images we are not so concerned with making an image 'look good'. Instead our objective is to allow all features of interest to be revealed with maximum fidelity. Artifacts arising from the tone mapping process must also be minimised because our ability to recognise the presence of artifacts may be greatly limited. What may, or may not, be 'normal' in a non-photographic image will not be necessarily obvious. In addition we do not necessarily have a good sense of the scale of features of interest. Thus it is important that any tone mapping algorithm allows user control over the scale of features that are enhanced.

Scientific images may have image statistics and feature types that are very different from photographic images. The process by which scientific images may be formed can be very varied. For example contrast in magnetic resonance images arise from protons returning to equilibrium states at varying rates in different body tissues. Contrast in CT images arise from differing X-ray attenuations. Geophysical aeromagnetic images represent the magnetic susceptibility of rocks in the earth. Unlike photographic images the range of values in a scientific image may cross the origin, or the range of values may lie a considerable distance away from the origin. For example aeromagnetic images may contain values ranging from around 30,000nT (nanotesla) to 70,000nT, where an anomaly as small as 10nT may be of interest. Thus it may be quite inappropriate to apply tone mapping techniques developed for photographic images to scientific images, assuming that the algorithm can even accept the range of data that is in the input image.

This paper presents a new tone mapping algorithm that works in the frequency domain. It makes no assumptions about the formation of the image or its range of values. It ensures the fidelity of features are maintained by preserving the local phase of features. The scale of the features that are highlighted by the algorithm are controlled via high-pass filtering.

Traditionally scale has been considered in terms of low-pass filtering however under this approach the locations of features can vary with scale, this is not the case under high-pass filtering [5]. Under high-pass filtering the relative magnitudes of features can vary but not their locations.

II. Prior Work

As mentioned earlier, work on tone mapping has primarily been concerned with HDR photographic images and rendering of computer graphics scenes. The techniques can be roughly divided in a number of broad categories though there can be considerable overlap in the methods.

Histogram modification:
Ward et al. [6] developed a method based on modifying a luminance histogram. Rather than trying to achieve a uniform histogram as is done with histogram equalisation they devise a histogram *adjustment* algorithm that limits the maximum local contrast that would otherwise be obtained via histogram equalisation. This prevents the formation of areas having unnaturally high local contrast. The overall aim is to preserve perceived contrast. Qiu et al. [7] cast the problem of finding an appropriate histogram mapping function as an optimisation problem. They devise an objective function that seeks to optimally divide the image histogram into local sections that are either scaled linearly, or histogram equalised.

Models of visual adaptation:
Early work includes that of Ferwerda et al. [8] who developed a model of visual adaptation based on psychophysical data. Light and dark areas of an image are modified according to a model of local adaptation. Reinhard et al. [9] develop a local dodging-and-burning operator where the scale of the required dodging-and-burning is determined locally. Krawczyk et al. [10] build a model based on lightness perception. The image is decomposed into frameworks/regions of common illumination. An anchor that provides a mapping between luminance and the perceived grey shade is computed for each framework. From this an appropriate tone mapping for each framework is derived. Johnson and Fairchild [11], [12] develop an image appearance model, iCAM, that models chromatic and luminance adaptation processes and also models local contrast/surround effects. This was one of the best ranked algorithms in Ledda et al.'s 2005 evaluation of tone mapping algorithms [13].

Decomposition into layers:
The approach of decomposing an image into two layers, a base layer of primarily low frequency information, and a detail layer is a common approach. The magnitude of the base layer is reduced in some way before being recombined with the detail layer. This approach corresponds to a simple functional model of visual adaptation. An early example is homomorphic filtering [14] where it is assumed that the low and high frequency components of an image correspond to its illumination and reflectance components respectively. The amplitude spectrum of the illumination component is reduced and, optionally, the reflectance component is amplified before the image is reconstructed. Ashikhmin [15] attempts to build a simple functional model of the human visual system. Local adaptation levels are computed across an image to which a tone mapping function is applied. Image detail information is then reapplied to the adapted image. Ashikhmin notes that some care is needed to avoid introducing false features in the adaptation image. Bilateral filters [16] are used by Durand and Dorsey [17] as edge-preserving filters to decompose an image into a low frequency base layer and a detail layer. The base layer then has its contrast reduced and the detail layer is then recombined. Note that this approach implicitly assumes that that only step edges are the features of interest that should be preserved in the smoothing process.

Gradient attenuation:
The gradient domain approach of Fattal et al [18] attenuates the magnitudes of large gradients in the luminance gradient field and then reconstructs the low dynamic range image by integrating the modified gradient field by solving a Poisson equation. Gradients at all scales are computed via a Gaussian pyramid and the gradient attenuation coefficients are propagated through to the finest scale to determine the overall gradient attenuation function. Mantiuk et al. [19] extend the gradient approach by imposing constraints on the contrasts over the whole image. This avoids the possibility of reversing polarity of contrast in the gradient attenuation process.

Non-photographic algorithms:
For geological and geophysical data, for which the algorithm developed in this paper is directed, high dynamic range images are often dealt with via histogram equalisation. The equalised image is then typically displayed via a rainbow colour map. This can lead to perceptual distortions due to the varying contrast adjustments that are applied throughout the data range in the equalised data, perceptual biases associated with different colours, and non-uniform perceptual contrasts between colours in the colour map [20].

For aeromagnetic data the Automatic Gain Control algorithm developed by Rajagopalan [21], [22] is sometimes used. This algorithm attempts to generate an image where the waveforms have constant amplitude. Within a local window, that is moved across the image, the input signal amplitude is estimated, say via the root mean square, and a gain value applied to normalise this amplitude. A difficulty with this algorithm is that the signal can be over-normalised and perception of the overall variations in the signal can be lost. However, the algorithm does allow the scale over which the signal normalisation is applied to be controlled by varying the size of the analysis window.

The tilt derivative given by

$$\tan^{-1}\left(\frac{\text{vertical component of the signal gradient}}{|\,\text{horizontal component of the signal gradient}\,|}\right)$$

is another normalised representation of magnetic data that can be used [23]. While the tilt derivative has some useful properties it too suffers from the problem of over-normalising of the data.

III. Phase Preserving Tone Mapping

In general, tone mapping algorithms devised for photographic images have been developed on the basis of statistics of natural images, feature types found in natural images, and models of the human visual system. The algorithms may require image segmentation and/or decomposition of the image into base and detail layers. Often there are many parameters that have to be set.

For non-photographic images we want to avoid making particular assumptions about the image type and the features that may be present. Even in natural images most feature types are not simple step edges [24]. We also want to avoid the need for segmentation or decomposition into base and detail layers because such operations may have no real meaning with respect to the image. The number of parameters that need to be set should be minimised, and where they do exist they should have a clear physical meaning.

The primary assumption that is made in the new algorithm is that the local phase values of the image must be preserved in the tone mapped output. Phase is important to the human visual system. Oppenheim and Lim [25] show that the amplitude spectrum of an image can be modified considerably, even swapped with that from another image, and the features from the original image will still be seen clearly as long as the phase information is preserved.

The approach adopted is to decompose the image into its local phase and amplitude values. The amplitude values are attenuated via some function and the image then reconstructed using the original phase values and the attenuated amplitude values. A point to emphasise here is that the *local* phase and amplitude values are used and manipulated, not the global ones that you would obtain via a Fourier transform.

To obtain the local phase and amplitude monogenic filters [26], [27] are used. Monogenic filters are formed by combining a radial band-pass or high-pass filter with its Riesz transform. The Riesz transform forms a 2D equivalent of the Hilbert transform. It is made up of two components. If we define two filters in the 2D frequency domain u_1, u_2

$$H_1 = i\frac{u_1}{\sqrt{u_1^2 + u_2^2}} \qquad H_2 = i\frac{u_2}{\sqrt{u_1^2 + u_2^2}}$$

then the spatial representation of the vector $\mathbf{H} = (H_1, H_2)$ defines the convolution kernel of the Reisz transform. These two filters represent quadrature phase shifting operations in the two orthogonal directions of the image. To obtain local phase and amplitude information the image, I, is convolved with the band-pass or high-pass filter f and the two Reisz transform filtered versions of f, $h_1 f$ and $h_2 f$. This provides three outputs, $I * f$, $I * h_1 f$ and $I * h_2 f$, where $*$ denotes convolution.

If some overloading of notation can be permitted, for brevity, these convolution results will also be referred to as $f, h_1 f$ and $h_2 f$. Figure 1 illustrates the process. At a point in the image the three convolution outputs can be thought of as forming a vector in 3-space. The output from convolution with the band-pass filter f corresponds to the vertical coordinate, and the convolutions with the Reisz transform filters $h_1 f$ and $h_2 f$ specify the two horizontal coordinates. The vertical axis can be thought of the real component of the signal and the two horizontal axes represent the two complex valued, phase shifted versions of the signal in the two orthogonal image axis directions.

The local amplitude at image location (x, y) is given by

$$A(x,y) = \sqrt{f(x,y)^2 + h_1 f(x,y)^2 + h_2 f(x,y)^2}\ .$$

The local phase is given by

$$\phi(x,y) = \mathrm{a}tan2(f(x,y), \sqrt{h_1 f(x,y)^2 + h_2 f(x,y)^2})$$

and orientation given by

$$\theta(x,y) = \mathrm{a}tan2(h_2 f(x,y), h_1 f(x,y))\ .$$

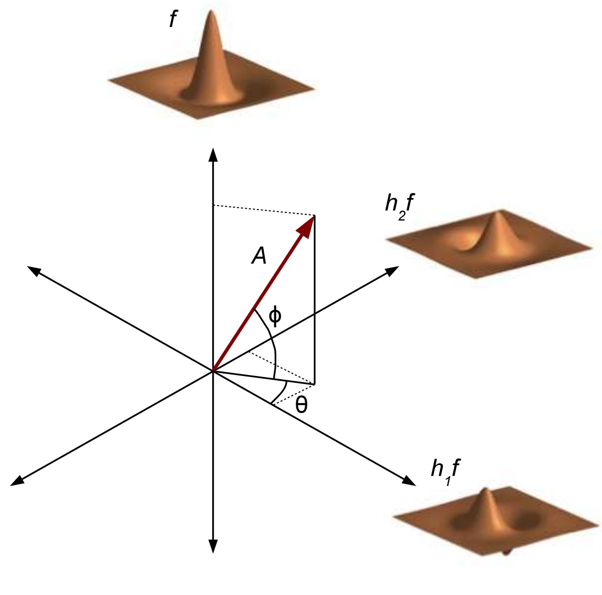

Fig. 1. Obtaining phase and amplitude from the outputs of monogenic filters. The band-pass filter is defined by f, and $h_1 f$ and $h_2 f$ are the two Reisz transform filters. A is the local amplitude, ϕ the local phase, and θ the local orientation.

Dynamic range reduction of the image is simply achieved by applying a range reducing function to the amplitude and then reconstructing using the original phase. Referring to Figure 1 this can be thought of reducing the length of the vector in 3-space while maintaining its direction, and then projecting it back onto the vertical, real, axis. In the work presented here the amplitude range reduction has been achieved using the logarithm of the amplitude, $\log(A + 1)$ or, in some cases, a nested logarithm ($\log(\log(A+1)+1)$) of the amplitude. Note that 1 is added to the amplitude values to avoid reversal of the signal for values less than 1.

The reconstructed, tone mapped image image values, $T(x, y)$ are given by

$$T(x, y) = \log(A(x, y) + 1). \sin(\phi(x, y))$$

IV. Scale variation and signal attenuation via high-pass filtering

Apart from the choice of amplitude range reduction function the main parameter in the algorithm is the choice of filter f. It must be a high-pass or band-pass filter because the Reisz transform, like the Hilbert transform, is not defined for a DC signal.

In this work we have found the use of high-pass filters most useful. For fine details of the image to be preserved it is important to retain all the high frequency components of the signal. Progressively attenuating the low frequency components of the signal achieves two things, the scale of the image is varied, and the dynamic range of the image is further reduced. As low frequency components are removed small scale features that would otherwise be swamped by the broader scale features are revealed. Unlike the situation if low pass filtering was to be used the locations of these small scale features remains stable under different levels of high-pass filtering, all that changes is their magnitude relative to features at broader scales. Clearly this is a desirable attribute of any analysis that varies over scale [5]. In addition, as mentioned earlier, the other thing achieved by high-pass filtering is that by removing low frequency components of the signal, which are typically large, the dynamic range of the image is further reduced.

For scientific images the scale of interest may be unknown beforehand. Indeed, the image may contain features of interest over a variety of unknown scales, as is typically the case with geophysical images. We have found it useful to generate a sequence of about ten images with the high-pass filter cutoff frequency increasing geometrically. The smallest cutoff spatial frequency in the series might be $1/w$, where w is the image size, and the largest might be, say, $30/w$. These images can be viewed individually, however, what we have found most effective is to use an image blender that allows one to interactively blend across the whole sequence of images. This allows one to simulate a continuum of processing scales. The intermediate blends of the images act as an approximations to the computational result that would have been obtained had the image been processed at that exact intermediate scale. Rapid interactive exploration of the image over multiple scales can be extremely instructive.

V. Results

Figure 2 shows a raw aeromagnetic image of the Yilgarn area in Western Australia along with its histogram. The raw data values range from about 50,000nT to about 70,000nT. Note how the histogram count values range over several orders of magnitude. In this figure the image values have been linearly scaled and shifted to a range of 0-255. Figure 3 shows a histogram equalised version of the image. Figures 4 to 6 show the results of the dynamic range compression algorithm with high-pass cutoff frequencies of 1/2000, 1/1000 and 1/200 respectively. The image size is 2492 x 2847 thus at the smallest cutoff frequency almost all the spatial frequencies in the image are fully represented. Note the progressive revealing of fine scale features as the cutoff frequency is increased. The inevitable halo artifacts also increase as cutoff frequency is increased. It is important to use a low order butterworth high-pass filter to minimise these effects. Interactively scanning through image scales via an image blender can also be useful in helping distinguish any confusing halo artifacts from genuine features. For comparison the Yilgarn data was also rendered using using the Mac OSX Preview application, which can render HDR files, this is shown in Figure 7. It was found that offsetting the image values so that the minimum value was 0, rather than 5×10^4 was required to obtain a reasonable result with the Preview application. Figure 8 shows the result obtained using MATLAB's HDR `tonemap` function [28]. This function is based on Ward's contrast limited histogram equalisation algorithm [6]. Neither of these alternate renderings show the fine structures that can be revealed by varying the high-pass cutoff frequency of the proposed algorithm.

While it is not suggested that the proposed algorithm is suited for photographic images it can be used on them. Figure 9 is an image with strong shadows and Figure 10 shows the output of the algorithm using a high-pass cutoff frequency of 1/200. The result may not be as aesthetically pleasing as what might be obtained using an algorithm optimised for photographic images, but is acceptable.

Out of interest the algorithm was also tested on Debevec's widely used radiance data of the Stanford Memorial Church [29]. For this image it was found that using a nested logarithmic attenuation of the local amplitudes of the radiance values was most appropriate. In addition, the logarithm of the radiance data was used as an approximation of brightness and used as input to the the dynamic range compression algorithm. The final images were also adjusted using a gamma value of 0.5. For comparison a rendering obtained using the Mac OSX Preview application is also shown. No detailed comparison can, or should, be made as with any HDR rendering package there are many parameters that can be adjusted. The interesting thing to note is that reasonable results can be obtained using the proposed algorithm which primarily has only one parameter, the cutoff frequency of the high-pass filter. The algorithm's primary aim is to make features visible, not to make them attractive. The results have been deliberately shown in grey scale to allow the tone mapping to be seen as distinct from colour reproduction issues which are not the topic of this paper.

VI. Conclusion

This paper presents a new tone mapping algorithm designed for non-photographic images. The algorithm works in the frequency domain and ensures the fidelity of features are maintained by preserving the local phase values within the image data. It makes no assumptions about the formation of the image, the feature types it contains, or its range of values. Thus, unlike algorithms designed for photographic images, this algorithm can be applied to a wide range of scientific images. No image segmentation, or decomposition into base and detail layers is needed. The algorithm only has one major parameter, the cutoff frequency of the high-pass filter used. This parameter controls the scale of the features that are highlighted by the algorithm. Preparation of a sequence of images with geometrically scaled high-pass cutoff values allows exploration of the features over a wide range of scales within the image. MATLAB code is available for those wishing to replicate the results presented here [30].

Acknowledgement

The author would like to thank Fugro Airborne Surveys Pty Ltd. for permission to use their aeromagnetic data.

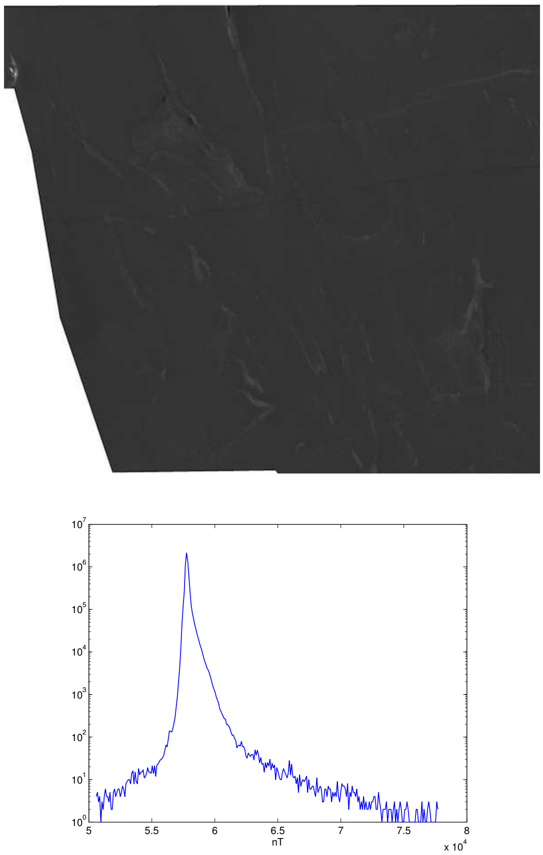

Fig. 2. Raw aeromagnetic data of the Yilgarn and logarithmic plot of its histogram values. This data is the property of Fugro Airborne Surveys Pty Ltd.

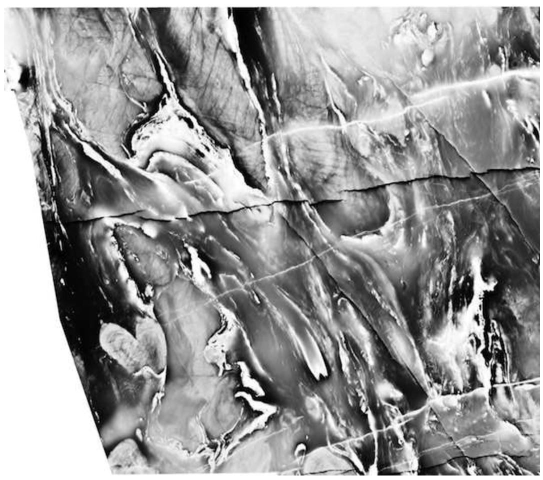

Fig. 3. Histogram equalised aeromagnetic image of the Yilgarn.

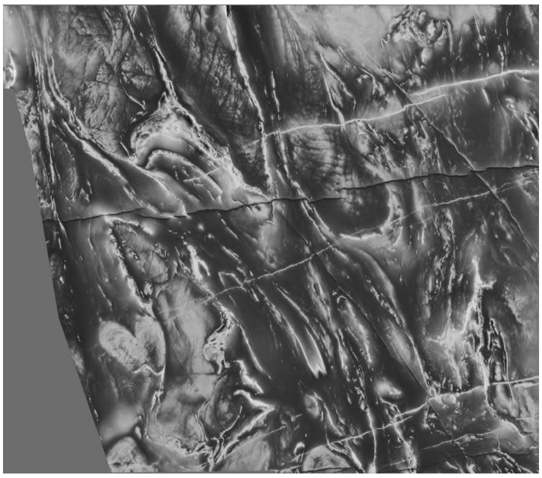

Fig. 4. Rendering of the Yilgarn aeromagnetic data using the proposed tone mapping algorithm. High-pass cutoff frequency 1/2000.

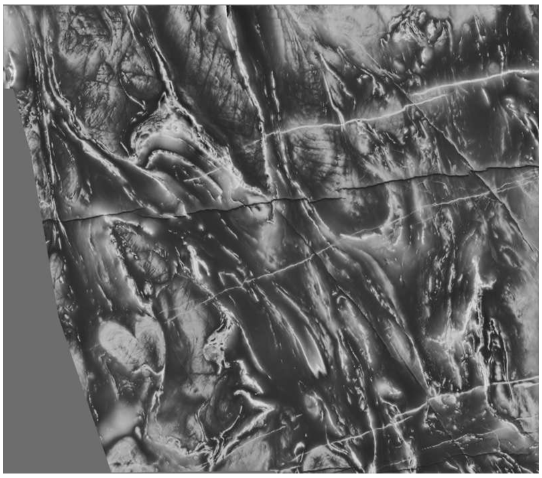

Fig. 5. Rendering of the Yilgarn aeromagnetic data using the proposed tone mapping algorithm. High-pass cutoff frequency 1/1000.

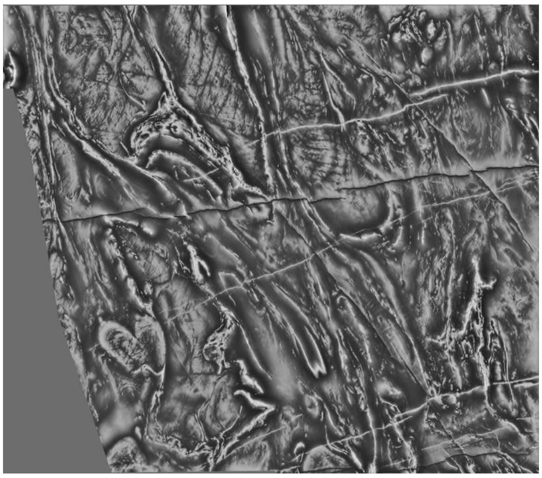

Fig. 6. Rendering of the Yilgarn aeromagnetic data using the proposed tone mapping algorithm. High-pass cutoff frequency 1/200.

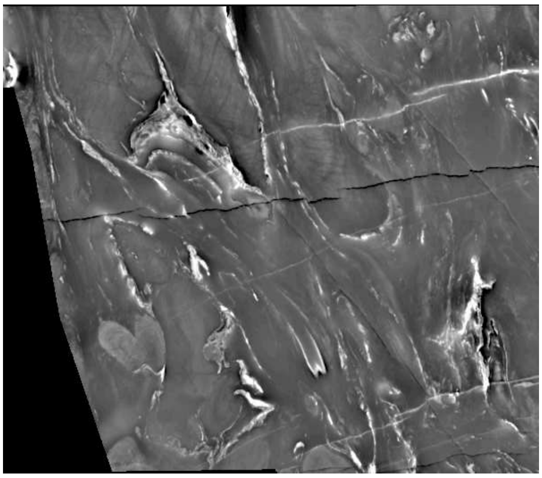

Fig. 7. Yilgarn HDR data rendered with Mac OSX Preview application.

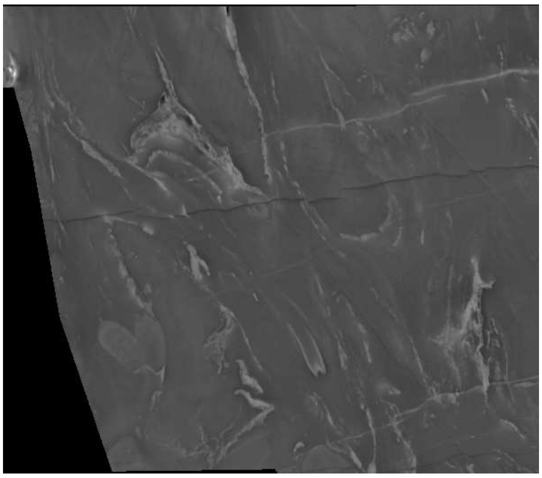

Fig. 8. Yilgarn data rendered with MATLAB's HDR `tonemap` function.

REFERENCES

[1] P. E. Debevec and J. Malik, "Recovering high dynamic range radiance maps from photographs," in *SIGGRAPH 97*, 1997.

[2] D. J. Field, "Relations between the statistics of natural images and the response properties of cortical cells," *Journal of The Optical Society of America A*, vol. 4, no. 12, pp. 2379–2394, December 1987.

[3] ——, "What the statistics of natural images tell us about visual coding," in *SPIE 1077*, 1989, pp. 269–276, los Angeles, California.

[4] E. Reinhard, G. Ward, S. Pattanaik, P. Debevec, W. Heidrich, and K. Myszkowski, *High Dynamic Range Imaging Acquisition, Display, and Image-Based Lighting*. Morgan Kaufmann, 2010.

[5] P. Kovesi, "Image features from phase congruency," *Videre: A Journal of Computer Vision Research*, vol. 1, no. 3, 1999.

[6] G. J. Ward, H. Rushmeier, and C. Piatko, "A visibility matching tone reproduction operator for high dynamic range scenes," *IEEE Transactions on Visualization and Computer Graphics*, vol. 3, no. 4, December 1997.

[7] G. Qiu, J. Guan, J. Duan, and M. Chen, "Tone mapping for HDR image using optimization - a new closed form solution," in *Proc. ICPR 2006, 18th International Conference on Pattern Recognition*, vol. 1, 2006, pp. 996–999.

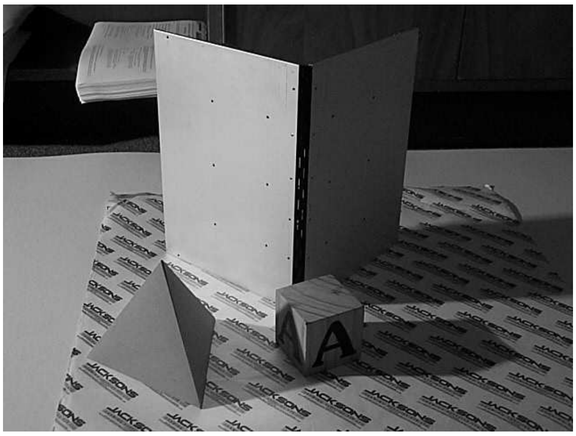

Fig. 9. Image with strong shadows.

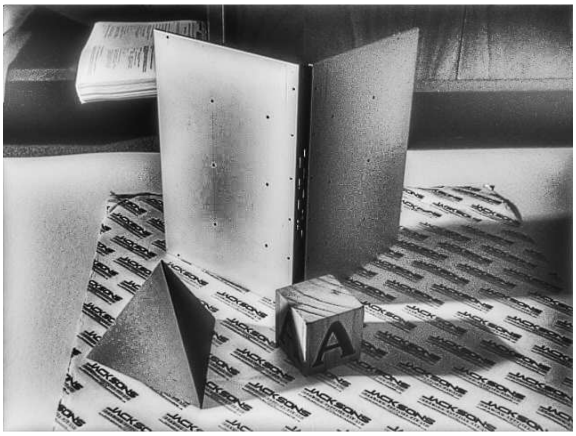

Fig. 10. Image rendered with the proposed tone mapping algorithm. High-pass cutoff frequency 1/200.

[8] J. A. Ferwerda, S. N. Pattanaik, P. Shirley, and D. P. Greenberg, “A model of visual adaptation for realistic image synthesis,” in *SIGGRAPH 96*, 1996.

[9] E. Reinhard, M. Stark, P. Shirley, and J. Ferwerda, “Photographic tone reproduction for digital images,” in *ACM Transactions on Graphics, 21(3), July 2002 (Proceedings of SIGGRAPH 2002)*, 2002.

[10] G. Krawczyk, K. Myszkowski, and H.-P. Seidel, “Computational model of lightness perception in high dynamic range imaging,” in *Human Vision and Electronic Imaging XI, IS&T/SPIE’s 18th Annual Symposium on Electronic Imaging (2006)*, B. E. Rogowitz, T. N. Pappas, and S. J. Daly, Eds., 2006.

[11] G. Johnson and M. Fairchild, “Rendering HDR images,” in *Proceedings of IS & T/SID 11th Color Imaging Conference*, 2003, pp. 36–41.

[12] J. Kuang, G. ohnson, and M. Fairchild, “iCAM06:a refined image appearance model for HDR image rendering,” *Journal of Visual Communication and Image Representation*, vol. 18, no. 5, pp. 406–414, 2007.

[13] P. Ledda, A. Chalmers, T. Troscianko, and H. Seetzen, “Evaluation of tone mapping operators using a high dynamic range display,” in *ACM SIGGRAPH 2005*. ACM, New York, 2005, pp. 640–648.

[14] J. T. G. Stockham, “Image processing in the context of a visual model,” *Proceedings of the IEEE*, vol. 60, pp. 828–842, 1972.

[15] M. Ashikhmin, “A tone mapping algorithm for high contrast images,” in *Eurographics Workshop on Rendering*, P. Debevec and S. Gibson, Eds., 2002, pp. 1–11.

Fig. 11. Stanford Memorial Church rendered with the proposed tone mapping algorithm. High-pass cutoff frequency 1/100. With this cutoff frequency halo effects are quite evident. Radiance data from Debevec [29].

Fig. 12. Stanford Memorial Church rendered with the proposed tone mapping algorithm. High-pass cutoff frequency 1/200.

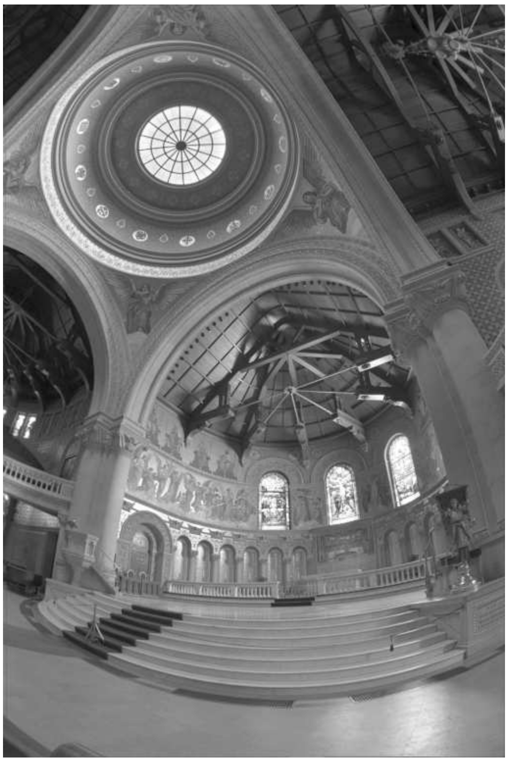

Fig. 13. Stanford Memorial Church rendered using the Mac OSX Preview application.

[16] C. Tomasi and R. Manduchi, "Bilateral filtering for gray and color images," in *Proc. IEEE Int. Conf. on Computer Vision*, 1998, pp. 836–846.

[17] F. Durand and J. Dorsey, "Fast bilateral filtering for the display of high-dynamic-range images," in *ACM SIGGRAPH 2002*, 2002.

[18] R. F. Dani, Lischinski, and M. Werman, "Gradient domain high dynamic range compression," in *ACM SIGGRAPH 2002*, 2002.

[19] R. Mantiuk, K. Myszkowski, and H.-P. Seidel, "A perceptual framework for contrast processing of high dynamic range images," *ACM Transactions on Applied Perception 3, 3*, pp. 286–308, 2006.

[20] D. Borland and R. M. T. II, "Rainbow color map (still) considered harmful," *IEEE Computer Graphics and Applications*, pp. 14–17, March/April 2007.

[21] S. Rajagopalan, "The use of 'automatic gain control' to display vertical magnetic gradient data," in *5th ASEG Conference*, 1987, pp. 166–169.

[22] S. Rajagopalan and P. Milligan, "Image enhancement of aeromagnetic data using automatic gain control," *Exploration Geophysics*, no. 25, pp. 173–178, 1995.

[23] A. Salem, S. Williams, D. Fairhead, R. Smith, and D. Ravat, "Interpretation of magnetic data using tilt-angle derivatives," *Geophysics*, vol. 73, no. 1, pp. 1–10, 2008.

[24] P. Kovesi, "Edges are not just steps," in *Proceedings of ACCV2002 The Fifth Asian Conference on Computer Vision*, 2002, pp. 822–827.

[25] A. V. Oppenheim and J. S. Lim, "The importance of phase in signals," in *Proceedings of The IEEE 69*, 1981, pp. 529–541.

[26] M. Felsberg and G. Sommer, "A new extension of linear signal processing for estimating local properties and detecting features," in *DAGM Symposium, Kiel*, 2000.

[27] ——, "The monogenic signal," *IEEE Transactions on Signal Processing*, vol. 49, no. 12, pp. 3136–3144, December 2001.

[28] "MathWorks MATLAB Image Processing Toolbox function `tonemap`." [Online]. Available: http://www.mathworks.com.au/help/toolbox/images/ref/tonemap.html

[29] P. Debevec, "Recovering high dynamic range radiance maps from photographs." [Online]. Available: http://www.pauldebevec.com/Research/HDR/#radiancemaps

[30] P. Kovesi, "MATLAB and Octave functions for computer vision and image processing." [Online]. Available: http://www.csse.uwa.edu.au/~pk/research/matlabfns/