

# COLOUR-DEPENDENT LINEAR VECTOR IMAGE FILTERING

Stephen J. Sangwine<sup>†</sup>, Todd A. Ell<sup>†‡</sup>, Barnabas N. Gatssheni<sup>†</sup>

<sup>†</sup> University of Essex, Department of Electronic Systems Engineering,  
Wivenhoe Park, Colchester, CO4 3SQ, UK.

<sup>‡</sup> 5620 Oak View Court, Savage, Minnesota, USA.

Email: s.sangwine@ieee.org, t.ell@ieee.org, bngatsheni@iee.org

## ABSTRACT

In previous work the authors have presented initial results from a new method of linear vector filtering applied to colour images. The method depends on resolution of colour pixel vectors into directions parallel and perpendicular to a chosen colour, followed by application of separate filters to the resolved image components in order to filter the chosen colour only (for example with a low-pass filter). In this paper we review this scheme and then present a new and detailed analysis of the limitations of the scheme as currently defined. Understanding of these limitations reveals for the first time where progress must be made with new linear operations before linear vector filters can be fully realized for practical use.

## 1. INTRODUCTION

In our previous work on colour image filtering [1] we have shown that it is possible to implement linear *vector* filters, with application to the processing of colour images. For example, it has been shown that a colour image may be selectively low-pass filtered so that only pixels of a certain *colour of interest* (COI) are filtered, leaving other pixels largely untouched. In the past this has been done by first defining pixels of interest in the image with an outline mask, or by colour thresholding, and then filtering only the selected pixels. In our work, all the pixels in the image are filtered (by convolution), but the *filtering itself* is inherently selective.

Filters of this type may be designed mathematically by resolving all the pixels within a colour image into directions in colour space using a standard parallel/perpendicular vector resolution, then applying a scalar filter to only one of the separated image components, and finally adding the two images back together to obtain the filtered result. In [1] we showed that the implementation of the filter need not actually include the parallel/perpendicular separation, because it is possible to define this separation algebraically, and then merge the separation step into the convolutions. This means that the filter may be implemented in the Fourier domain (using a suitably defined hypercomplex Fourier transform) such as those based on quaternions<sup>1</sup> which two of the authors have worked upon [2, 3]). The advantage of implementation in the Fourier domain is faster computation for large mask sizes, just as with real or complex convolution. Vector filtering based on hypercomplex convolution was first reported in [4] and in expanded form in [5]. For completeness we reproduce here the

definition of hypercomplex convolution:

$$g(n, m) = \sum_{x=-X}^X \sum_{y=-Y}^Y h_L(x, y) f(n-x, m-y) h_R(x, y) \quad (1)$$

where  $g(n, m)$  is the filtered image, and  $f(x, y)$  is the original image, both of dimension  $N \times M$ ;  $h_L(x, y)$  and  $h_R(x, y)$  are the left and right masks respectively, both of dimension  $(2X+1) \times (2Y+1)$ . The images are represented by pure quaternion pixels and there are two masks with quaternion coefficients, one on the left of the image and one on the right. Quaternion multiplication is not commutative in general, and the left and right coefficients permit masks to be defined that implement linear operations on the pixel vectors, such as rotation. These operations are implemented numerically using quaternion arithmetic, and are amenable to algebraic manipulation using quaternion algebra. The convolution process is standard in image processing, and the only difference with hypercomplex convolution is that the arithmetic is hypercomplex and there are coefficients on the left and right of the pixel values.

In the next section we review briefly the work presented in [1] and then explain its limitations. We also present colour ramp images in which all the pixels are co-directional in colour space, but vary in magnitude. In section 3 we present some analysis of the filtering scheme which explains the limitations and suggests some approaches to obtaining better colour-selectivity.

## 2. PREVIOUS WORK

Figure 2 shows the derivation of the filtering scheme presented in our previous work [1]. We first choose a colour of interest (COI). The intention is to filter pixels in the image which are close to the COI, in some sense. For example, we might apply a low-pass filter to pixels which are close to the COI, so that regions in the image with colours near to the COI are smoothed, while the rest of the image is left largely untouched.

We have chosen to work in an offset RGB colour space in which the origin is at the centre of the RGB cube, that is the origin is at mid-grey with normalized RGB values  $(\frac{1}{2}, \frac{1}{2}, \frac{1}{2})$ , rather than at black with normalized RGB values  $(0, 0, 0)$ . This means that all pixel values are vectors directed away from mid-grey. To convert an image from standard normalized RGB values to this 'offset' RGB space, we simply have to subtract a constant (or DC) value  $(\frac{1}{2}, \frac{1}{2}, \frac{1}{2})$  from all the pixels. A full description and discussion of this grey-centered RGB colour space appears in [6]. It should be noted that subtracting an offset from all the pixel values is a non-linear opera-

Todd A. Ell is a Visiting Fellow at the University of Essex. Barnabas N. Gatssheni was formerly with the University of Essex, and his work was funded by the UK Engineering and Physical Sciences Research Council under grant GR/M 45764.

<sup>1</sup>A brief summary of quaternions is included in Appendix A.




Figure 1: Example colour ramp images. Top to bottom: greyscale, red/cyan, green/magenta, blue/yellow, pink/turquoise. Vector directions:  $(1,1,1)$ ,  $(1,-1,-1)$ ,  $(-1,1,-1)$ ,  $(-1,-1,1)$ ,  $(1,0,0)$ , respectively

tion. It corresponds in signal processing to conversion from unipolar samples to bipolar samples.

Figure 1 shows several colour ramp images. All the pixels in a given colour ramp image are aligned in one direction in colour space. The magnitude of the pixel vector varies uniformly from left to right along the image, passing through zero at the centre of the image (where the colour is always mid-grey). Thus the colours to the left of the centre of the image are opponent colours to those on the right. In terms of the perceptual quantities *hue*, *saturation* and *intensity*, these images have constant hue from centre to left edge and from centre to right edge (and the two hues are opponent hues, with a difference of  $180^\circ$ ); saturation varies from a maximum at the left edge to zero at the centre and back to a maximum at the right edge; intensity varies monotonically from left to right edge. Given this representation of an image (which is only slightly different from the standard RGB representation), we can *resolve* an image into two images using the usual concept of resolving a vector into a chosen direction and a perpendicular plane. This is done using quaternion arithmetic as shown in Appendix A. The chosen direction is defined by our COI relative to mid-grey. All pixels in the image are resolved into a component parallel to the COI, and a component perpendicular to the COI. Obviously, a pixel with a value which is exactly that of the COI will have no component in the plane perpendicular to the COI, and a pixel with a value which is exactly perpendicular to the COI will have no component in the direction of the COI. Other pixels will fall somewhere in between these extremes. An obvious special intermediate case is that of a pixel whose value subtends an angle of  $45^\circ$  to the COI, since this will be resolved into equal components in the direction of the COI, and in the plane perpendicular to the COI.

The implementation of the filter consists of the sum of two hypercomplex convolutions computed in parallel, and the decomposition step shown in the figure is merged into the hypercomplex multiplications. This is possible because the parallel/perpendicular resolution can be expressed algebraically and this algebraic representation can then be combined algebraically with the convolution coefficients. The details were given in [1].

### 3. LIMITATIONS

The filtering scheme reviewed in the previous section suffers from three limitations which were not previously understood, and which have not been published in previous work. These limitations are:




Figure 2: Colour selective filtering scheme used in [1].

- The colour selectivity is limited. Discrimination between colour vectors close to the COI and those at up to  $45^\circ$  from the COI is limited. This is a fundamental problem whose solution is required before more colour-specific filters may be designed, and in what follows, we show why the lack of discrimination occurs, and that a seemingly obvious way to improve it does not in fact make any difference.
- When the COI is close to the grey direction (roughly ‘luminance’), the lack of sharp colour selectivity means that luminance is also filtered, even in the case of pixels which are not close to the COI. Geometrically, this is the same limitation as in the previous point, although its perceptual effect in the filtered image is different because the human visual system has much greater acuity to luminance than chrominance.
- The opponent colour to the COI is not separable from the COI using a linear operation. (Opponent colour examples are: black – white; red – cyan; green – magenta; blue – yellow.) This is not an insoluble problem, because the required non-linear operation is easy to implement, but because the resulting filter includes a non-linear operation, it will not be possible to implement the entire filter in the Fourier domain. A more accurate categorization of a filter which included separation of the COI from its opponent colour would be *quasi-linear* rather than *non-linear*.

All three of these limitations are discussed in depth below.

#### 3.1 Colour selectivity

The limited colour selectivity achieved by resolving pixel vectors into the direction of the COI and the plane perpendicular to the COI is very simply explained. When a vector  $p$  is resolved into the COI direction (specified by a unit vector  $u$ ), the resolved vector is given by  $u(u \cdot p)$ , where  $\cdot$  denotes the ‘dot’ or scalar product. More specifically, the resolved vector is given by  $u|p|\cos\theta$  where  $\theta$  is the angle between  $u$  and  $p$ . Since the magnitude of the resolved vector depends on the cosine of the angle between the vector  $p$  and the COI direction, there is limited ‘rolloff’ with angle, since the cosine function changes only slowly with angle until the angle exceeds  $45^\circ$ . (We use the term ‘rolloff’ here by analogy with the attenuation of a linear filter, except that what we are trying to express is the way that the selectivity of the filter varies as the pixel vectors increase in angle from the COI direction. What we need is a sharper rolloff with angle than we obtain with the scheme reviewed in this paper.)

It might seem that a solution to this problem could be obtained by resolving the pixels so that the COI direction is in the plane perpendicular to the direction of the unit vector  $u$ , and then resolving the pixels in the perpendicular plane




Figure 3: Colour selective filtering scheme with two decompositions and filtering of the perpendicular separated image component.




Figure 4: Vector  $p$  resolved into the direction of a colour of interest (COI) by choosing, left: the parallel direction to coincide with COI; right: the perpendicular direction to coincide with COI.  $\theta$  is the angle between  $p$  and the parallel direction.

using a second direction perpendicular to both the COI and to  $u$ . Figure 3 illustrates a filter based on this scheme. The attraction of this scheme might be thought to be that the components of the pixel vectors resolved into the perpendicular plane are determined by the sine function of the angle  $\theta$ , thus giving exactly the improved selectivity desired. However, it turns out that this is a fallacy and Figure 4 shows why: it makes no difference whether the COI is in the parallel direction or the perpendicular plane – in both cases the angle between the vector  $p$  and the COI is the same and it is the cosine of this angle that determines the magnitude of the resolved vector. We include this point to make clear that the problem is a non-trivial one, despite its apparent simplicity.

### 3.2 Luminance rejection

The second limitation discussed above is closely related to that discussed in the previous subsection. The problem is that when a colour direction includes a significant luminance component, the weak colour selectivity achieved using vector decomposition means that many pixels in the separated image have a luminance component even though their chrominance in the chosen COI is minimal. This can be seen in Figure 5 which is a separation of the colour image ‘lena’ such that the COI is in the direction of the colour of the feathers hanging from Lena’s hat. It can be seen that although the feathers have been separated from the other colours in the image, we also have an almost greyscale image for all other areas of the image, and if we filter this separated image with a low-pass filter (intending to ‘soften’ the feathers), we will



Figure 5: Separation of the lena image in the colour direction corresponding to the blue feathers, showing strong luminance element in the separated image, and also opponent colour (the yellow areas).

also blur the rest of the image noticeably because of the luminance element of the remaining pixels in the separated image.

### 3.3 Opponent colour separation

The third problem noted above is that of opponent colours. Ideally, it would be desirable to be able to separate out a chosen colour (defined by a COI direction) and filter just pixels of this colour. However, if we use mid-grey as the origin of our colour space (and there are good reasons for doing so), we cannot separate the COI from its opponent colour (vectors in the opposite direction away from mid-grey) without using a non-linear operation. Figure 5 is an example of the problem and clearly shows the yellow colour which is opponent to the colour of the feathers (the COI).

The required operation is simple enough, and corresponds to extracting the sign of each pixel vector (analogous to extracting the sign bit from a signed integer sample in digital signal processing). However, there appears to be no linear algebraic way to perform this extraction and until this problem is solved in a linear algebraic fashion this step cannot be combined into other linear steps as we have been able to do with the parallel/perpendicular separation.

## 4. CONCLUSION

The problem of colour-dependent linear vector image filtering has been shown to be non-trivial. The paper has presented a new analysis of the difficulty of achieving colour-specificity using resolution of pixel vectors into the colour space direction of a colour of interest, and has also shown that problems exist for colour directions close to the luminance or grey axis of colour space, because in the case of these colours it is difficult to separate the colour of interest

from the luminance. Further work remains to be done, and it is suggested that some form of colour vector amplification may be the solution to the problem. We have presented a *non-linear* scheme for vector amplification in [7], but we seek, as in this paper, a linear solution.

The alternative is to resort to filtering schemes which rely on hue, saturation, intensity; or use a chrominance/luminance colour space, as in [8], for example. Such schemes are non-linear, or at best quasi-linear, and do not permit implementation in the Fourier domain, nor do they permit the application of the very wide body of knowledge developed in the field of linear systems. It is a major motivation of our work that this body of knowledge may be generalized to linear vector filtering, and thus lead to a systematic exploration of filtering and processing possibilities, which has not been achieved with non-linear filtering schemes.

The discussion in this paper presents our work so far on the problem of devising and designing linear vector filters for colour images, and although we have not yet found a good linear solution to the problem of colour selectivity, we have made some progress with understanding the problem itself.

### A. QUATERNIONS

The quaternions (also referred to in this paper as *hypercomplex numbers*<sup>2</sup>) were discovered by Hamilton in 1843. They combine by the normal rules of algebra with the exception that multiplication is not commutative. A quaternion has four components, one real and three imaginary. The usual notation, extended from that of the complex numbers is  $q = w + xi + yj + zk$  where  $w, x, y$  and  $z$  are real, and  $i, j$  and  $k$  are complex operators which obey the following rules:

$$\begin{aligned} i^2 = j^2 = k^2 = ijk = -1 \\ ij = k \quad jk = i \quad ki = j \\ ji = -k \quad kj = -i \quad ik = -j \end{aligned}$$

The conjugate of  $q$  (denoted here by an overbar) is  $\bar{q} = w - xi - yj - zk$  and its modulus is  $|q| = \sqrt{w^2 + x^2 + y^2 + z^2}$

A quaternion with zero real part is called a *pure* quaternion, and a quaternion with unit modulus is called a *unit* quaternion. The imaginary part of a quaternion has three components and may be associated with a 3-space vector. For this reason, it is sometimes useful to consider the quaternion as composed of a vector part and a scalar part, thus:  $q = S(q) + V(q)$ , where the scalar part,  $S(q)$ , is the real part ( $w$  in our notation above), and the vector part is a composite of the three imaginary components,  $V(q) = xi + yj + zk$ .

A pure quaternion, or vector,  $v$  may be decomposed about a direction specified by a vector/pure quaternion  $\mu$  into components  $v_{\parallel}$  and  $v_{\perp}$  such that  $v = v_{\parallel} + v_{\perp}$  and  $v_{\perp} \perp \mu$  and  $v_{\parallel} \parallel \mu$ . The parallel ( $v_{\parallel}$ ) and perpendicular ( $v_{\perp}$ ) component of each pixel is given by [9]:

$$v_{\parallel} = \frac{1}{2}(v - \mu v \mu) \quad v_{\perp} = \frac{1}{2}(v + \mu v \mu) \quad (2)$$

<sup>2</sup>Strictly, the term *hypercomplex number* refers to a more general case of  $n$  components,  $n \geq 2$ , so that, while quaternions are hypercomplex, not all hypercomplex numbers are quaternions.

### REFERENCES

- [1] S. J. Sangwine, B. N. Gatscheni, and T. A. Ell. Linear colour-dependent image filtering based on vector decomposition. In *Proceedings of EUSIPCO 2002, XI European Signal Processing Conference*, volume II, pages 274–277, Toulouse, France, 3–6 September 2002. European Association for Signal Processing.
- [2] S. J. Sangwine and T. A. Ell. The discrete Fourier transform of a colour image. In J. M. Blackledge and M. J. Turner, editors, *Image Processing II Mathematical Methods, Algorithms and Applications*, pages 430–441, Chichester, 2000. Horwood Publishing for Institute of Mathematics and its Applications. Proceedings Second IMA Conference on Image Processing, De Montfort University, Leicester, UK, September 1998.
- [3] S. J. Sangwine and T. A. Ell. Hypercomplex Fourier transforms of color images. In *IEEE International Conference on Image Processing (ICIP 2001)*, volume I, pages 137–140, Thessaloniki, Greece, 7–10 October 2001. Institute of Electrical and Electronics Engineers.
- [4] S. J. Sangwine. Colour image edge detector based on quaternion convolution. *Electronics Letters*, 34(10):969–971, 14 May 1998.
- [5] S. J. Sangwine and T. A. Ell. Colour image filters based on hypercomplex convolution. *IEE Proceedings – Vision, Image and Signal Processing*, 147(2):89–93, April 2000.
- [6] S. J. Sangwine and T. A. Ell. Gray-centered RGB color space. In *Second European Conference on Color in Graphics, Imaging and Vision (CGIV 2004)*, Technology Center AGIT, Aachen, Germany, 5–8 April 2004. The Society for Imaging Science and Technology.
- [7] S. J. Sangwine, B. N. Gatscheni, and T. A. Ell. Vector amplification for color-dependent image filtering. In *IEEE International Conference on Image Processing (ICIP 2003)*, volume 2, pages 129–132, Barcelona, Spain, 14–17 September 2003. Institute of Electrical and Electronics Engineers.
- [8] Luca Lucchese and Sanjit K. Mitra. Filtering color images in the xyY color space. In ICIP2000 [10], pages 500–503.
- [9] T. A. Ell and S. J. Sangwine. Hypercomplex Wiener-Khintchine theorem with application to color image correlation. In ICIP2000 [10], pages 792–795.
- [10] *IEEE International Conference on Image Processing (ICIP 2000)*, Vancouver, Canada, 11–14 September 2000. Institute of Electrical and Electronics Engineers.