

Biologically inspired framework for spatial and spectral velocity estimations

Xuefeng Liang,^{1,*} Peter W. McOwan,¹ and Alan Johnston²

¹*School of Electronic Engineering and Computer Science, Queen Mary University of London, London E1 4NS, UK*

²*CoMPLEX & Cognitive, Perceptual and Brain Sciences, University College London, London WC1H 0AP, UK*

*Corresponding author: *xliang@dcs.qmul.ac.uk*

Received June 17, 2010; revised January 27, 2011; accepted February 3, 2011;
posted February 8, 2011 (Doc. ID 129451); published March 31, 2011

The multichannel gradient model (McGM) is an established, biologically plausible framework for the robust extraction of image velocity. Here we describe the McGM extension into color space and report the resulting performance improvement. Our new model, in contrast to existing approaches that process color channels separately, incorporates spectral energy measures to form a local description of the stimulus chromatic spatio-temporal structure from which we can recover both spatial and spectral velocities. We present a range of comparative experiments on synthetic and natural test data that demonstrate that our new method reduces errors and is more robust over a range of viewing environments. © 2011 Optical Society of America

OCIS codes: 330.4150, 330.5020.

1. INTRODUCTION

The recovery of the motion vector field (optical flow) from the time varying pattern of light on the retina is a fundamental human capability. Extracting optic flow is also a key problem in computer vision. So far, numerous studies have addressed this problem with gradient-based methods [1,2], phase correlation methods [3], region matching methods [4], and hierarchy methods [5,6]. Differential-based methods (e.g., Lucas–Kanade [7], Horn–Schunck [8]) underpin some of the most popular approaches to motion estimation and have been extensively researched. However, most of the existing differential models are only designed to work with gray-scale intensity images.

It is well known that color is a powerful cue that plays a key role in a broad range of computer vision applications, and color information can be incorporated to improve robustness in optic flow calculations. Ohta [9] first introduced color information into motion estimation. He used the additional color channel data to help resolve the aperture problem by solving an otherwise overdetermined matrix. Subsequently, he evaluated the method in RGB space with an extensive test bed of real world data [10] and found that his method did not always improve the accuracy of motion estimation. Barron and Klette [11] evaluated a similar color motion framework, but introduced additional features (weighting the importance of each color channel etc). They found the accuracy of optical flow estimation improved slightly, but incorporating the saturation channel had a deleterious effect for images encoded in HSV space. Many methods [12–15] for finding invariant color properties rely on a photometric model and on assumptions about the physical variables involved. For example, hue is known to be insensitive to surface orientation, illumination direction, intensity, and specular highlights, under white light illumination. Normalized RGB is a useful object measure for matte, dull surfaces illuminated by white light due to its insensitivity to the intensity of illumination. All these methods select two or

three conditioned channels, and use a least square or over-determined matrix technique to compute the velocity vector. However, we do not always have perfect white illuminated images in real-world applications. Zickler *et al.* [16] derived photometric invariants from illuminant-dependent subspaces of RGB space. This color subspace can be used to improve the estimation of optic flow. This strategy is particularly helpful for motion recovery in specular regions where gray-scale algorithms frequently perform poorly.

In this paper, we introduce a novel framework for discovering motion fields based on an assumption of constant color. It extends the series of Gaussian derivatives into color space and integrates them into the multichannel gradient model (McGM) [17,18]. Applied to gray-scale images initially, the original McGM model relies on image intensity consistency constraints using an enhanced gradient technique based on an interpretation of the computational function of receptive cell profiles found in the primary visual cortex. The basic idea of McGM is as follows. It is generally accepted that simple cells in the primary visual cortex act like spatio-temporal linear filters. If, at a position in the image, the ratio of the temporal derivative and spatial derivative is calculated, then local velocity is recovered in a simple formula $\frac{dx}{dt} = \frac{(\frac{\partial I}{\partial t})}{(\frac{\partial I}{\partial x})}$, where I is the image. However we can replace the image brightness with some other measure that does not change as the image moves. A Taylor expansion provides a very rich description of image structure at each point in the visual field and it also allows the approximation of the image brightness values at adjacent points in space and time. The Taylor representation requires a bank of linear filters, taking derivatives in two spatial directions, x and y , and in time t . We just replace the I , in the ratio above, with a vector containing the values of the various orders of derivative structure (mean level, slope, second-order variation, etc., up to order n), predicted within and aggregated over a cuboid volume around the point at which we measure the derivatives, and then condition the ratio of temporal and spatial derivatives of this elaborated representation using a

least squares approach. We now show it is possible to extend the original McGM further by incorporating color information, through modeling the spectral sensitive cortical cell response profiles and incorporating these new measures in a direct and consistent extension of our existing space–time framework. Thus, the framework can gracefully integrate space, color, and time information and then approximate the local spatio-temporal-spectral image structure by a truncated Taylor expansion where the color space derivatives provide additional motion measurement. These color derivatives enhance the robustness of motion estimation, especially when gray-level intensity fails. Rather than having the aforementioned color motion estimation algorithms process color in separate channels, we consider the color signal as a whole, incorporating spectral derivatives that allow us to introduce invariance to change in luminance level. This strategy allows us to more reliably recover motion than can be achieved using a simple color space transformation from RGB space. In addition, the spectral component of our model can be used for spectral velocity (color change) estimation. Thus, spatial and spectral velocity can be estimated through our single color McGM framework without any additional calculations. We are unaware of any other optic flow algorithm that is able to do this. We illustrate the improved performance of the proposed color McGM in motion estimation for both synthetic and natural color image sequences.

In this paper, we first present, in Section 2, the framework for representing color using color derivatives that underpins our proposed approach. Section 3 describes the method used for building the spatio-temporal-spectral Taylor expansion, and the mathematical framework of color image structure and spectral and spatial motion velocity extractions. This new color McGM model is then evaluated in Section 4 with a range of experimental stimuli and natural video sequences. Results show that the new approach outperforms the original McGM formulation in both gray-scale and various color spaces for the cases considered. Finally, we summarize and conclude the paper in Section 5.

2. COLOR REPRESENTATION

A. Perceptual Configuration of Color

One of the most remarkable and important insights in color appearance comes from the study of color opponency, which refers to a powerful psychological relationship between different hues [19]. While some mixtures of hues can coexist as a single color sensation, others cannot. For instance, orange is composed of red and yellow. However, we never experience a color sensation that is simultaneously red and green, similarly with blue and yellow. These two hue pairs, red–green and blue–yellow, are called *opponent colors*. This relationship was first championed by Hering [20] and was later firmly established by quantitative psychophysical and physiological studies [21,22]. Hurvich and Jameson’s hue-cancellation experiments resulted in curves for the yellow–blue and red–green opponent mechanisms that closely approximate the green (double lobed) and blue (triple lobed) curves shown in Fig. 1. The light–dark luminance mechanism, which can also be measured psychophysically [23,24], approximates the red (single lobe) curve in Fig. 1.

The wavelength distribution of the retinal input is encoded by three cone types (peaking at long, medium and short

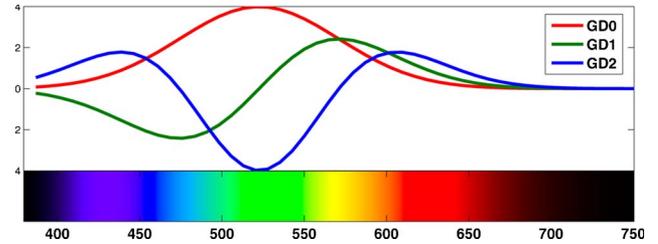


Fig. 1. (Color online) The zeroth, first, and second Gaussian derivatives along the wavelength axis are similar to the luminance, blue–yellow, and red–green weighting functions found in the human visual system [21–24].

wavelengths, i.e. red, green, and blue respectively) with broad action spectra. The hue-cancellation functions can be approximated by summing and differencing the action spectra [25]. Significantly, the relationship between these color opponent functions can also be approximated by a process of differentiation. We use a Gaussian function of wavelength to model the achromatic signal—the red curve shown in Fig. 1. The first derivative of this Gaussian function of wavelength corresponds to a blue–yellow opponent signal (the green curve shown in the figure), and the second derivative of this function corresponds to the red–green opponent signal (the blue curve shown in the figure).

B. Spectral Structure of Color

Taking a lead from the general form of the color opponent functions [21,22], we employ a Gaussian and its derivatives and formulate a scale-space framework [26] to allow a better approximation of the spectral energy distribution. Let $E(\lambda)$ be the reflected energy distribution of the incident light, where λ is the wavelength. Let $G(\lambda; \lambda_0, \sigma)$ be a Gaussian density at position (mean) λ_0 with spectral scale (standard deviation) σ . We say the function

$$M(\lambda_0; E, \sigma) = \int E(\lambda)G(\lambda; \lambda_0, \sigma)d\lambda \quad (1)$$

measures the spectrum E with an aperture $G(\lambda; \lambda_0, \sigma)$, and results in the integrated spectral energy M . The observed energy \hat{M} , at infinitely small perturbations, can be approximated by the second-order Taylor expansion in the neighborhood of λ_0 equal to

$$\begin{aligned} \hat{M}(\lambda_0 + \delta; E, \sigma) &= M(\lambda_0; E, \sigma) + \delta M_{\lambda}(\lambda_0; E, \sigma) \\ &\quad + \frac{\delta^2}{2} M_{\lambda\lambda}(\lambda_0; E, \sigma) + O(\delta^3), \end{aligned} \quad (2)$$

in which

$$\begin{aligned} M_{\lambda}(\lambda_0; E, \sigma) &= \int E(\lambda)G_{\lambda}(\lambda; \lambda_0, \sigma)d\lambda, \\ M_{\lambda\lambda}(\lambda_0; E, \sigma) &= \int E(\lambda)G_{\lambda\lambda}(\lambda; \lambda_0, \sigma)d\lambda, \end{aligned} \quad (3)$$

where the aperture functions G , G_{λ} , and $G_{\lambda\lambda}$ denote derivatives of the Gaussian with respect to λ , which correspond to the sensitivities shown in Fig. 1. Thus, the spectral measurements M , M_{λ} , and $M_{\lambda\lambda}$ are interpreted as the responses of intensity, yellow–blue, and red–green receptive fields, respectively.

In the color McGM, we will truncate the expansion to include up to first order in the basic functions. This allows us to differentiate again with respect to wavelength to reach second order [see Eqs. (4) and (6)]. This constraint respects the trichromatic nature of human vision [20]. The spectral filter parameters are fixed at $\lambda_0 \simeq 520$ nm and $\sigma \simeq 55$ nm [27]. Note that we do not separately consider spectral scale, spatial scale, and temporal scale, but integrate the spectral intensity and its first- and second-order derivatives with appropriate spatio-temporal derivatives together to form an integrated framework for color scale space. This underpins our proposed color McGM. Geusebroek *et al.* [27] use these same filters to define a bank of color invariants.

3. COLOR McGM

Neuromorphic systems emulate key elements of biological nervous systems by modeling their physical architectures and putative computations, and can provide extremely useful insights for building technical systems. McGM [17,18,28–30] follows the neuromorphic approach and is based on a gradient model [31] for motion estimation. It has proved very effective both in explaining observer performance [17,18,32], predicting visual experience in motion-based visual illusions, and providing a mathematical strategy that avoids computationally intensive operations, such as matrix inverse [7] or iterative methods [8,33], that are not biologically justified.

A. Color Image as a Taylor Expansion

It is generally accepted that simple cells in the primary visual cortex act like spatio-temporal linear filters. Following Koenderink and van Doorn [34] and Koenderink and Kappers [26], we consider these simple cells to approximate Gaussian derivatives of various orders. Blurring and differentiation of images can be accomplished by these differential operators. Thus, the outputs of appropriate simple cells can provide various orders of partial derivatives of the image, allowing a truncated Taylor approximation of the image in the neighborhood of a point in the space–time spectrum. A Taylor expansion provides a very rich description of image structure at each point in the visual field, allowing the approximation of image energy values at adjacent points in space, time, and spectral wavelength. The Taylor representation requires a bank of linear filters, taking derivatives in two spatial directions x and y , in time t , and in wavelength λ . The spatial filters extracting higher derivatives of the image are tuned to higher spatial frequencies. Using more terms allows more accurate representation of image structure. Johnston and Clifford [17] showed that the temporal filters of the human visual system can be characterized as causal differentials of a Gaussian in log time. The spectral filters [35] convert RGB values to intensity, yellow–blue, and red–green receptive field outputs that provide the measurement of spectral energy distributions and color invariants.

An image in a video sequence is usually expressed as $I(t)$. As we discussed previously, the spectral energy distribution, M , rather than the RGB value, denotes a color element in this paper. The spectral energy about a point $P = (x, y, t, \lambda)$ in the color image can be approximated locally from the Taylor expansion [36]:

$$\hat{M}_{(x,y,t,\lambda)}(p,q,r,s) = \sum_{i=0}^a \sum_{j=0}^b \sum_{k=0}^c \sum_{l=0}^d \frac{p^i q^j r^k s^l \partial^{(i+j+k+l)} E(x,y,t,\lambda)}{i!j!k!l!} \frac{1}{\partial x^i \partial y^j \partial t^k \partial \lambda^l}, \quad (4)$$

where (p, q, r, s) are the excursion parameters in two orthogonal spatial directions, time and spectrum. This expansion yields a set of gradient measures up to order a, b, c, d in the four dimensions.

Each term in the Taylor expansion can be treated as a piece of information, which then, taken together as a truncated Taylor expansion, represents the local image geometry. Using this mathematical formulation it is, therefore, possible to integrate the measures of space x, y , time t , and wavelength λ in a principled and unified way.

B. Spatial Velocity

Motion can be regarded as orientation in space–time. Suppose we have a vertical cyan bar moving continuously from left to right, sampled at three points in time [see Fig. 2(a)]. If we examine a space–time slice, then the motion of the bar is apparent from its orientation relative to the time axis [Fig. 2(b)]. Thus, motion results in the orientation/slope of the line segment ab relative to the axis [Fig. 2(c)].

The luminance or color content of the pattern simply serves to define a signal that remains invariant under translation. The relative changes in the signal with respect to space and time allow motion measurement. These changes can be represented by the terms in the Taylor expansion. Different order terms of the Taylor expansion encode different types of information, such as slopes (first-order change), curvature (second-order change), and elongation (third-order change) [34].

To reflect this geometric variation, for any given point (x, y, t, λ) , we group the terms of the same derivative order together in Eq. (5) to construct a vector \mathbf{k} , each component of which is a function of horizontal space x , vertical space y , time t , and wavelength λ , as well as the excursion parameters (p, q, r, s) . Specifically,

$$\mathbf{k}_{(x,y,t,\lambda)} = (k_0, k_1, \dots, k_n)^T, \quad (5)$$

where the numeric subscript denotes derivative order, $n = a + b + c + d$. For example, the first few terms are

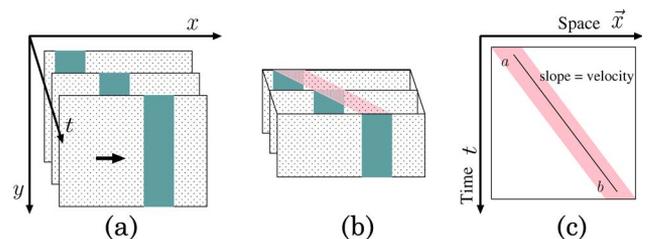


Fig. 2. (Color online) Motion as orientation in space–time. If we take a slice in the (\vec{x}, t) plane, we can see that the structure is oriented. So space–time orientation/slope indicates motion.

Zeroth Order Term

$$k_0 = E(x, y, t, \lambda);$$

First Order Term

$$k_1 = pE_x(x, y, t, \lambda) + qE_y(x, y, t, \lambda) \\ + rE_t(x, y, t, \lambda) + sE_\lambda(x, y, t, \lambda);$$

Second Order Term

$$k_2 = \frac{p^2}{2}E_{2x}(x, y, t, \lambda) + \frac{q^2}{2}E_{2y}(x, y, t, \lambda) + \frac{r^2}{2}E_{2t}(x, y, t, \lambda) \\ + \frac{s^2}{2}E_{2\lambda}(x, y, t, \lambda) + pqE_{xy}(x, y, t, \lambda) + prE_{xt}(x, y, t, \lambda) \\ + psE_{x\lambda}(x, y, t, \lambda) + qrE_{yt}(x, y, t, \lambda) + qsE_{y\lambda}(x, y, t, \lambda) \\ + rsE_{t\lambda}(x, y, t, \lambda);$$

Higher Order Term

⋮

We then differentiate each component image of the vector k corresponding to a given point (x, y, t, λ) with respect to space, time, and wavelength, ensuring that the result is truncated to the second order in each four-dimensional space, yielding four-dimensional functions that can be written as the entries of a matrix:

$$\mathcal{J} = Dk_{(x,y,t,\lambda)} = (k_x, k_y, k_t, k_\lambda) = \begin{pmatrix} k_{0,x} & k_{0,y} & k_{0,t} & k_{0,\lambda} \\ k_{1,x} & k_{1,y} & k_{1,t} & k_{1,\lambda} \\ \vdots & \vdots & \vdots & \vdots \\ k_{n,x} & k_{n,y} & k_{n,t} & k_{n,\lambda} \end{pmatrix}, \quad (6)$$

where D is the derivative operator, and subscripts denote the partial derivative and order. Note that one gets the same result if we differentiate the Eq. (6) with respect to x, y, t, λ ; or p, q, r, s . We can think of the (p, q, r, s) associated with the Taylor expansion as determining a parameter space that can be mapped into an n -dimensional space in which each coordinate axis delimiting the space serves to represent the value of the n th-order derivative of the local structure of the chromatic image surface [30]. We drop the reference to the excursion parameters here to simplify the notation. Each column of the matrix represents the Taylor expansion differentiated with respect to x, y, t , or λ and each row contains differentials of the same order. In practice, each derivative of the color image, represented by a component in the matrix in Eq. (6), is computed by integration with the appropriate Gaussian derivative operator, following the ideas introduced in Subsection 2.B.

If we regroup each column in matrix the matrix in Eq. (6) into two sets with respect to the derivative of the spectrum, the components represented by all derivatives correspond in an obvious way to the vectors shown in the vectors in Eqs. (7) and (8) that involve the expected spectral information. For instance, k_x is

$$\underbrace{(x_1y_0t_0\lambda_0, \dots, x_{i+1}y_jt_k\lambda_0)}_{\text{Intensity}(D_{S_0})}, \underbrace{x_1y_0t_0\lambda_1, \dots, x_{i+1}y_jt_k\lambda_1)}_{\text{Yellow-Blue}(D_{S_1})}, \quad (7)$$

where D_S denotes the spectral derivative, and subscripts $i = 0, \dots, 5, j, k = 0, 1$ are two spatial and the temporal deriva-

tive orders, respectively. It is similar for k_y and k_t . For k_λ , we have

$$\underbrace{(x_0y_0t_0\lambda_1, \dots, x_iy_jt_k\lambda_1)}_{\text{Yellow-Blue}(D_{S_1})}, \underbrace{x_0y_0t_0\lambda_2, \dots, x_iy_jt_k\lambda_2)}_{\text{Red-Green}(D_{S_2})}. \quad (8)$$

These vectors underpin our color motion estimation framework and the color invariants.

From these measurements, we compute a simple matrix product:

$$\mathcal{J}^T \mathcal{J} = \begin{pmatrix} k_x \cdot k_x & k_x \cdot k_y & k_x \cdot k_t & k_x \cdot k_\lambda \\ k_y \cdot k_x & k_y \cdot k_y & k_y \cdot k_t & k_y \cdot k_\lambda \\ k_t \cdot k_x & k_t \cdot k_y & k_t \cdot k_t & k_t \cdot k_\lambda \\ k_\lambda \cdot k_x & k_\lambda \cdot k_y & k_\lambda \cdot k_t & k_\lambda \cdot k_\lambda \end{pmatrix}. \quad (9)$$

Specifically, each component in the matrix in Eq. (9) $k_X \cdot k_Y = \sum_{i=0}^n k_{i,X} k_{i,Y}$, where $X, Y \in \{x, y, t, \lambda\}$.

This matrix is then integrated over a spatio-temporal-spectral region [30]:

$$C = \int_{-s}^s \int_{-r}^r \int_{-q}^q \int_{-p}^p \mathcal{J}^T \mathcal{J} dx dy dt d\lambda \\ = \begin{pmatrix} x \cdot x & x \cdot y & x \cdot t & x \cdot \lambda \\ y \cdot x & y \cdot y & y \cdot t & y \cdot \lambda \\ t \cdot x & t \cdot y & t \cdot t & t \cdot \lambda \\ \lambda \cdot x & \lambda \cdot y & \lambda \cdot t & \lambda \cdot \lambda \end{pmatrix}. \quad (10)$$

From the elements of the matrix in Eq. (10), image speed can be obtained in two orthogonal spatial directions by computing the well-conditioned ratios:

$$\nu_x = \left(\frac{\lambda \cdot t + x \cdot t}{\lambda \cdot x + x \cdot x} \right), \quad \nu_y = \left(\frac{\lambda \cdot t + y \cdot t}{\lambda \cdot y + y \cdot y} \right). \quad (11)$$

Although individual orders may locally be zero, the aggregated denominator component generally retains a nonzero value for the following reason: where structure exists to be measured, odd and even value derivatives will have maximum and minimum values at different points in their cycle, e.g., a zero in a first-order derivative will normally not coincide with a zero in the second-order derivatives.

If we merely consider the first term of the vector k [Eq. (5)], ν_x in Eq. (11) can be simplified to

$$\nu_x = \frac{\frac{\partial E}{\partial \lambda} \frac{\partial E}{\partial t} + \frac{\partial E}{\partial x} \frac{\partial E}{\partial t}}{\frac{\partial E}{\partial \lambda} \frac{\partial E}{\partial x} + \frac{\partial E}{\partial x} \frac{\partial E}{\partial x}},$$

from which one can more clearly see that this ratio computes velocity.

The vectors in Eqs. (7) and (8) show that x, y, t contain the intensity and yellow-blue information of the color signals, and λ contains both opponent color pairs. These additional spectral derivative terms guarantee the robustness of motion estimation. As discussed above, velocity can be considered as orientation in space-time. For any image descriptor, we recover velocity from the ratio of change with respect to space and time. Our gradient technique provides a least squares estimate of this ratio based on (D_{S_0}, D_{S_1}) and (D_{S_1}, D_{S_2}) of the vector k [see the vectors in Eqs. (7) and (8)]. We plot these terms schematically in Fig. 3. Note that the least squares estimate of the ratio depends mainly upon the derivatives with the large values. Near-zero terms contribute little to

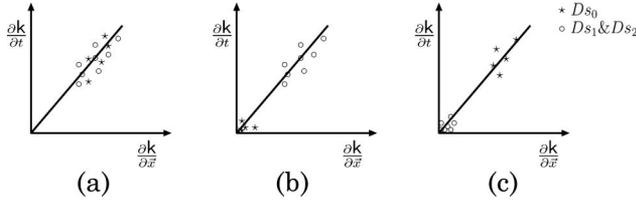


Fig. 3. The color McGM primitives have intensity (Ds_0), yellow–blue (Ds_1) and red–green (Ds_2) for velocity estimation in space-time. The coordinates of the points represent the values of matched pairs of terms combined through the least squares calculation that differ only in that one has an extra temporal derivative and one has an extra spatial derivative. (a) illustrates the normal situation of a moving object leading to both intensity and color changes. So, both intensity (Ds_0) and color opponency (Ds_1, Ds_2) contribute the orientation estimate. (b) illustrates the situation with different colors but a constant intensity. Thus, intensity changes are small and do not make a substantial contribution. The orientation is mainly estimated from chromatic change. (c) illustrates the situation for constant color but varied local luminance. In this case, the chromatic change is small. The orientation estimate relies heavily on the intensity changes.

the estimate. Normally, the change in an image involves both intensity and color. As can be seen from the definition of speed [Eq. (11)], all three groups of terms, Ds_0 , Ds_1 , and Ds_2 contribute to the least squares calculation [see Fig. 3(a)]. However, some of them will fail in the velocity estimation in two extreme cases. In the first case, a loss of information may occur in the RGB to gray-scale transformation, such as, the same gray-value intensity may arise from different colors. When the gray levels of foreground and background are relatively constant across space and time, components involving Ds_0 will be small [see Fig. 3(b) and the third experiment in Subsection 4.B for illustration]. Fortunately, our model still can estimate the speed using components involving Ds_1 and Ds_2 instead. In the second case, suppose the white light from a spotlight is moved over a colored background; then the local change of intensity will be much more salient than the change in color. The least squares estimate will, therefore, be dominated by components involving Ds_0 [see Fig. 3(c)]. Components involving Ds_1 and Ds_2 will not make a substantial contribution. The calculation ensures the resulting speed estimate will be dominated by the larger measurements and insensitive to null measurements, leading to robust motion estimation [28]. In addition, combination of just the Ds_1 and Ds_2 derivatives can measure chromatic change independent of any change in overall luminance.

We further consider computing these speed measures for a range of special orientations corresponding to the orientation columns found in the primary visual cortex [18], and introduce the speed $\hat{\nu}$ and the inverse speed $\hat{\nu}^{-1}$ measures. Please refer to Appendix A for details.

C. Spectral Velocity

The components in matrix (10) also allow us to measure the color change. From this we can estimate spectral velocity. In this formulation the spectral velocity is defined as the rate of change of hue at a spatial position \hat{x}_i , and this measure can be easily derived from the spectral element by computing the well-conditioned ratio:

$$\nu_s = \left(\frac{\lambda \cdot \dot{\lambda}}{\lambda \cdot \lambda} \right). \quad (12)$$

Note that wavelength is a one-dimensional distribution and so the direction of ν_s is only positive or negative. We define hue changes toward a longer wavelength to be positive and toward a shorter wavelength to be negative. Since Eq. (12) does not include changes in the local image representation across space and time, the spatial gradients and their orientation do not affect the calculation of spectral velocity. Thus, spectral velocity can be obtained independently of spatial velocity.

4. EXPERIMENTS

To demonstrate the performance of the enhanced color McGM, two aspects of the proposed framework have been experimentally analyzed on both synthetic and natural test data. They are spectral velocity and spatial velocity estimation.

A. Spectral Velocity Evaluation

The experiments on spectral velocity estimation aim to demonstrate the quality of the proposed model with respect to color change measurement and color constant measurement under changing illumination. We first created a test sequence of paired frames with gradually increasing wavelength from 430 to 610 nm, calibrated with respect to the color matching functions of CIE_1964 [25]. The increment step is 7.2 nm (see Fig. 4). All pairs of adjacent frames are processed according to Eq. (12). Because the color McGM model input is an RGB image, the result of calculation must be mapped back to wavelength. Figure 5 then shows the model's calculated magnitude of spectral velocity for these test stimuli. The average absolute speed error (ASE) is measured at 0.1277 nm, and the standard deviation (Std.) is 0.1398. Color change directions are all correctly predicted. Because our spectral Gaussian derivative is designed for an idealized universal RGB camera, the performance could be further improved by optimizing the model with respect to the camera input. The appropriate CIE color matching function selection should also be taken into account to improve results.

A more important test is to evaluate the color constant measurement under varied illumination conditions. Color reference swatches (red, green, and blue) are placed in an average daylight environment (color temperature 6500 K) [37]. The varied illumination conditions are simulated by changing the lens aperture of the camera but fixing all other parameters. The camera used to create these test sequences was a Nikon D40 (AF-S 18-55). The camera shutter speed was 1/30 s, apertures were from $f/5.6$ to $f/18$, the ISO was 200, and white

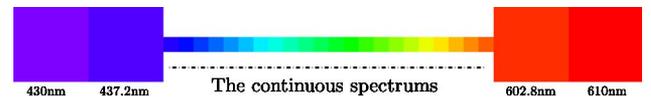


Fig. 4. (Color online) Synthetic color image sequence (25 paired frames) with gradually increasing wavelength.

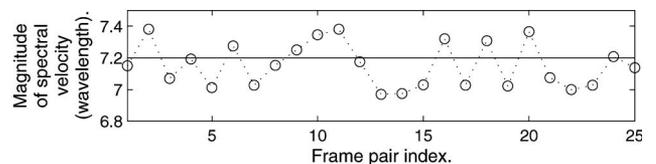


Fig. 5. Model calculated speed of wavelength change between adjacent frame pairs.

Table 1. Computed ν_s and Hue Value Change of Color References Under Different Illuminations

| Aperture Change | Red | | Green | | Blue | |
|-----------------|---------|---------|---------|---------|---------|--------|
| | ν_s | hue | ν_s | hue | ν_s | hue |
| $f/5.6-f/6.3$ | -0.3634 | -0.5596 | 0.3584 | -0.9596 | 0.4038 | 1.2467 |
| $f/6.3-f/7.1$ | -0.2037 | -1.3169 | 0.3245 | 2.2829 | 0.1972 | 2.2034 |
| $f/7.1-f/8$ | -0.0916 | -1.9479 | 0.1673 | 0.6472 | 0.1089 | 2.0654 |
| $f/8-f/9$ | -0.0723 | -1.2776 | 0.1382 | -0.4371 | 0.0691 | 0.8703 |
| $f/9-f/10$ | -0.0791 | -2.1564 | 0.0958 | 0.2623 | 0.0727 | 2.0516 |
| $f/10-f/11$ | -0.0819 | -0.9497 | 0.0883 | 1.9152 | 0.0709 | 1.5833 |
| $f/11-f/13$ | -0.1016 | -1.0196 | 0.1228 | 0.5586 | 0.0838 | 1.8325 |
| $f/13-f/14$ | -0.0956 | -0.7760 | 0.0908 | 1.6864 | 0.0678 | 1.5664 |
| $f/14-f/16$ | -0.0778 | -0.7253 | 0.0689 | 2.5314 | 0.0613 | 0.6220 |
| $f/16-f/18$ | -0.0361 | -0.4780 | 0.0636 | 0.9293 | 0.0347 | 0.3754 |

balance was daylight. When applying Eq. (12) to image pairs with two adjacent aperture settings, the ideal calculated result should be equal to zero. The aforementioned algorithms [11–15] work in HSV space to recover color constancy [hue (H) channel]. To allow comparison, the proposed McGM framework is evaluated against color constancy in the H channel alone. Results are given in Table 1, where the result values are degree of hue [0, 360].

As expected, the proposed framework offers a better color constancy measurement than that afforded by use of the H channel alone. But the model performs less well under high illumination, as the model has no prior knowledge of the specific camera characteristics. For the H channel, as the commonly used conversion from RGB space to HSV space [38] does not use a white reference, the hue value is unlikely to represent the true color. The real constant hue can only be recovered by using the CIE color matching functions [25] with an accurate white reference. However, it is difficult to obtain the precise white reference from a normal camera, especially from a self-adjusting video camera. We believe from the data that, under the experimental constraints imposed by the equipment used, our model works reasonably effectively.

B. Spatial Velocity Evaluation

The original McGM can only estimate the spatial velocity, and was previously tested against other standard gradient-based models in the literature and found to outperform those for a series of test sequences [29]. In this paper, to evaluate the performance of the enhanced color McGM, four different comparison experiments were performed. We first created an image sequence where each image is formed from a shifted color (hue) sine wave represented in HSV color space with an image size of 201×201 . In this test stimulus, the saturation and value channels are set to maxima [see Fig. 6(a)]. Spatial frequency is 2 cycles/image, speed is 4 pixels/frame, direction is 30° to top left. The scaling of filters to images is based on known biology [17]. As we can recover a dense representation of optic flow (velocity vector) at all points in the stimulus, rather than use a spatially sampled arrow diagram, we show calculated pixel speed as an intensity with respect to the surrounding zero speed border, and output direction of motion as a color with respect to the surrounding color wheel. In Fig. 6(a), the upper two images are two consecutive frames of the shifting color sine wave video; the bottom left image shows the computed speed (magnitude); the bottom right image shows the computed direction of motion. These color

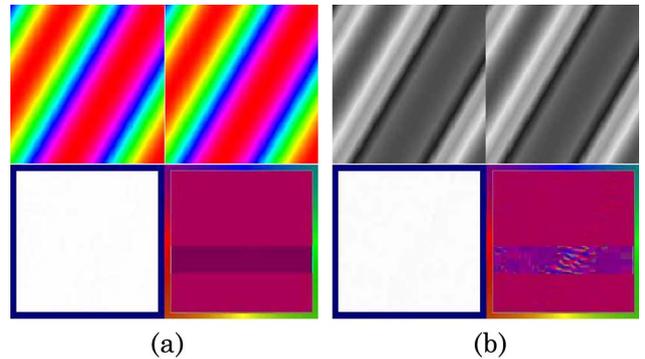


Fig. 6. (Color online) Performance comparison of our proposed color McGM with the original McGM on grating sequences. The selected viewing rectangles are at the same position in both sequences; the direction results within these rectangular viewing windows are contrast enhanced for illustrative purposes to show the difference in performance. (a) shows the result from a color image sequence using the color McGM and (b) shows the result from a directly converted grayscale image derived from the color sequence using the original McGM.

images were converted to corresponding gray-scale images to allow comparison [see Fig. 6(b)]. The color McGM and original version were then evaluated on color and gray test data, respectively. Figure 6(a) shows that the calculated speed and direction are uniform for the color McGM. By contrast, there are small errors within the very smooth regions of the direction results in Fig. 6(b). We select rectangles at the same position in both direction results, and stretch the contrast to make a more visible comparison. The numerical comparison of model performance can be found in Table 2.

In the first experiment, the color was smoothly modulated by a sinusoidal function. To examine performance for more noisy data, the second experiment employed color random dot kinematograms (RDKs) with the stimulus parameter speed, 2 pixels/frame; direction, up to down; image size, 300×300 . These comprise a random noise background with a block of noise texture rigidly translated between frames. Because the color and brightness change sharply between two adjacent dots, motion estimation is more challenging, yet human observers report sharp delineations of the moving patch [39]. Figure 7 shows that the color McGM outperforms the original version for this stimulus. The new model can compute speed and direction more accurately than the original. This could be explained by that fact that constant color provides a more stringent constraint than constant brightness alone. Table 2 gives the detailed comparison.

In the above experiments, different colors have different gray values. Therefore, intensity-based methods also can sense the motion. We created an image sequence where a

Table 2. Numerical Comparison of the Color McGM and the Original McGM^a

| | ASE | SE Std. | ADE | DE Std. |
|---------------|--------|---------|--------|---------|
| Color grating | 0.0016 | 0.0022 | 0.0581 | 0.0368 |
| Gray grating | 0.0037 | 0.0054 | 1.6245 | 2.0119 |
| Color RDK | 0.0829 | 0.1127 | 2.9304 | 4.4763 |
| Gray RDK | 0.1250 | 0.1808 | 6.1258 | 10.1934 |
| Color traffic | 0.0171 | 0.0145 | 2.7037 | 2.8139 |
| Gray traffic | 0.1039 | 0.2938 | 5.5821 | 10.7043 |

^aASE, average absolute speed error; SE, speed error; ADE, average absolute direction error; DE, direction error.

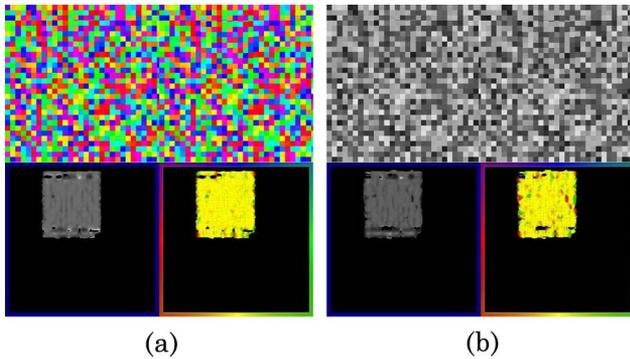


Fig. 7. (Color online) Performance comparison of the proposed color McGM and the original McGM on RDKs. (a) is the result of color RDK using the color McGM. (b) is the result of the corresponding gray-scale RDK using the original McGM.

cyan bar (R:20, G:228, B:235) moved from right to left on a background (R:26, G:230, B:212). The foreground and background colors are different but they result in the same gray value: 167. It is impossible to see motion on the basis of gray-level intensity, because there is no difference on the gray scale. Figure 8(b) shows the new color McGM still works well. The upper two images are the first and 23rd frames of the video sequence; the velocity results of the whole sequence are shown at the bottom. Moreover, when Eq. (12) was applied on the image pair in Fig. 8(a), our model measures a 5.087 nm difference between them.

In the fourth experiment, the stability of the proposed method was investigated with respect to varying illumination. We gradually reduced the values in the saturation and value channels, but kept the hue value unchanged (see Fig. 9). Many computer vision approaches become problematic when the image darkens. This issue is also reflected in the performance of our model; here the computational reason for reduced performance is that the sampling in the luminance and color domain is becoming coarser, which effectively increases the noise. In Fig. 10, we show that both speed and direction error standard deviations increase with decreasing illumination. However, the color McGM is seen to outperform the original McGM for certain illumination conditions, and also degenerates more gracefully with reducing illumination. This confirms our finding from the second experiment in Subsection 4.A: our color constancy renders the underlying method more accurate with respect to changes of illumination. In both cases, the mean computed speed and direction are relatively robust.

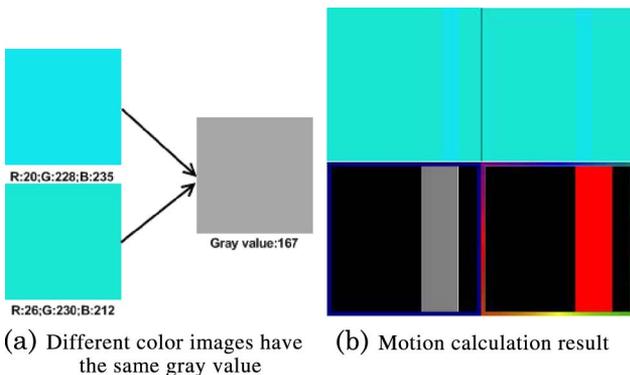


Fig. 8. (Color online) Performance evaluation of the proposed color McGM on image patterns with different colors that result in the same gray value.

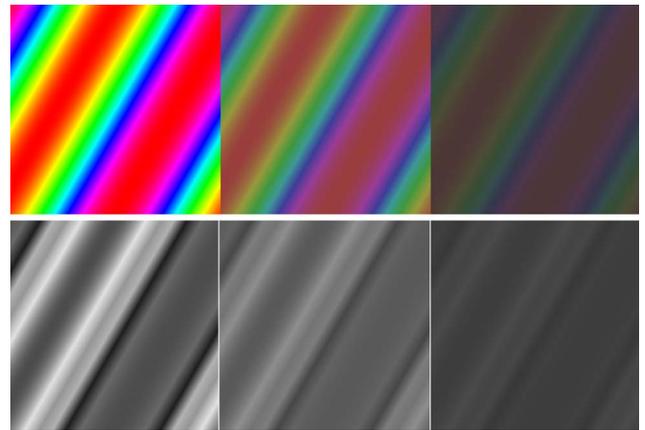


Fig. 9. (Color online) Grating illumination is gradually reduced. The top images show a gradual reduction in illumination color; the bottom images show the corresponding reductions in gray-scale images.

We believe this shows that the incorporation of color measurements in our new model provides for a more robust strategy than simply relying on the brightness.

Finally, we evaluate our algorithm on an important class of input; real video data. The upper two images of Fig. 11(a) are frames 45 and 47 of the color traffic scene 1, with a single moving vehicle. The computed results are shown below. Figure 11(b) shows analogous gray-scale frames and their computed result using the basic McGM approach. As one can see, the calculated velocity using the color McGM is more accurate with a small standard deviation and error (see Table 2). By observation, we can see the shape of the car is well preserved. Figure 12 demonstrates the model's performance for a more complicated traffic scene (scene 2). As compared to the color McGM, the gray-scale-based model delivers more error and a larger standard deviation. Some details should be noted: the color McGM calculates the velocity within uniformly moving areas in a more coherent way. Moving areas appear more uniform, indicating that all points in the object area are the same direction.

C. Comparison with Other Methods

Most color-based optical flow methods [10–15] apply the principle of the gray-scale-based framework in each channel of the corresponding color space, and then derive the final motion vectors by averaging the results of all channels or by summing weighted contributions. This often pragmatic approach is contrary to the color McGM presented here, which employs a principled mechanism to combine the measures and integrates space, color, and time information. As the original McGM was previously tested against other gray-scale-based models in the literature [29], here we evaluate the color McGM across some appropriate set of color space models [RGB, normalized RGB (NRGB), HSV, and spherical (r, ϕ, θ)]. Because the arithmetic chromaticity space usually performs a bit better than the geometric chromaticity space, we use arithmetic mean for NRGB in this paper. First, we consider the color traffic scenes presented in Subsection 4.B. The corresponding results in terms of the absolute error and error standard deviation are listed in Table 3. As can be observed, the HSV space performs worst. The reason for this might be that it is transformed from RGB, and only the H channel supports the color constancy assumption. Normalized RGB representation performs at a level very close to the color McGM because

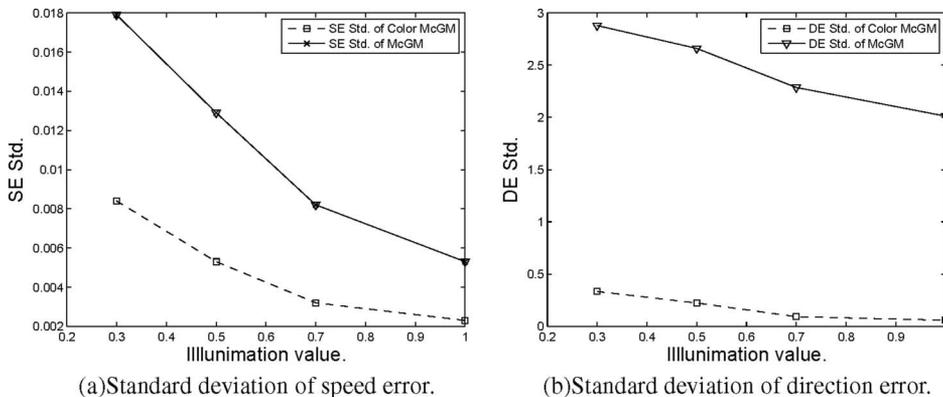


Fig. 10. Comparing degenerations of the color and noncolor McGM while illumination decreases. (a) Standard deviations of speed error increase with decreasing illumination. (b) Standard deviations of direction error increase with decreasing illumination. Note that the color McGM error increases more slowly than does the corresponding noncolor graph.

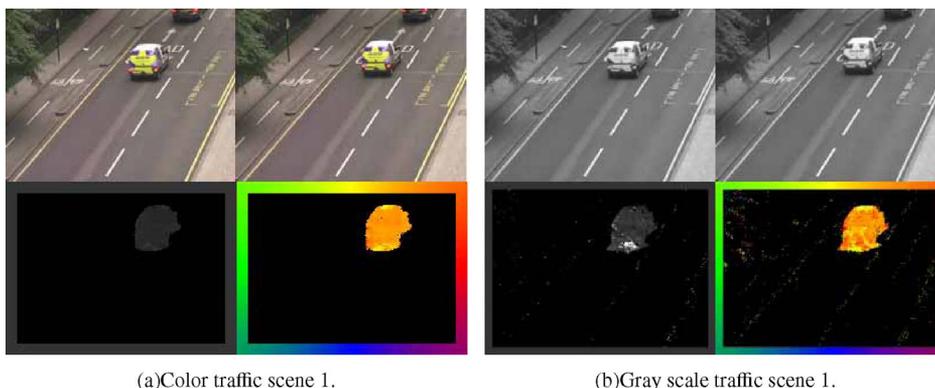


Fig. 11. (Color online) Performance comparison of the proposed color McGM and the original McGM on benchmark traffic scene data. (a) shows the result for car velocity using the color McGM. (b) is the result of using original McGM on the corresponding gray-scale video sequence and shows clearly a greater level of noise in the output.

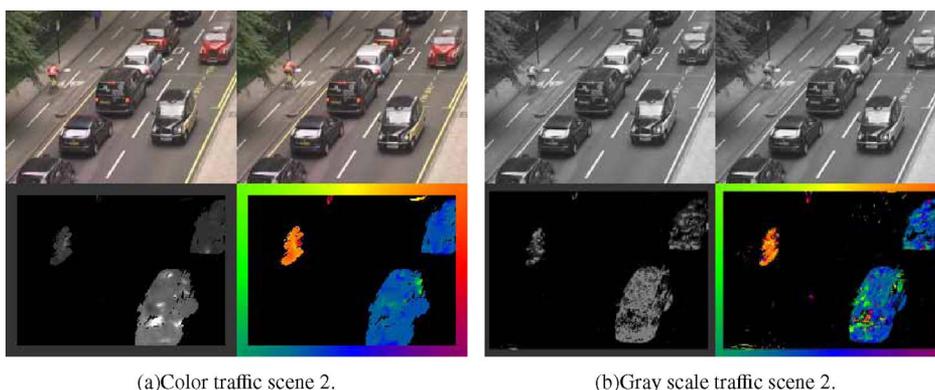


Fig. 12. (Color online) Performance comparison of more complicated traffic scene. (a) is the result using the color McGM. (b) is the result using the original McGM. Note the cars in the left lane are static.

Table 3. Comparison of the Color McGM and Applying Original McGM in the Different Color Spaces for the Color Traffic Scene^a

| | ASE | SE Std. | ADE | DE Std. |
|-------------------|--------|---------|--------|---------|
| Color McGM | 0.0171 | 0.0145 | 2.7037 | 2.8139 |
| McGM in RGB | 0.0728 | 0.0967 | 3.9304 | 5.7630 |
| McGM in NRGB | 0.0179 | 0.0154 | 2.7737 | 2.9117 |
| McGM in HSV | 0.1081 | 0.3048 | 5.6227 | 11.0031 |
| McGM in Spherical | 0.0269 | 0.0244 | 3.0547 | 3.9025 |

^aIn each color space, the final motion vectors are the average of original McGM results in all channels.

Table 4. Comparison of the Color McGM and Applying Original McGM in the Different Color Spaces for the Shifted Color Sine Wave with Illumination Changes

| | ASE | SE Std. | ADE | DE Std. |
|--------------------------------------|--------|---------|--------|---------|
| Color McGM | 0.0032 | 0.0051 | 0.6115 | 0.5068 |
| McGM in RGB | 0.0378 | 0.0499 | 13.341 | 19.0453 |
| McGM in NRGB | 0.0041 | 0.0064 | 0.7003 | 0.5917 |
| McGM in HSV(hue) | 0.0121 | 0.0182 | 1.9245 | 3.0139 |
| McGM in Spherical (ϕ, θ) | 0.0036 | 0.0054 | 0.6457 | 0.5095 |

the normalization strategy yields an invariant expression. Spherical space representation performs better than HSV, but worse than NRGB because ϕ and θ are invariant with respect to shadow and shading, but r is not.

To investigate the color invariant properties of the color McGM and corresponding color invariants in each color space, we test them on the shifted color sine wave with changing illumination. Table 4 shows that the color McGM, the color invariants (ϕ, θ) in spherical space, and the NRGB achieve expected good performance. But the standard RGB completely failed under varying illumination. As the hue measure in HSV is a single channel, the computation error may result in worse performance than for the spherical space and NRGB. Both experiments provide proof that the spectral component in the color McGM can measure the color invariant, and this can lead to it outperforming the velocity estimation of other related methods. Although the NRGB and spherical space (ϕ, θ) also achieve a similar spatial velocity estimation performance, they cannot estimate the spectral velocity in a single framework. However, our new color McGM can do this.

D. Computational Speed Test

Computational speed is an issue in practical applications. To explore this, we compared the computing time of the expanded color McGM against the coarse-to-fine iterative method of a leading rival approach [33]. The evaluation was performed on a platform comprising a Intel Duo-Core 2.00 GHz, 2 Gbytes RAM, Windows XP, and MATLAB R2008b. The input images of size 300×300 pixels were first grouped into different moving patterns (global motion, local motion, etc.). We used $m = 24$ orientations for the color McGM (in line with the biological data on the resolution of orientation columns in the visual cortex [17,18]). For the color McGM, the computing time was around 6.3 s on average. In contrast, the method in [33] employs a coarse-to-fine iteration to achieve similar accuracy using 11–20 iterations (depending on the input images), with each iteration taking 1.2 s on average. Therefore, the color McGM can achieve a reasonable accuracy in competitive time scales with no iteration. Indeed, the method is embarrassingly parallel, so how fast it runs scales linearly with the amount of computing power available. So, in practice, the color McGM is more effective than an iterative algorithm that cannot be parallelized because each iteration depends upon the previous computation. In the future, it may be possible to improve the color McGM algorithm time performance further by using specialized hardware, as there currently exists a real-time implementation of the original McGM deployed on a field-programmable gate array device [40].

5. CONCLUSIONS

In this paper, a novel biologically inspired color neuromorphic motion estimation approach, the color McGM, is proposed. This model reflects known structures in the primary visual cortex, and represents the local spectral structure using a truncated Taylor series expansion in space, time, and spectrum. The terms of the Taylor expansion are measured by a series of Gaussian derivative operators, modeling cortical simple cells. The expansion terms with spectral derivatives include not only intensity, but also yellow–blue and red–green

information, which directly underpins the proposed model. This delivers robustness with respect to variation in stimulus luminance level. Spectral velocity (color change) also inherits this useful robustness property when incorporated within our single framework color McGM. Other computational methods find this type of robustness challenging. In this paper, we have shown how we can incorporate this spectral structural descriptor into the existing successful McGM framework. In effect this allows the new color McGM to operate on the basis of constant color rather than the constant brightness constraint of the earlier model. In comparison to most existing color optical flow estimation methods that process color channels separately, our framework considers color information as a single entity that encodes the spectral structure of the color signal. We examined our new model's performance in two measures: spectral velocity and spatial velocity. The results and analysis presented show that the color McGM provides for a much better recovery of motion, even in noisy environments, and also degenerates more gracefully than the non-color version. In addition, the biologically inspired nature of the new model precludes the need for commonly employed coarse-to-fine iterative steps, and is, therefore, arguably more computationally efficient.

APPENDIX A: EXTRACTING SPEED AND INVERSE SPEED

It is convenient to introduce a notation for speed \hat{v} and inverse speed \check{v} vectors, $\nu = (\nu_{\parallel}, \nu_{\perp})$, whose components are parallel speed and orthogonal speed corresponding to the primary direction. This vector is computed at m different orientations θ around a point in the image. We do this by rotating our coordinate system (which also rotates the derivative filters) to the m directions. The raw measures of speed ν_x and ν_y are also conditioned by including the measures of the image structure $\frac{x \cdot y}{x \cdot x}$ and $\frac{x \cdot y}{y \cdot y}$ that are infinite for directions parallel to isobrightness contours. However, we can define well-conditioned directional speed vectors:

$$\begin{aligned}\hat{v}_{\parallel} &= \sqrt{\frac{2}{m}} \left[\nu_x \left(1 + \left(\frac{x \cdot y}{x \cdot x} \right)^2 \right)^{-1} \right], \\ \hat{v}_{\perp} &= \sqrt{\frac{2}{m}} \left[\nu_y \left(1 + \left(\frac{x \cdot y}{y \cdot y} \right)^2 \right)^{-1} \right],\end{aligned}\quad (\text{A1})$$

in which ν_x and ν_y , the horizontal and vertical velocities in Eqs. (11) and (A1), are normalized by a factor depending on the number of directions, m . By including multiplications involving terms computing the orientation of image structure as a function of direction, e.g., $\frac{x \cdot y}{x \cdot x}$, we can ensure that, for one-dimensional spatial stimuli, Eqs. (A1) deliver a sinusoidal function of direction, the amplitude of which is directly related to speed.

We can also calculate inverse speed, \check{v} , from the elements in the matrix in Eq. (10):

$$\check{v}_{\parallel} = \sqrt{\frac{2}{m}} \left(\frac{\lambda \cdot x + t \cdot x}{\lambda \cdot t + t \cdot t} \right), \quad \check{v}_{\perp} = \sqrt{\frac{2}{m}} \left(\frac{\lambda \cdot y + t \cdot y}{\lambda \cdot t + t \cdot t} \right). \quad (\text{A2})$$

The inverse speed is evaluated using different terms from those used to compute speed, and so constitutes an additional measurement.

Koenderink and van Doorn [34] showed that a velocity field can be decomposed into a translation component and differential components: divergence and curl, plus two components of affine shear. For a pure translation, the sum of both the components of \hat{v} over 2π rad will be zero. A local divergence will result in a nonzero sum of \hat{v}_{\parallel} and a local curl will result in a nonzero sum of \hat{v}_{\perp} over 2π rad. To recover the translation component and remove the differential components, we can force the integral of the directional speed functions to be zero by extracting the fundamental Fourier coefficients. This is achieved by projecting onto fiducial sine and cosine functions. We construct normalized cosine and sine vectors:

$$\mathbf{F}(\theta) = (F_{\parallel}(\theta), F_{\perp}(\theta)) = \sqrt{\frac{2}{m}}(\cos(\theta), \sin(\theta)). \quad (\text{A3})$$

This matrix forms both a fiducial reference frame, in terms of angle θ , for the computation of direction of motion and allows for the extraction of the fundamental Fourier coefficients of the directional speed functions. Speed squared is computed as a ratio of determinants:

$$S^2 = \frac{\begin{vmatrix} \hat{v}_{\parallel} \cdot F_{\parallel} & \hat{v}_{\parallel} \cdot F_{\perp} \\ \hat{v}_{\perp} \cdot F_{\parallel} & \hat{v}_{\perp} \cdot F_{\perp} \end{vmatrix}}{\begin{vmatrix} \hat{v}_{\parallel} \cdot \check{v}_{\parallel} & \hat{v}_{\parallel} \cdot \check{v}_{\perp} \\ \hat{v}_{\perp} \cdot \check{v}_{\parallel} & \hat{v}_{\perp} \cdot \check{v}_{\perp} \end{vmatrix}}, \quad (\text{A4})$$

where, for example, $\hat{v}_{\parallel} \cdot F_{\parallel}$ is the scalar product of the first column of \hat{v} and the second column of F . The denominator takes the value of 1 for rigid motion of simple patterns and can vary from point to point in the image to compensate for conditions in which the computed speed and inverse speed are not exact inverses. The denominator can be zero when $\hat{v}_{\parallel} = c\hat{v}_{\perp}$, which would be the case for a pure divergence; however, if this relation holds, the numerator is also zero and, as in the similar situation noted above, the indeterminacy is resolved by rule. It is the assumption that direction is coded in the visual system as a phase angle by pairs of cells encoding the projection on fiducial sine and cosine functions, respectively. However, for the purpose of the simulations, direction is computed explicitly as

$$\Theta = \tan^{-1} \frac{(\hat{v}_{\parallel} + \check{v}_{\parallel}) \cdot F_{\parallel} - (\hat{v}_{\perp} + \check{v}_{\perp}) \cdot F_{\perp}}{(\hat{v}_{\parallel} + \check{v}_{\parallel}) \cdot F_{\perp} + (\hat{v}_{\perp} + \check{v}_{\perp}) \cdot F_{\parallel}}. \quad (\text{A5})$$

ACKNOWLEDGMENTS

We thank UK Engineering and Physical Sciences Research Council (EPSRC) for funding the work.

REFERENCES

- J. W. Brandt and W. Jonathan, "Improved accuracy in gradient-based optical flow estimation," *Int. J. Comput. Vis.* **25**, 5–22 (1997).
- P. C. Chung, C. L. Huang, and E. L. Chen, "A region-based selective optical flow back-projection for genuine motion vector estimation," *Pattern Recognit.* **40**, 1066–1077 (2007).
- E. D. Castro and C. Morandi, "Registration of translated and rotated images using finite Fourier transforms," *IEEE Trans. Pattern Anal. Machine Intell.* **PAMI-9**, 700–703 (1987).
- D. Kalivas and A. Sawchuk, "A region matching motion estimation algorithm," in *Comput. Vision Graphics Image Process. Image Underst.* **54**, 275–288 (1991).
- H. Liu, R. Chellappa, and A. Rosenfeld, "Fast two-frame multi-scale dense optical flow estimation using discrete wavelet filters," *J. Opt. Soc. Am. A* **20**, 1505–1515 (2003).
- T. Amiaz, E. Lubetzky, and N. Kiryati, "Coarse to over-fine optical flow estimation," *Pattern Recognit.* **40**, 2496–2503 (2007).
- B. Lucas and T. Kanade, "An iterative image registration technique with an application to stereo vision," in *IJCAI'81 Proceedings of the 7th International Joint Conference on Artificial Intelligence* (Morgan Kaufmann, 1981), Vol. 2, pp. 674–679.
- B. Horn and B. Schunck, "Determining optical flow," *Artif. Intell.* **17**, 185–203 (1981).
- N. Ohta, "Optical flow detection by color images," in *Proceedings of IEEE International Conference on Image Processing* (IEEE, 1989), Vol. 4, pp. 801–805.
- N. Ohta and S. Nishizawa, "How much does color information help optical flow computation," *IEICE Trans. Inf. Syst.* **E89-D**, 1759–1762 (2006).
- J. Barron and R. Klette, "Quantitative color optical flow," in *Proceedings of IEEE International Conference on Pattern Recognition* (IEEE, 2002), Vol. 4, pp. 251–255.
- K. Aires, A. Santana, and A. Medeiros, "Optical flow using color information: preliminary results," in *Proceedings of the 2008 ACM Symposium on Applied Computing* (ACM, 2008), pp. 1607–1611.
- R. Andrews and B. Lovell, "Color optical flow," in *Proceedings of Workshop on Digital Image Computing* (2003), pp. 135–139.
- P. Golland and A. Bruckstein, "Motion from color," *Comput. Vision Image Underst.* **68**, 346–362 (1997).
- Y. Mileva, A. Bruhn, and J. Weickert, "Illumination-robust variational optical flow with photometric invariants," in *Proceedings of the 29th DAGM Conference on Pattern Recognition* (Springer-Verlag, 2007), pp. 152–162.
- T. Zickler, S. Mallick, D. Kriegman, and P. Belhumeur, "Color subspaces as photometric invariants," *Int. J. Comput. Vis.* **79**, 13–30 (2008).
- A. Johnston and C. Clifford, "A unified account of three apparent motion illusions," *Vision Res.* **35**, 1109–1123 (1995).
- A. Johnston, P. W. McOwan, and C. P. Benton, "Robust velocity computation from a biologically motivated model of motion perception," *Proc. R. Soc. London Series B* **266**, 509–518 (1999).
- B. A. Wandell, in *Foundations of Vision* (Sinauer, 1995), Chap. 9.
- E. Hering, *Outlines of a Theory of the Light Sense* (Harvard University, 1964).
- L. M. Hurvich and D. Jameson, "An opponent-process theory of color vision," *Psych. Rev.* **64**, 384–404 (1957).
- D. Jameson and L. M. Hurvich, "Some quantitative aspects of an opponent-colors theory. i. chromatic responses and spectral saturation," *J. Opt. Soc. Am.* **45**, 546–552 (1955).
- J. B. Derrico and G. Buchsbaum, "A computational model of spatiochromatic image coding in early vision," *J. Visual Commun. Image Represent.* **2**, 31–38 (1991).
- A. B. Poirson and B. A. Wandell, "The appearance of colored patterns: pattern-color separability," *J. Opt. Soc. Am. A* **10**, 2458–2470 (1993).
- M. D. Fairchild, in *Color Appearance Models*, 2nd ed. (Wiley, 2005), Chap. 8.
- J. Koenderink and A. Kappers, *Color Space* (Utrecht University, 1998).
- J.-M. Geusebroek, R. Boomgaard, A. W. Smeulders, and H. Geerts, "Color invariance," *IEEE Trans. Pattern Anal. Machine Intell.* **23**, 1338–1350 (2001).
- A. Johnston, P. W. McOwan, and H. Buxton, "A computational model of the analysis of first and second order motion by simple and complex cells," *Proc. R. Soc. London Series B* **250**, 297–306 (1992).
- P. W. McOwan, C. Benton, J. Dale, and A. Johnston, "A multi-differential neuromorphic approach to motion detection," *Int. J. Neural Syst.* **9**, 429–434 (1999).
- A. Johnston, P. W. McOwan, and C. Benton, "Biological computation of image motion from flows over boundaries," *J. Phys.* **97**, 325–334 (2003).

31. S. Baker and I. Matthews, "Lucas-Kandade 20 years on: a unifying framework," *Int. J. Comput. Vis.* **56**, 221–255 (2004).
32. A. Johnston, C. Benton, and P. W. McOwan, "Induced motion at texture-defined motion boundaries," *Proc. R. Soc. London Series B* **266**, 2441–2450 (1999).
33. T. Brox, A. Bruhn, N. Papenberg, and J. Weickert, "High accuracy optical flow estimation based on a theory for warping," in *Proceedings of European Conference on Computer Vision* (Springer, 2004), Vol. 4, pp. 25–36.
34. J. Koenderink and A. van Doorn, "Receptive field assembly pattern specificity," *J. Visual Commun. Image Represent.* **3**, 1–12 (1992).
35. J.-M. Geusebroek, R. Boomgaard, A. W. Smeulders, and T. Gevers, "Color constancy from physical principles," *Pattern Recognit. Lett.* **24**, 1653–1662 (2003).
36. S. Lang, *Calculus of Several Variables* (Springer-Verlag, 1987).
37. J. Schanda, in *Colorimetry: Understanding the CIE System* (Wiley, 2007), Chap. 3.
38. "HSV Color Space," http://en.wikipedia.org/wiki/HSL_and_HSV#cite_note12.
39. O. J. Braddick, "Segmentation versus integration in visual motion processing," *Trends Neurosci.* **16**, 263–268 (1993).
40. G. Botella, M. Rodriguez, A. Garcia, and E. Ros, "Neuromorphic configurable architecture for robust motion estimation," *Int. J. Reconfigurable Comput.* **2008**, 384–404 (2008).