Lookup-Table Based Gradient Field Reconstruction

Graham D. Finlayson, *Member, IEEE*, David Connah,
and Mark S. Drew, *Member, IEEE*

Abstract

In computer vision there are many applications where it is advantageous to process an image in the gradient domain and then re-integrate the gradient field: important examples include shadow removal, lightness calculation and data fusion. A serious problem with this approach is that the reconstruction step often introduces artefacts – commonly, smoothed and smeared edges – to the recovered image. This is a result of the inherent ill-posedness of re-integrating a non-integrable field. Artefacts can be diminished, but not removed, by using complex to highly complex re-integration techniques.

Here we present a remarkably simple (and on the face of it naive) algorithm for reconstructing gradient fields. Suppose we start with a multichannel original, and from it derive a (possibly one of many) 1D gradient field; for many applications the derived gradient field will be non-integrable. Here we propose a lookup-table based map relating the multichannel original to a reconstructed scalar output image whose gradient best matches the target gradient field. The idea, at base, is that if we learn how to map the gradients of the multichannel original onto the desired output gradient, then using the Lookup-Table (LUT) constraint we effectively derive the mapping from the multichannel input to the desired, re-integrated, image output. While this map could take a variety of forms, here we derive the best map from the multichannel gradient as a (nonlinear) function of the input to each of the target scalar gradients. In this framework, reconstruction is a simple equation-solving exercise of low dimensionality.

One obvious application of our method is to the image-fusion problem, e.g. the problem of converting a colour or higher-D image into greyscale. We will show, through extensive experiments and

Copyright (c) 2010 IEEE. Personal use of this material is permitted. However, permission to use this material for any other purposes must be obtained from the IEEE by sending a request to pubs-permissions@ieee.org.

G.D. Finlayson and D. Connah are with the School of Computing Sciences, The University of East Anglia, Norwich, England NR4 7TJ, e-mail: {graham/dc}@cmp.uea.ac.uk.

M.S. Drew is with the School of Computing Science, Simon Fraser University, Vancouver, B.C. Canada V5A 1S6, e-mail: mark@cs.sfu.ca

complementary theoretical arguments, that our straightforward method preserves the target contrast as well as do complex previous re-integration methods *but without artefacts*, and with a substantially cheaper computational cost. Finally, we demonstrate the generality of the method by applying it to gradient field reconstruction in an additional area, the shading recovery problem.¹

Index Terms

Gradient, re-integration, contrast, lookup table, image fusion, shape-from-shading

I. INTRODUCTION

There are many problems in computer vision whose solution essentially follows a three-step work flow. First, gradient fields are derived from input image(s); second, gradients are manipulated (to impose a desirable property) and fused; and then, third, a new image is found by reconstructing the derived gradient field. For example, in lightness computation small gradients (assumed to be caused by illumination) are set to zero; the result, after re-integration, is an image free of illumination shading [1]. Similarly, in shadow removal and intrinsic image calculation, shadow edges and non-material boundaries are identified and either set to zero or inpainted prior to re-integration [2], [3]. In shape from shading and photometric stereo the gradient field (surface orientation) at each point is found and then the depth map results from reconstructing the gradient field [4], [5]. Other applications where gradient reconstruction appears include image compositing, noise removal and image fusion and enhancement [6], [7], [8].

In this paper as a first exemplar application we consider the gradient domain data fusion method of Socolinsky and Wolff (SW) [9]. In that method a single output gradient vector is calculated for each pixel (from the multichannel structure tensor) and this is re-integrated to give a greyscale scalar image that captures the contrast of the multichannel original. However, the gradient field is typically non-integrable and when the image field is solved for using a Poisson solver or more recent robust techniques [10], [11], serious image artefacts remain, often characterised by smearing of edges or bending in flat image regions.

Fig. 1 shows an example where re-integrating a non-integrable gradient with existing methods results in large visible artefacts. The image in Fig. 1A has non-negligible noise levels, which results in local indeterminacy in the calculation of the one-dimensional gradient vector. The least-squares based Poisson solver output [12] in Fig. 1B cannot overcome the resulting variability, and large bending artefacts result.

¹The contents of this paper constitute in part Patent Filing GB0914603.6, “Image Reconstruction Method”, Date Lodged: 20 August 2009.

Fig. 1. A colour image with a strongly non-integrable contrast, and the output of three different re-integration methods.

Fig. 1C shows the output from a recent technique due to Reddy et al. [11], which was designed for the shape-from-shading problem. This method, following on from earlier work in [13], [10] assumes that local integrability errors are caused by underlying image structure and not random noise, an assumption that clearly fails for this particular image. The latest method of Reddy *et al.* attempts to mitigate these problems by minimising the L_1 error, rather than the L_2 error and hence making the output more robust to noisy outliers. Nonetheless, when we apply the method to this difficult problem we can see that the output offers little improvement over the least-squares Poisson solver. This result, that minimising the L_1 norm generally offers little improvement over the L_2 norm for the image-fusion problem in the images we have experimented with, was also reported in [14]; nonetheless it was shown to have demonstrable advantages when solving the shape-from-shading problems considered in [11].

Finally we present an output for our new method, shown in Fig. 1D, which clearly outperforms previous approaches, eliminating re-integration artefacts.

For the new, artefact-free re-integration method presented in this paper we investigate the theoretical motivation of the method by deriving a theorem, presented in §3. This theorem recognises that if we are given multiple, integrable, gradient fields, where each is derived at a different spatial scale (and the

full range of scales are included), then if we reintegrate to deliver a coherent image (i.e. one whose multi-scale gradient field matches the given gradient fields exactly), that image must be formed by an LUT mapping of the original image. In practice, a gradient field may only be provided at a single scale, or if a multi-scale gradient is provided it will rarely be integrable at all scales. Nonetheless, the use of an LUT mapping is both a novel and, it turns out, surprisingly effective algorithm for gradient field reconstruction. The algorithm, described in detail §3, employs an LUT mapping from the range of values for the original image (e.g., colour-space) to the range (here, for illustration, a single greyscale) of the output image; i.e. a transform that maps colour pixel values from the original image to a single greyscale in the output via a lookup-table approach (although these greyscale values need not be unique). This mapping is chosen to minimise the difference between the gradient of the final output image and the gradient derived by SW's method. While the implementation could take different forms, the embodiment here generates a polynomial expansion of the original colour image, and maps this to greyscale. In this form, the re-integration becomes a simple equation solving exercise, which is both simpler and more effective than antecedent re-integration methods.

Our gradient field reconstruction formulation is quite general and can be applied to other reconstruction problems so long as that we have access to the multichannel original as well as the putative output greyscale gradient field. As well as for colour-to-grey, we additionally present a re-integration example from shape-from-shading, and we will demonstrate that the lookup-table solution provides an output that compares favourably with other methods, and is both computationally fast and simple to implement.

In §2 of this paper we provide a short review of the gradient reconstruction problem in general, and go on to discuss how different problems can be solved via lookup-table mapping. The theoretical justification for the method is given in §3, along with guidelines for our preferred implementation. In §4 we demonstrate results for each application, with the main focus placed on the problem of converting colour to greyscale. For this data fusion problem, we also evaluate our algorithm in a preference experiment with human observers. Here, an observer must choose which output from a variety of algorithms is preferred. Not only is our lookup-table method shown to be preferred compared with the original SW method but it is also preferred compared with other colour to greyscale algorithms. The paper concludes in §5.

II. BACKGROUND

The problems associated with a least-squares approach to re-integrating a, usually non-integrable, gradient field are well known. Indeed, finding a re-integrated image which when differentiated is closest in a least-squares sense to the original gradient field is only one of many optimisations. In [15], shadow free

images are found by re-integrating a gradient field where shadow edges are set to 0. In that paper, an image was represented as a grid-graph and re-integration was formulated as finding an *appropriate* Hamiltonian path along which an image is re-integrated as if it were a 1D function. There, appropriateness was determined by a cost function which measured the likelihood of re-integration errors. More specifically, a means was presented for determining whether the same surface reflectance was present across a shadow edge (this is the assumption when setting shadow edges to 0). When this condition is met the gradient field is integrable. Their method had the disadvantage that some pixel locations were not visited by the re-integration path and as a result there were gaps in the re-integrated image. In-filling the gaps delivered reasonable results.

More generally, this idea of determining which parts of images are integrable and which are not is explored by Agrawal et al. [13]. Non-integrability is identified as gradient field locations with an above-threshold curl. They too view an image as a grid graph, where connectivity is a function of the curl present at an image location. Edges omitted because of non-integrability are estimated from those that are integrable. Once an integrable field has been recovered, a traditional Poisson solver can be used to re-integrate the gradient field. In the most recent incarnation, by Reddy *et al.*, the technique is further enhanced by increasing its robustness to outliers by minimising the L_1 norm error. Nonetheless, as shown in the output in Fig. 1C for the data-fusion problem, while this can correct some artefacts where the underlying image structure is contained in the non-negative portion of the curl, there are some artefacts for which the method is not effective.

A. Gradient domain data fusion

There has been recent interest in the problem of converting colour (or, in general multi-dimensional) image data into greyscale [16], [9], [17], [18], [19], [20], [21]. The basic goal is to preserve in the greyscale output the colour-contrast present in the original image. The most well known method is that of Gooch et al. [17], who use a gradient descent method to minimise the difference between the contrast in the output image and the original colour-contrast. They also employ three user inputs to define the neighbourhood over which colour calculations are performed, to resolve sign ambiguities for isoluminant colours with different hues, and also to define how much chroma contrast is retained in the output image. Grundland et al. [19], Rasche et al. [20], and Kim et al. [21] employ similar optimisation methods, but rather than maintaining local contrasts they take entirely global approaches to the optimisation. Finlayson and Qiu [22] propose a related method where they define contrast globally using the covariance matrix of the image colours: this allows them to use a simple, closed form, optimisation to retrieve a variance-

preserving greyscale.

One of the best performing (and mathematically elegant) approaches is that proposed by SW [9]. Their method begins by defining the colour-contrast as the Di Zenzo [23] colour gradient. Let $\rho(a, b)$ denote an image with vector values at each location (a, b) . The gradient in the x and y directions are vector quantities denoted $\nabla_x(a, b)$ and $\nabla_y(a, b)$. Suppose we wish to determine the vector gradient in the direction $\mathbf{w} = [u \ v]^t$ ($\mathbf{w}$ has unit length). If $\mathbf{M} = [\nabla_x(a, b) \ \nabla_y(a, b)]$, then $\nabla_{\mathbf{w}} = \mathbf{M} \mathbf{w}$. The matrix $\mathbf{D}(a, b) = \mathbf{M}^t \mathbf{M}$ is called the structure tensor and captures the relationship between directionality and magnitude of the vector field: $\mathbf{w}^t \mathbf{D} \mathbf{w}$ is the squared magnitude of the gradient in direction $\mathbf{w}$. To a first *approximation* the colour tensor can be represented as a single unsigned gradient vector (x and y derivative) at each pixel – its magnitude is the square root of the largest eigenvalue of $\mathbf{D}$ and its orientation is given by the corresponding eigenvector. The sign can be determined by the gradients of brightness in the original image (where the sign is indeed uniquely determined). Because the gradient field here is a **Tensor derivative approximation**, we write the gradient derived from the eigenvectors of $\mathbf{D}$ as ∇J :

$$\nabla J = Tda(\rho(a, b), \rho(a + 1, b), \rho(a, b + 1)) = Tda(\rho) \quad (1)$$

Note that the gradient in each channel is approximated using a local finite-difference operator; hence the gradient at pixel (a, b) is dependent upon its lower and right neighbours only.

The resulting gradient field, which is typically non-integrable, can be re-integrated to minimize squared error, to give a colour contrast preserving greyscale encoding G [16]. That is, the scalar image G is sought which minimizes the following in the least-squares sense:

$$\min_G \|\nabla G - \nabla J\| \quad (2)$$

Eq. (2) can be solved by assuming the Poisson equation $\nabla^2 G = \text{div}(\nabla J)$, or by other minimisation methods where $\|\cdot\|$ represents a norm other than L_2 .

In Fig. 2 we show a simple but challenging input colour image, 2A, and the resulting greyscale 2B recovered by the SW method. Notice the large amount of smearing and “bending” around the edges. This is due to the non-integrability of the gradient field and the least-squares recovery. We point out to the reader that in Socolinsky’s thesis [14] he discusses a reconstruction technique that maps the level sets in the luminance image to the level sets in the output. Significantly, the amount of smearing in the output is reduced (see Fig. 2C). However, the method requires a degree of quantisation of the input and is laborious to compute for complex images.

In the lookup-table approach we also aim to solve the minimisation in Eq. (2). But in our approach we constrain the scalar image G to be a lookup-table mapping from the multichannel original. To capture this idea, we can write:

$$\min_f ||\nabla f(\boldsymbol{\rho}) - \nabla J|| \quad (3)$$

where $f(\cdot)$ denotes the lookup-table map (our specific implementation of which is given in § 3). Practically, this constraint means that any given input colour (or multidimensional vector) will always be mapped to the same output greyscale value. The main advantage is in the prevention of “halo”, “ringing”, or “bending” artefacts such as that seen in Figure 2B: in general, such artefacts are most conspicuous in regions that are homogeneous (or close to homogeneous) in the original image. When using an LUT mapping, any homogeneous region in the original image is always mapped to a homogeneous region in the output, thus preventing such spatial artefacts from occurring – as seen in Figure 2D. Furthermore, provided that the LUT mapping is also smooth, regions that are almost homogeneous in the original will also be mapped to almost homogenous outputs (i.e. with no additional artefacts). We also note from Fig. 2D that the output is clearly comparable to the level sets implementation of SW in this case but, as we will demonstrate, our method is much simpler to implement and can readily be scaled to more complex images.

A discussion of the colour-to-greyscale problem should include other recent variations on SW’s method. Neumann et al. calculate local colour contrast based upon distances in the coloroid colour space, and use a least-squares method to re-integrate the field [24]. Alsam and Drew [25] replace the Di Zenzo tensor with the maximum gradient direction and re-integrate using an iterative solver with gamut constraints. While these methods produce good outputs for a number of images, neither addresses the artefacts associated with the least-squares re-integration method.

We also point out that there are, of course, other colour-to-greyscale mappings that involve a lookup-table transform. The most basic, $(R+G+B)/3$, is clearly a lookup-table map, and other methods, notably those of Gooch et al.[17], Rasche et al. [20], and most recently Kim et al. [21], employ lookup-table strategies. These methods differ conceptually from our approach: firstly, our method is more general in that it can be applied to other re-integration problems (as we will demonstrate in the remainder of this section); secondly, in relation to the data-fusion problem, we are concerned with the re-integration of the gradient field derived by SW’s method – which can be derived just as easily for any M -channel input image as it can for a (3-channel) colour image. The other colour-to-greyscale methods, on the other hand,

Fig. 2. 2A shows a deceptively simple colour image. 2B is the reconstruction using the SW plus Poisson – note the “bending” effect in the upper-left square. 2C shows Socolinsky’s much more complex level-sets method. 2D shows the new, fast, lookup-table based method.

are devised only for that particular problem, and do not naturally generalise to M -channel input images.

Here we also remind the reader that the image-fusion problem has a long and varied history of proposed solutions, from methods based upon dimensionality reduction techniques [26], wavelets based approaches [27], or by applying complex models of human visual processing [28]. Since the purpose of this paper is to describe our gradient reintegration method, we restrict our discussion to those methods that explicitly involve a gradient field reintegration.

B. Recovering Shading from Colour Images

Let us now demonstrate the utility of the method in another sphere: Shape from shading methods generally assume that intensity in an image varies with the shape of the object. Of course intensity also changes at material boundaries, and so it is of interest to find a method for disentangling material from shape changes [29]. Below we begin by briefly recapitulating one method, that set forth in [3]. First consider colour image formation, with $\rho = \{R, G, B\}$:

$$\rho(a, b) = \sigma I \int E(\lambda) S(\lambda) \mathbf{q}(\lambda) d\lambda \equiv \sigma I \mathbf{c} \quad (4)$$

In this equation, σ is shading and I is light intensity, both spatially varying, E is illumination, S is surface reflectance, and $\mathbf{q}$ and $\mathbf{c}$ are respectively the device spectral sensitivities and the intrinsic colour

response. Ultimately, here we are interested in the shading-intensity field σI since this term is a required input to shape from shading algorithms. Given that both I and σ are scalar functions, σI is a scalar function.

Let L be scalar “luminance” (or brightness), $L = \sum_{k=1}^3 \rho_k$ and let $\kappa = \sum_{k=1}^3 c_k$ be lightness. The κ image is often piecewise constant and only changes with material change. The chromaticity χ is defined as the pixel RGB vector ρ without magnitude: $\chi = \rho / L$. Now, we can write the following relations:

$$\begin{aligned} J &= \log(\sigma I); \\ \log L &= \log(\sigma I) + \log \kappa. \end{aligned} \tag{5}$$

i.e. J is the logarithm of the shading-intensity field and L is the logarithm of luminance.

Suppose large changes in chromaticity are found by thresholding the chromaticity: Since $\sum_{k=1}^3 \chi_k \equiv 1$, we use only the first two components $\{r, g\} = \{\chi_1, \chi_2\}$ of χ to form a chromaticity edge map:

$$E = \left\{ (\|\nabla r\|^2 + \|\nabla g\|^2)^{1/2} > t \right\} \tag{6}$$

with t a user-defined threshold value and E is *true* where there is an edge and *false* otherwise. Assuming that large changes in shading and intensity do not coincide with large colour changes, this thresholding operation gives a reasonable indicator of colour changes [3].

We now have the basic tools for recovering the shading-intensity field. We assume that the gradient field of the shading image can be calculated as:

$$\nabla J = \text{zeros}(\nabla \log L, E) \tag{7}$$

where the function $\text{zeros}(\cdot)$ assigns the left hand side gradient and sets it to zero where the threshold image, E , is *true*. Thus the meaning of Eq. (7) is that the gradient field ∇J is equal to all gradient information in the log brightness image, but with large changes omitted. Re-integrating returns J and then exponentiating gives us the shading (σI). Thus we have recovered all of the shading-intensity field from the luminance and the edge mask. Then, using Eq. (5) we can solve for the lightness field κ .

In Fig. 3A we show the simple picture of a cereal box from [3]. Fig. 3B shows the recovered κ image and Fig. 3C the recovered intensity-shading field using Poisson. Lookup-table based reconstruction, shown in 3D is completely comparable, demonstrating the utility of the lookup-table approach in this domain. We revisit this shading-recovery problem with respect to lookup-table based reconstruction in §3.

III. LOOKUP-TABLE BASED RECONSTRUCTION

In the lookup-table based approach we propose that:

Fig. 3. 3A shows the original colour image. The recovered albedo is shown in 3B, calculated from the corresponding intensity-shading field 3C re-integrating using Poisson. 3D shows the comparable lookup-table based reconstruction

$$||\nabla f(\rho) - \nabla J|| \approx 0 \quad (8)$$

where $f(\cdot)$ denotes a lookup-table based map and ∇J is the desired output gradient. That is, we assume that our reconstructed image is an onto map from the multichannel original, and consequently that the gradient of the mapping approximates the gradient of the reconstruction. The idea, at base, is that if we learn how to map the input set of gradients onto the desired output gradient, then with the lookup-table reconstruction assumption we have effectively also derived the mapping from a function of the multichannel input to the desired, re-integrated, image output.

Intuitively, we can argue that enforcing a lookup-table approach (at least to some degree) is a good thing. Fig. 4A shows an image of alternating red and blue stripes. We calculate the gradient field according to the SW method and arrive at Fig. 4B. Clearly the image looks wrong. But, the construction is in fact exact: the tensor gradient field of 4A is integrable and is exactly reconstructed in 4B. The figure looks unnatural because, perceptually, we would like the same colour to map to the same greyscale. 4C shows the more natural looking lookup-table based greyscale encoding.

We can address this unnaturalness above by enforcing integrability over all *scales*. One way to do this is to compare image locations which are an offset apart. We use the notation $\nabla G_{a,b,j,k}$ to denote the gradient calculated as: $[G(a+j) - G(a) \quad G(b+k) - G(b)]^t$. This idea is used in the Retinex dynamic range compression algorithm as a means to incorporate information from all scales in an image[30]. In terms of gradients calculated from the colour tensor we use the notation $\nabla J_{a,b,j,k} = Tda(\rho(a,b), \rho(a+j,b), \rho(a,b+k))$ where the neighbouring pixels in each of the colour channels are also an offset (j,k) away. To allow gradients to be calculated without worrying about offsets that go outside the image extent we assume our image is periodic so that offsets can wrap around. A periodic image is made by replicating

Fig. 4. Fig 4A. Simple colour image. 4B is an exact reconstruction of 4A using the Socolinsky and Wolff plus Poisson method: its gradient just equals the colour gradient. Clearly, the encoding as a ‘ramp’ does not look realistic. The lookup-table based gradient field reconstruction looks much more perceptually plausible.

an image first by mirror reflecting at the right hand border of an image and then reflecting horizontally at the bottom of the image. This results in a new image 4 times the size of the original which is now periodic. This “flip-flip” operation is common when using Fourier methods to manipulate the gradient field of the image.

Lookup-Table Theorem for Image Fusion: Suppose that G is recovered from $\nabla J = Tda(\rho)$ such that $\nabla G = \nabla J$ (the gradient field is integrable) and that $\nabla G_{a,b,j,k} = Tda(\rho(a,b), \rho(a+j,b), \rho(a,b+k))$. Then $G = f(\rho)$. **Proof:** Suppose that $\rho(a,b) = \rho(u,v)$ where $u = a - j + 1$ and $v = b - k + 1$. Then by simple substitution $J_{a,b,1,1}$ (a normal tensor derivative) is the same as $J_{u,v,j,k}$. Further, $J_{a,b,1,1} = J_{u,v,j,k} = G_{a,b,1,1} = G_{u,v,j,k}$. This implies that $G(a+1,b) - G(a,b) = G(a+1,b) - G(u,v)$ which is true if and only if $G(u,v) = G(a,b)$. Similarly, $G(a,b+1) - G(a,b) = G(a,b+1) - G(u,v)$ implies $G(u,v) = G(a,b)$. Thus, the same colour in an image is mapped to the same greyscale in G : there is a lookup-table map taking ρ to G . ■

We make three observations about this theorem. First, at heart it says something simple: if we are allowed to compare two colour pixels at different locations and assess their contrast (which seems reasonable) and a greyscale is recovered exactly then colour is an onto map to greyscale. Second, while for most real images a greyscale will not exist that matches contrast at all scales, we will show that using a lookup-table based mapping nonetheless results in contrast-preserving artefact-free images. Finally, although the proof is tailored to the data fusion problem from a multichannel image, the lookup-table method can also be applied to other gradient-domain problems.

A. Implementation

There are a number of ways we might build a lookup-table based map. Here we simply take the original multichannel image (be it RGB or higher-D) and make a 2nd-order polynomial expansion P : i.e., rather than RGB at each point we also have RG , GB , RB , R^2 , G^2 and B^2 . We now calculate the gradient in each of these 9 channels and then find the linear combination of these closest to the target gradient field: we solve for

$$\min_{\mathbf{z}} \|\sum_{i=1}^9 \nabla P(\boldsymbol{\rho})_i z_i - \nabla J\|. \quad (9)$$

The coefficient vector $\mathbf{z}$ which minimizes Eq. (9) (the linear combination of the gradients of the polynomial terms which is closest to the target gradient) is found by linear regression. To be more explicit; once we have calculated the x and y gradients for each of the 9 channels to give ∇P , we can resize this 9-channel gradient-image into an $2M \times 9$ matrix $\mathbf{P}_{xy}$, where M is the number of pixels. The coefficient vector $\mathbf{z}$ is then given by

$$\mathbf{z} = (\mathbf{P}_{xy}^T \mathbf{P}_{xy})^{-1} \mathbf{P}_{xy}^T \mathbf{J}_{xy} \quad (10)$$

where $\mathbf{J}_{xy}$ contains the target gradients. If there are fewer image colours than there are terms in the polynomial expansion, the matrix $(\mathbf{P}_{xy}^T \mathbf{P}_{xy})$ is not invertible, and the Moore-Penrose pseudo-inverse is used instead.

This formulation also allows the inclusion of gradients calculated at multiple scales. If we have $\mathbf{J}_{xy}^i$ for scales $i = 0, \dots, N$ (scale zero is the original image), these gradients can be added as additional elements to the vector $\mathbf{J}_{xy}$; i.e. $\mathbf{J}_{xy} = [\mathbf{J}_{xy}^0, \dots, \mathbf{J}_{xy}^N]$. In the image fusion problem, for example, these could be generated by smoothing $\boldsymbol{\rho}$, downsampling to give $\boldsymbol{\rho}^i$ for the i -th scale, and recalculating $Tda(\boldsymbol{\rho}^i)$ at each point. These multiscale derivatives can be introduced into Eq. (10) by defining $\mathbf{P}_{xy}^i = \nabla P(\boldsymbol{\rho}^i)$ as the calculated gradients at scale i , and setting $\mathbf{P}_{xy} = [\mathbf{P}_{xy}^0, \dots, \mathbf{P}_{xy}^N]$.

Note that even redefining $\mathbf{P}_{xy}$ and $\mathbf{J}_{xy}$ to incorporate multiple scales, the dimension of the coefficient vector $\mathbf{z}$ is unchanged, and once calculated, this is ultimately applied to $\boldsymbol{\rho}$ without recourse to the blurred, then downsampled, images $\boldsymbol{\rho}^i$. The final effect of including multiple scales will therefore be to weight the gradient of the final image to better match the gradients of the included scales. When only the finest gradients are included ($\mathbf{J}_{xy}^0$), the minimisation in Eq. (10) is driven solely by local-contrasts. The most useful scale for a given application will be dependent upon the scale of the features that are most important for that application.

In a final comment about Eq. (10) we note that ∇P need not be limited to 2nd order polynomial functions; we have, for example, experimented with using higher-order polynomials, Tikhonov regularization, and other functional expansions. We also point out that the dimensionality of the gradient image ∇P will clearly change when we have a different input dimension; e.g. with a second order polynomial expansion a 7-channel multispectral image will generate 28 polynomial channels, while a 2-channel chromaticity image will generate just 5.

Colour to Greyscale

In the colour-to-greyscale example, the gradient field is constructed from the colour differential tensor. Thus the target gradient ∇J is calculated according to Eq. (1) and the lookup-table based minimization is Eq. (9, where now ρ is indeed 3-vector colour and ∇J is greyscale gradient capturing colour contrast. The recovered greyscale image is then simply calculated as:

$$G = \sum_{i=1}^9 P(\rho)_i z_i \quad (11)$$

Fig. 4C shows the lookup-table reconstruction of the colour gradient image calculated for 4A. In this and future examples we include only the finest scale in our calculations, to keep a fair comparison with other methods (which only have access to the smallest scale). Note that the reconstruction is much more realistic than solving using SW.

Recovering Shading from Colour Images

As another example, we recover the shading field from a colour image, using two stages. First, we recover κ (the scalar colour albedo). Second we substitute κ in (5) to recover the shading field σI . To begin, we calculate the log-intensity gradient from the image, $\nabla \log L$, and we set to zero locations where we think there is **not** a material edge. The thresholded log-intensity gradient is written as $zeros(\nabla \log L, \neg E)$. This gradient field should be the same as the lightness. Using a lookup-table based approach we can now solve:

$$\min_{\chi} ||\sum_{i=1}^9 \nabla P(\chi)_i z_i - zeros(\nabla \log L, \neg E)|| \quad (12)$$

We now calculate κ using:

$$\kappa = \sum_{i=1}^9 P(\chi)_i z_i \quad (13)$$

Substituting into (5) we can recover the shading field. Fig. 3D shows a lookup-table reconstruction of the shading image. It is remarkably close to that returned by solving the Poisson equation.

Advantages, and a Proviso

As a general remark, we note here that one of the key advantages of the lookup-table based method is that it works in the *forward* direction in the sense that our reconstruction depends only on the derivatives calculated from the original multichannel image mapping to those that are desired (e.g., thresholded greyscale derivatives). Traditional reconstruction works in the *backward* direction and attempts to re-integrate the derivatives.

Finally, we point out to the reader that any integration process essentially leaves a constant of integration undefined. This constant is generally heuristically defined. For colour to greyscale we add a scalar so that the mean of the greyscale image equals the mean of the colour input. In shading recovery we multiply by a factor so the average shading field equals the average image brightness.

IV. EXPERIMENTS

In Fig. 5A we show a painting by Monet with the Luminance encoding shown in 5B. Notice that the contrast between the sun and sky is gone. Indeed, this image is one of the examples typically used to test colour to greyscale algorithms [31]. In 5C we show the output of the least-squares reconstruction to the gradient field derived from the image colour tensors according to the method of SW. Here their method has worked well and the contrast between sun and background is preserved in the output. Fig. 5D shows the output of our lookup-table based gradient field reconstruction method. Here again we have greyscales that preserve the contrast details of the original.

In our second experiment we take a further 3 images shown in the top row of Fig. 6. These are a hillside landscape, a parrot image, and a segmented parrot (6A, 6B and 6C). All 3 images are known to cause problems for the least-squares reconstruction of the SW method, the outputs of which are shown in 6D, 6E and 6F. Here we see clear “bending” artefacts around the tail of the segmented parrot (6F), which appear in a less severe form in the original parrot image (6E) while the landscape in 6D has a visible loss of contrast in the bottom-left quadrant. We applied the Agrawal et al. 2006 method [10] to re-integrate the gradient field delivered by SW’s method and this results in outputs 6G, 6H and 6I. Note there is a definite improvement over the SW output for the two parrot images (6H and 6I), although these images are still somewhat “washed-out”. However, the method fails badly in 6G, where severe “blocking” artefacts can be seen in the clouds. In the reconstructions of our lookup-table method (6J, 6K and 6L) these artefacts are not visible, although we note that this does come with the cost of reducing the “pop-out” of some regions, e.g. the chest of the segmented parrot.

In Fig. 7, we show a further 3 images that are known to cause problems for traditional Luminance encodings of colour images. Figs 7A, 7B and 7C show the colour originals and 7D, 7E and 7F the

Fig. 5. 5B shows the usual greyscale of 5A. Notice the sun as all but disappeared. 5C and 5D show the greyscale produced by the Socolinsky and Wolff and our lookup-table colour gradient field reconstruction.

Luminance encodings. In all three cases important details have been lost. In 7G, 7H and 7I we show the results of our lookup-table method. Again we have natural greyscale images, without artefact, and which maintain the details present in the colour original. Further, our experience is that our lookup-table based gradient field reconstruction method produces a natural looking greyscale image for colour counterparts which also maintains the visible colour contrasts.

Finally, we demonstrate that this method is a general image fusion method rather than simply a colour-to-greyscale technique. In Fig. 8 we show two registered images: the first (A) is a 3-channel colour image and the second (B) is a 1-channel near-infrared (NIR) image [32]. Together these make up a 4-channel multispectral image. Unlike other colour-to-greyscale methods, the method of SW naturally generates a gradient field for this four-channel image since e.g. SW's method does not rely on special directions in 3-space. In the figure we then show four different re-integrations using antecedent methods. In (C) is a least-squares Poisson solver, in (D) is the method of Agrawal et al. 2006 [10], in (E) is the later method

Fig. 6. 6A, 6B and 6C are colour originals. 6D, 6E and 6F are the greyscales obtained by SW. Then, 6G, 6H and 6I are the greyscales returned using Agrawal et al.'s re-integration [10]. The reconstructions 6D through 6I show visible reconstruction artefacts. These are removed by our lookup-table method: 6J, 6K and 6L.

of Reddy et al. 2009 [11]; finally, (F) shows the output of our Lookup Table method. In (D) we can see large integrability artefacts around the bush on the right-hand side of the image in particular, and also around tower in the centre of the image, while in (E) these are reduced slightly, but still visible. In (F), a further improvement is noticeable, but the image is still “washed-out”. The new method in (F) does not

Fig. 7. 7A, 7B and 7C are known problem images for traditional luminance greyscale encodings: 7D, 7E and 7F. Our lookup-table based colour to greyscale transform, 7G, 7H and 7I, produces good reconstructions for all 3 images

produce any of the above artefacts.

Human preference study

In a final experiment, to evaluate the performance of the new algorithm we compare it with other existing algorithms in a psychophysical evaluation. We use a 2-alternative-forced-choice (2AFC) paradigm, where pairs of colour-to-grey renderings are shown side by side and observers are asked to choose their preferred reproduction. We compare our method with three other algorithms: the luminance transform, which is the standard method and displays an estimate of the brightness at each pixel; the Gooch et al. transform[17], which also aims to maintain contrast in the output image and is considered the state-

Fig. 8. 8A and 8B show a colour image and a corresponding image captured in the NIR (taken from [32], which together make a 4 channel image. Figures 8C-F show reintegration examples corresponding to the vector field induced by 8A and 8B (see text for details).

of-the-art method ²; and the original algorithm of SW, where the gradient field is re-integrated using a

²This algorithm contains three parameters, which we set at values of $\alpha = 10$, $\mu = FULL$ (each pixel is compared with all others) and $\Theta = 45^\circ$ for all images.

least squares solution to Poisson’s equation [12]³. To compare the methods we used 10 colour images, composed of 6 natural scenes and 4 graphics images, and 10 naive observers, where each observer repeated a given judgement 4 times.

In Fig. 9 we show the raw scores for each of the methods, where the raw score represents the total number of times that a method is preferred, summed over all observers, repetitions and images. It is important to note that the errorbars in this figure do not explicitly represent uncertainty in the scores, but rather overlapping errorbars indicate scores that are not significantly different at the 5% level. To compute the significance levels we use the method outlined by David [34]. In this method one computes a critical score-difference at a given confidence level, and when a difference between two scores exceeds this critical level the difference is deemed to be significant. Here the significant difference was found to be 73, thus the errorbars have an extent of 73 units. The results of the experiment show that the lookup-table method significantly outperforms the other three algorithms. We also find that this result is more marked for graphics images than for natural scenes, where the three contrast-preserving methods have the same order, but the differences between them are not significant.

A. Shading Recovery

The lookup-table based recovery of the shading-intensity field of 3A is shown in 3D. Clearly, we have removed the effect of the coloured lettering and are ready to carry out shape from shading. Now, we wished to examine our method more quantitatively. Indeed, it would also be interesting to compare the shading-field result with some ground-truth field, instead of to the least-squares Poisson solution.

So we will construct a synthetic example. Let us apply the colour in Fig 3A to a same-sized shading field derived from a depth map (Mozart bust); the result is shown in 10A. We have a ground-truth shading field, 10B, to compare to (shading is derived from an illuminant from angles $\{slant, tilt\} = \{56^\circ, 76^\circ\}$, applied to the depth map).

We compare the actual shading field shown in 10B with the shading field recovered from the lookup-table based reconstruction shown in 10C. (That from the Poisson method is very similar, with correlation coefficient 0.9901.) Comparing to the correct, synthetic shading, in this case, the lookup-table method does slightly better: the original method produces shading with a correlation coefficient of 0.9827 compared to the ground truth, whereas the lookup-table method’s value is 0.9966. Mesh plots of the actual and lookup-table based recovered shading fields are shown in 10D and 10E.

³The SW method can generate a greyscale image whose values exceed the range [0 1]; in this case we use the linear gamut mapping algorithm described in [33]

Fig. 9. Psychophysical evaluation of four methods for colour-to-grey: Lookup-table based method significantly outperforms antecedent methods

V. CONCLUSIONS

There are many problems in computer vision that are solved in the gradient domain. An image is differentiated, the gradients are manipulated, and an output image is found by re-integration. However, the idea, though simple and elegant, is sometimes difficult to implement in practice because the manipulated gradient field is non-integrable: there is no image that when differentiated results in the gradient field at hand. Least-squares re-integration can result in large smearing artefacts[15]. While methods exist for correcting gradient fields [13] they often work best when the field is almost integrable.

In this paper we argue that for multichannel images the underlying gradient field can be re-integrated using a lookup-table map to the target gradient field. In particular we examine the problem of recovering the shading field for a colour image and that of coding multichannel image contrast in a greyscale image. Both these problems are amenable to lookup-table based gradient field reconstruction. Experiments demonstrate that we can transform colour images to greyscale without introducing artefacts and also recover the shading field from colour images. Moreover, our results compare favourably with antecedent

Fig. 10. 10A shows the chromaticity image of the Raisin Bran cereal box modulated by the Mozart shading field. 10B and 10C show respectively the actual ground-truth shading and the shading recovered by our lookup-table method. Figs 10D and 10E show the shading fields as a mesh plot.

gradient field reconstructions in user studies. Further, our lookup-table method is both fast and simple to implement.

VI. ACKNOWLEDGEMENTS

This work was supported by the Engineering and Physical Sciences Research Council under Grant EP/E012248/1 and Grant EP/E121159/1, and by a Discovery Grant from the Natural Sciences and Engineering Research Council of Canada (M.S.D). We would also like to thank Dikpal Reddy for making his L_1 minimisation code available for testing.

REFERENCES

- [1] A. Blake, “Boundary conditions for lightness computation in Mondrian world,” *Comp. Vision, Graphics, and Image Proc.*, vol. 32, pp. 314–327, 1985.
- [2] G. Finlayson, S. Hordley, C. Lu, and M. Drew, “On the removal of shadows from images,” *IEEE Trans. Patt. Anal. Mach. Intell.*, vol. 28, pp. 59–68, 2006.
- [3] B. V. Funt, M. S. Drew, and M. Brockington, “Recovering shading from color images,” in *ECCV-92: Second European Conference on Computer Vision*, G. Sandini, Ed., 1992, pp. 124–132.
- [4] B. K. P. Horn, *Robot Vision*. MIT Press, 1986.
- [5] R. Woodham, “Photometric method for determining surface orientation from multiple images,” *Optical Engineering*, vol. 19, pp. 139–144, 1980.
- [6] P. Pérez, M. Gangnet, and A. Blake, “Poisson image editing,” *ACM Trans. Graph.*, vol. 22, no. 3, pp. 313–318, 2005.
- [7] B. Karacali and W. Snyder, “Noise reduction in surface reconstruction from a given gradient field,” *Int. J. Comp. Vision*, vol. 60, pp. 25–44, 2004.
- [8] R. Raskar, A. Ilie, and J. Yu, “Image fusion for context enhancement and video surrealism,” in *NPAR '04: 3rd Int. Symp. on Non-photorealistic Animation and Rendering*, 2004, pp. 85–152.
- [9] D. Socolinsky and L. Wolff, “Multispectral image visualization through first-order fusion,” *IEEE Trans. Im. Proc.*, vol. 11, pp. 923–931, 2002.
- [10] A. Agrawal, R. Raskar, and R. Chellappa, “What is the range of surface reconstructions from a gradient field?” in *ECCV 2006*, 2006, pp. 578–591.
- [11] D. Reddy, A. Agrawal, and R. Chellappa, “Enforcing integrability by error correction using l1-minimization,” in *CVPR*, 2009, pp. 2350–2357.
- [12] R. T. Frankot and R. Chellappa, “A method for enforcing integrability in shape from shading algorithms,” *PAMI*, vol. 10, pp. 439–451, 1988.
- [13] A. Agrawal, R. Chellappa, and R. Raskar, “An algebraic approach to surface reconstruction from gradient fields,” in *ICCV05*, 2005, pp. I: 174–181.
- [14] D. A. Socolinsky, “A variational approach to image fusion,” Ph.D. dissertation, The Johns Hopkins University, April 2000.
- [15] C. Fredembach and G. Finlayson, “Hamiltonian path based shadow removal,” in *16th British Machine Vision Conf, BMVC2005*, 2005, pp. 970–980.
- [16] D. Socolinsky and L. Wolff, “A new visualization paradigm for multispectral imagery and data fusion,” in *CVPR*, 1999, pp. I:319–324.
- [17] A. Gooch, S. Olsen, J. Tumblin, and B. Gooch, “Color2gray: salience-preserving color removal,” *ACM Trans. Graph.*, vol. 24, no. 3, pp. 634–639, 2005.
- [18] R. Bala and R. Eschbach, “Spatial color-to-grayscale transform preserving chrominance edge information,” in *12th Color Imaging Conference: Color, Science, Systems and Applications.*, 2004, pp. 82–86.
- [19] M. Grundland and N. Dodgson, “Decolorize: Fast, contrast enhancing, color to grayscale conversion,” *Pattern Recogn.*, vol. 40, pp. 2891–2896, 2007.
- [20] K. Rasche, R. Geist, and J. Westall, “Re-coloring images for gamuts of lower dimension.” *Computer Graphics Forum*, vol. 24, pp. 423 – 432, 2005.
- [21] Y. Kim, C. Jang, J. Demouth, and S. Lee, “Robust color-to-gray via nonlinear global mapping,” *SIGGRAPH ASIA 2009*, vol. 29, no. 5, Dec. 2009.

- [22] M. Qiu, G. Finlayson, and G. Qiu, “Contrast maximizing and brightness preserving color to grayscale image conversion,” in *CGIV08*, 2008, pp. 347–351.
- [23] S. Di Zeno, “A note on the gradient of a multi-image,” *Comp. Vision, Graphics, and Image Proc.*, vol. 33, pp. 116–125, 1986.
- [24] L. Neumann, M. Cadik, and A. Nemcsics, “An efficient perception-based adaptive color to gray transformation,” in *Workshop on Computational Aesthetics in Graphics, Vis. and Imaging*, 2007, pp. 73–80.
- [25] A. Alsam and M. Drew, “Fast colour2grey,” in *16th Color Imaging Conf.*, 2008, pp. 342–346.
- [26] J. Tyo, A. Konsolakis, D. Diersen, and R. Olsen, “Principal-components-based display strategy for spectral imagery,” *IEEE Trans. on Geosciences and Remote Sensing*, vol. 41, pp. 708–718, 2003.
- [27] P. J. Burt and E. H. Adelson, “Merging images through pattern decomposition,” *Proc. SPIE 575*, pp. 173–181, 1985.
- [28] A. Waxman, A. Gove, D. Fay, J. Racamoto, J. Carrick, M. Seibert, and E. Savoye, “Color night vision: Opponent processing in the fusion of visible and ir imagery,” *Neural Networks*, vol. 10, pp. 1–6, 1997.
- [29] J. M. Rubin and W. A. Richards, “Color vision and image intensities: When are changes material,” *Biological Cybernetics*, vol. 45, pp. 215–226, 1982.
- [30] J. Frankle and J. McCann, “Method and apparatus for lightness imaging,” in *U.S. Patent 4,384,336*, May 1983.
- [31] M. Drew, D. Connah, G. Finlayson, and M. Bloj, “Improved colour to greyscale via integrability correction,” in *Human Vision and Electronic Imaging XIV*, 2009.
- [32] C. Fredembach and S. Süsstrunk, “Colouring the near infrared,” in *16th Color Imaging Conf.*, 2008, pp. 176–182.
- [33] D. Connah, G. Finlayson, and M. Bloj, “Coding contrast as brightness to convert colour images to grey-scale,” *Perception*, vol. 36, pp. 303–303, 2007.
- [34] H. David, *The Method of Paired Comparisons*. Charles Griffin and Co., London, 1969.

Professor Graham Finlayson gained his PhD from Simon Fraser University in 1995 and was appointed to a Professorship in Computing Sciences at the University of East Anglia in 1999. He has to date published over 200 articles and conference papers. In 2003 Finlayson was awarded the Philip Leverhulme prize for science. Finlayson was granted a Royal Society-Wolfson Merit Award in recognition of his research contribution to the fields of colour in computer vision and colour imaging in October 2008. In November 2009, he received the Davies medal from the Royal Photographic Society for his contribution to digital imaging. In November 2009 he was elected a Fellow of the Society of Imaging Science and Technology.

Dr. David Connah received both his BSc degree (Biology) and MSc (Machine Perception and Neuro-computing) from Keele University, in 1997 and 1998 respectively. In 2004 he gained a PhD in Colour Science from the Colour and Imaging institute, University of Derby followed by two years as a post-doctoral researcher in the Norwegian Colour Lab, based at Gjøvik University College. Since 2005 has worked in the Colour Group at the University of East Anglia with Prof. Graham Finlayson. His research interests include reflectance recovery, automated classification problems, and image fusion.

Mark S. Drew is a Professor in the School of Computing Science at Simon Fraser University in Vancouver, Canada. His background education is in Engineering Science, Mathematics, and Physics. His interests lie in the fields of image processing, color, computer vision, computer graphics, multimedia, and visualization. He has published over 150 refereed papers in journals and conference proceedings. Dr. Drew is the holder of a U.S. patent in digital color processing, a U.S. patent application in color computer vision, and two U.S. and U.K. patent applications in image reconstruction and image accenting. He is the co-author of the book “Fundamentals of Multimedia” published by Prentice-Hall in 2004.