Tensor Based Feature Detection for Color Images

J. van de Weijer and Th. Gevers
Intelligent Sensory Information Systems ,University of Amsterdam, The Netherlands

Abstract

Extending differential-based operations to color images is hindered by the multi-channel nature of color images. The derivatives in different channels can point in opposite directions, hence cancellation might occur by simple addition. The solution to this problem is given by the structure tensor for which opposing vectors reinforce each other.

We review the set of existing tensor based features which are applied on luminance images and show how to expand them to the color domain. We combine feature detectors with photometric invariance theory to construct invariant features. Experiments show that color features perform better than luminance based features and that the additional photometric information is useful to discriminate between different physical causes of features.

Introduction

Feature detection is an important tool in computer vision [1]. Differential based features for detecting events such as edges, corners, salient points, are used abundantly in a variety of applications such as matching, object recognition, tracking [2] [3] [4]. Although the majority of images is in color format nowadays, the computer vision community still uses luminance based feature extraction. Obviously, extensions of these feature detection techniques to the color domain is desired.

A pioneering work on extending edge detection to color images was proposed by DiZenzo [5]. The paper addresses the problem of opposing vectors for different color channels. Opposing vectors occur on edges where for one channel the signal decreases while for another the signal increases. A simple addition of the opposing derivative signals of the different channels reduces the total derivative strength. DiZenzo solves this problem by proposing the tensor based gradient for which opposing vectors reinforce one another. Sapiro and Ringach [6] further investigated the local structure tensor and the interpretation of its eigenvalues within the context of anisotropic diffusion.

Similar equations as found by DiZenzo [5] were presented by Kass and Witkin [7], who proposed an orientation estimator for orientated patterns (e.g. fingerprint images). Oriented patterns differ from non-oriented ones in that they consist of ridges which have a differential structure of opposing vectors within a small neighborhood. Just as for color images, the solution was found in tensor mathematics and a set of tensor based features for oriented patterns was proposed [8] [9].

In addition to the apparent loss of information which occurs by converting an RGB image to a luminance image, photometric information is also lost. In [10] Shafer introduces the dichromatic reflection model. The theory provides a physical model which allows for discrimination of different photometric events in the image, such as changes caused by shadows, shading or specularities. On the basis of this model, others provided algorithms invariant to various photometric events [11] [12]. Combining the photometric invariance theory with geometric operations has been investigated by Geusebroek et al. [13]. More recently, van de Weijer et al. [9] proposed the quasi-invariants, which is a set of derivatives invariant to photometric changes. These derivative filters have the advantage over existing methods that they remain stable over the entire RGB-space.

In this paper, we start from the observation made by DiZenzo that tensors are suited to combine first order derivatives for color images. The first contribution that we collect a number of existing features which are based on the structure tensor and show how to extend these to color images. The second contribution is that we combine these features with the photometric derivatives which allows for photometric invariant feature detection.

Extending Differential Based Operations to Color Images

The extension of differential based operations to color images can be done in various ways. The main challenge here is how to project differential structure back to a scalar representation. For the first order differential structure of color images this has been explored in [5] and [6]. Here, we describe several consideration which will result in the color tensor framework given in section 3.

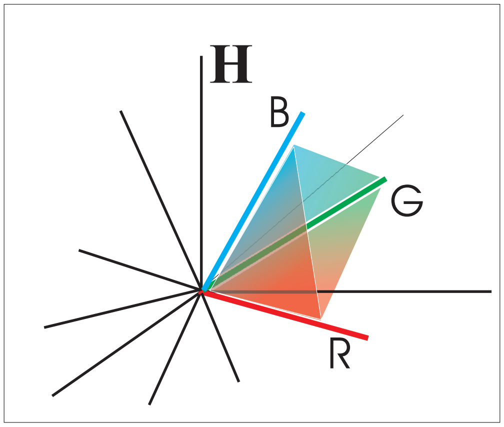

(a)

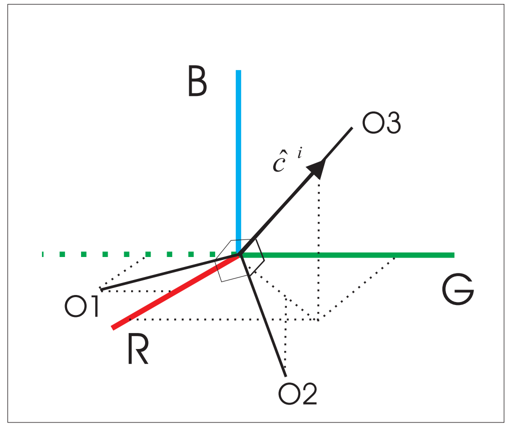

(b)

Figure 1: a) The subspace of measured light RGB in the Hilbert space of possible spectra. b) The RGB coordinate system and an alternative orthonormal color coordinate system which spans the same subspace.

Mathematical Viewpoint

As pointed out in [5], simply adding the differential structure of different channels may result in a cancellation even when evident structure exists in the image. E.g. for a red-green edge the derivatives in the red and green channel point in opposing directions. Instead of adding the direction information (defined on $[0, 2\pi\rangle$) of the different channels, it is more appropriate to sum the orientation information (defined on $[0, \pi\rangle$). In the example of the red-green edge the derivatives in the red and green channel have opposing direction, but the same orientation. A well known mathematical method for which vectors in opposite directions reinforce each other is provided by tensor mathematics. Tensors describe the local orientation rather than the direction. More precise, the tensor of a vector and the tensor of the same vector rotated over $180°$ are equal.

Photometric Viewpoint

A good reason for using color images is the photometric information which can be exploited. Photometric invariance theory provides invariants for different photometric variations. Well known results are photometric invariant colorspaces such as normalized RGB or HSI. Opposing derivative vectors are common for invariant colorspaces. Actually, for normalized RGB the summed derivative is per definition zero. Hence, the structure tensor is indispensable for computing the differential structure of photometric invariant representations of images.

Physical Viewpoint

Next to the photometric invariance discussed above, we will look into invariance with respect to coordinate transformations. For color images, values are represented in the RGB coordinate system. The ∞-dimensional Hilbert space is sampled with three probes which results in the red, green and blue channels (see fig. 1). For operations on the color coordinate system to be physically meaningful they should be independent of orthonormal transformation of the three axes in Hilbert space. As an example of an orthonormal color coordinate system the opponent color space can be mentioned (see fig. 1). The opponent color space spans the same subspace as the subspace defined by the RGB-axes and hence physically meaningful features computed from both subspace should yield the same results. We will verify the color features to be invariant of the accidental choice of the color coordinate frame.

Tensor Based Feature Detection for Color Images

In this section we extend several tensor based features to color images. As stated before, the tensor basis ensures that vectors pointing in opposite direction reinforce each other. Further, the feature detectors are verified to be invariant for orthonormal rotations of the RGB-space.

Structure Tensor Based Features

In [5] Di Zenzo pointed out that the correct method to combine the first order derivative structure is by using a local tensor. Analysis of the shape of the tensor leads to an orientation and a gradient norm estimate.

Given an image f, the structure tensor is given by

$$\mathbf{G} = \begin{pmatrix} \overline{f_x^2} & \overline{f_x f_y} \\ \overline{f_x f_y} & \overline{f_y^2} \end{pmatrix} \quad (1)$$

where the subscripts indicates spatial derivatives and the bar $\bar{.}$ indicates the convolution with a Gaussian filter. As discussed in section 2 tensors can be added for different channels. For a multichannel image $\mathbf{f} = \left(f^1, f^2, ..., f^n\right)^T$

the structure tensor is given by

$$\mathbf{G} = \begin{pmatrix} \overline{\mathbf{f}_x^T\mathbf{f}_x} & \overline{\mathbf{f}_x^T\mathbf{f}_y} \\ \overline{\mathbf{f}_y^T\mathbf{f}_x} & \overline{\mathbf{f}_y^T\mathbf{f}_y} \end{pmatrix} \tag{2}$$

where superscript T indicates the transpose operation. For color images $\mathbf{f} = (R, G, B)^T$, this results in the color structure tensor

$$\begin{pmatrix} \overline{R_x^2 + G_x^2 + B_x^2} & \overline{R_xR_y + G_xG_y + B_xB_y} \\ \overline{R_xR_y + G_xG_y + B_xB_y} & \overline{R_y^2 + G_y^2 + B_y^2} \end{pmatrix}. \tag{3}$$

The color structure tensor describes the 2D first order differential structure at a certain point in the image. Eigenvalue analysis of the tensor leads to two eigenvalues which are defined by

$$\begin{aligned} \lambda_1 &= \tfrac{1}{2}\left(\overline{\mathbf{f}_x^T\mathbf{f}_x} + \overline{\mathbf{f}_y^T\mathbf{f}_y} + \sqrt{\left(\overline{\mathbf{f}_x^T\mathbf{f}_x} - \overline{\mathbf{f}_y^T\mathbf{f}_y}\right)^2 + \left(2\overline{\mathbf{f}_x^T\mathbf{f}_y}\right)^2}\right) \\ \lambda_2 &= \tfrac{1}{2}\left(\overline{\mathbf{f}_x^T\mathbf{f}_x} + \overline{\mathbf{f}_y^T\mathbf{f}_y} - \sqrt{\left(\overline{\mathbf{f}_x^T\mathbf{f}_x} - \overline{\mathbf{f}_y^T\mathbf{f}_y}\right)^2 + \left(2\overline{\mathbf{f}_x^T\mathbf{f}_y}\right)^2}\right) \end{aligned} \tag{4}$$

The direction of λ_1 indicates the prominent local orientation, which is equal to the orientation in the image with maximum color change,

$$\theta = \tfrac{1}{2}\arctan\left(\frac{2\overline{\mathbf{f}_x^T\mathbf{f}_y}}{\overline{\mathbf{f}_x^T\mathbf{f}_x} - \overline{\mathbf{f}_y^T\mathbf{f}_y}}\right). \tag{5}$$

The λ's can be combined to give the following local descriptors:

- $\lambda_1 + \lambda_2$ describes the total local derivative energy.
- λ_1 is the derivative energy in the most prominent direction.
- $\lambda_1 - \lambda_2$ describes the 'line'-energy (see [6]). The derivative energy in the prominent orientation is corrected for by energy contributed by noise, λ_2.
- λ_2 describes the amount of derivative energy perpendicular to the prominent local orientation which is important for example to select good features for tracking [4].

An often applied feature detector, which is based on the structure tensor in computer vision, is the Harris corner detector. The color Harris operator H on an image $\mathbf{f}$ can be computed with

$$H\mathbf{f} = \overline{\mathbf{f}_x^T\mathbf{f}_x}\;\overline{\mathbf{f}_y^T\mathbf{f}_y} - \overline{\mathbf{f}_x^T\mathbf{f}_y}^2 - k\left(\overline{\mathbf{f}_x^T\mathbf{f}_x} + \overline{\mathbf{f}_y^T\mathbf{f}_y}\right). \tag{6}$$

The elements of the tensor are proven to be invariant under spatial transformations, the invariance with respect to orthonormal transformation of the RGB-axis is easily checked, since $\frac{\partial}{\partial x}\mathbf{R}\mathbf{f} = \mathbf{R}\mathbf{f}_x$ and

$$\overline{\left(\mathbf{R}\mathbf{f}_x\right)^T\mathbf{R}\mathbf{f}_y} = \overline{\mathbf{f}_x^T\mathbf{R}^T\mathbf{R}\mathbf{f}_y} = \overline{\mathbf{f}_x^T\mathbf{f}_y} \tag{7}$$

Here $\mathbf{R}$ is a rotation operator on the channels of $\mathbf{f}$.

Adapted Structure Tensor Based Features

Similar equations to Di Zenzo's equations for orientation estimation are found by Kass and Witkin [7]. They studied orientation estimation for oriented patterns (e.g. fingerprint images). Oriented patterns are defined as patterns with a dominant orientation everywhere. For oriented patterns other mathematics are needed than for regular object images. The local structure of object images is described by a step edge, whereas for oriented patterns the local structure is described as a set of lines (roof edges). Lines have the property that they have opposing vectors on a small scale. Hence for geometric operations on oriented patterns, mathematical methods are needed for which opposing vectors enforce one another. This is the same problem as encountered for all color images and hence similar equations were found in both fields. Next to orientation estimation, a number of other estimators were proposed by oriented pattern research [14] [8] [15]. These operation are based on adaptations of the structure tensor and can also be applied to the color tensor.

The structure tensor of eq. 1 can also be seen as local projection of the derivative energy on perpendicular axes, namely $\mathbf{u}_1 = \begin{pmatrix} 1 & 0 \end{pmatrix}^T$ and $\mathbf{u}_2 = \begin{pmatrix} 0 & 1 \end{pmatrix}^T$.

$$G = \begin{pmatrix} \overline{\left(\mathbf{f}_{x,y}\mathbf{u}_1\right)^T\left(\mathbf{f}_{x,y}\mathbf{u}_1\right)} & \overline{\left(\mathbf{f}_{x,y}\mathbf{u}_1\right)^T\left(\mathbf{f}_{x,y}\mathbf{u}_2\right)} \\ \overline{\left(\mathbf{f}_{x,y}\mathbf{u}_1\right)^T\left(\mathbf{f}_{x,y}\mathbf{u}_2\right)} & \overline{\left(\mathbf{f}_{x,y}\mathbf{u}_2\right)^T\left(\mathbf{f}_{x,y}\mathbf{u}_2\right)} \end{pmatrix} \tag{8}$$

in which $\mathbf{f}_{x,y} = \begin{pmatrix} \mathbf{f}_x & \mathbf{f}_y \end{pmatrix}$ is a $n \times 2$ matrix. In [8] [14] local symmetry features based on Lie groups of transformations led to whole sets of perpendicular $\mathbf{u}_1$ and $\mathbf{u}_2$'s. They include feature extraction for circle, spiral and starlike structures. Here the star and circle detector is given as an example. It is based on $\mathbf{u}_1 = \frac{1}{\sqrt{x^2+y^2}}\begin{pmatrix} x & y \end{pmatrix}^T$ which coincide with the derivative pattern of a circular patterns and $\mathbf{u}_2 = \frac{1}{\sqrt{x^2+y^2}}\begin{pmatrix} -y & x \end{pmatrix}^T$ which denotes the perpendicular vector field which coincides with the derivative pattern of starlike patterns. These vectors can be used to compute the adapted structure tensor with eq. 8. Only the elements on the diagonal have non zero entries and are equal to

$$\begin{aligned} \lambda_1 = H_{11} &= \overline{x^2\mathbf{f}_x^T\mathbf{f}_x} + 2\overline{xy\mathbf{f}_x^T\mathbf{f}_y} + \overline{y^2\mathbf{f}_y^T\mathbf{f}_y} \\ \lambda_2 = H_{22} &= \overline{x^2\mathbf{f}_y^T\mathbf{f}_y} - 2\overline{xy\mathbf{f}_x^T\mathbf{f}_y} + \overline{y^2\mathbf{f}_x^T\mathbf{f}_x} \end{aligned}. \tag{9}$$

Here λ_1 describes the amount of derivative energy contributing to circular structures and λ_2 the derivative energy which describes a starlike structure. Similar to the prove given in eq. 7 the elements of eq. 9 can be proven to be invariant under transformations of the RGB-space.

Photometric Derivatives

One of the motivations for color image based feature extraction is the extra photometric information which is lost by converting to a luminance image. Here we describe a set of filters, the quasi-invariants [9] which are respectively invariant for shadow/shading and specular edges. The combination of the photometric invariance theory with the tensor based features of section 3 allows for photometric invariant feature detection.

The Dichromatic Reflection Model

The dichromatic model divides the reflection in the interface (specular) and body (diffuse) reflection component for optically inhomogeneous materials. We assume white illumination, i.e. smooth spectrum of nearly equal energy at all wavelengths, and neutral interface reflection. Further, we assume that shadows are not significantly colored. The RGB vector, $\mathbf{f} = (R, G, B)^T$, can be seen as a weighted summation of two vectors,

$$\mathbf{f} = e(m^b \mathbf{c}^b + m^i \mathbf{c}^i) \tag{10}$$

in which $\mathbf{c}^b$ is the color of the body reflectance, $\mathbf{c}^i$ the color of the interface reflectance (i.e. specularities or highlights), m^b and m^i are scalars representing the corresponding magnitudes of reflection and e is the intensity of the light source. For matte surfaces there is no interface reflection and the model further simplifies to

$$\mathbf{f} = e m^b \mathbf{c}^b \tag{11}$$

which is the well-known Lambertian reflection.

Photometric Quasi-Invariance

Here we describe the set of derivatives called the quasi-invariants and the photometric variants introduced in [9]. They are called so to distinguish them from the full invariants such as normalized RGB and hue. The quasi-invariants have the advantage that they lack the instabilities which are inherent to the full invariants, e.g. around the normalized RGB is unstable near zero intensity. Therefore the quasi-invariants are a better starting-point for the derivative based features proposed in section 3.

The quasi-invariants are computed by projecting the derivative on photometric axes in the RGB-space. These three axes are

1. the shadow-shading axis in which changes due to shadow and shading occur. This direction follows from the fact that changes of the shadow-shading (i.e. changes in em^b) for matte surface are parallel to the direction of $\mathbf{f}$ according to eq. 11.
2. the specular axis in which changes due to specularities occur. As can be seen from eq. 10 this direction is determined by $\mathbf{c}^i$ which is equal to the color of the lightsource. For a white illuminant $\mathbf{c}^i = \frac{1}{\sqrt{3}}(1,1,1)^T$.
3. the hue direction in which no changes due to specularities, shadow and shading occur. The hue direction, $\mathbf{b}$ can be found by taking the outer product between the shadow-shading and the specular direction.
$$\hat{\mathbf{b}} = \frac{\hat{\mathbf{f}} \times \hat{\mathbf{c}}^i}{\left|\hat{\mathbf{f}} \times \hat{\mathbf{c}}^i\right|}.$$
The hat, $\hat{.}$, is used to denote unit vectors.

The shadow-shading variant, $\mathbf{S}_x$ contains the derivative part which can be explained by shadow or a shading change. Subtracting this derivative from the total derivative results in the shadow-shading quasi-invariant $\mathbf{S}_x^c$,

$$\begin{array}{l} \mathbf{S}_x = \left(\mathbf{f}_x^T \hat{\mathbf{f}}\right) \hat{\mathbf{f}} \\ \mathbf{S}_x^c = \mathbf{f}_x - \mathbf{S}_x \end{array} . \tag{12}$$

The specular variant, $\mathbf{O}_x$ is found by a projection on the lightsource color direction $\hat{\mathbf{c}}^i$, and comprises of the derivative signal which can be caused by specularity. The specular quasi-invariant is found by subtracting the specular variant from the total derivative signal, and contains all the derivative signal which cannot be explained by specular reflections.

$$\begin{array}{l} \mathbf{O}_x = \left(\mathbf{f}_x^T \hat{\mathbf{c}}^i\right) \hat{\mathbf{c}}^i \\ \mathbf{O}_x^c = \mathbf{f}_x - \mathbf{O}_x \end{array} . \tag{13}$$

Finally, the specular shadow-shading quasi-invariant, $\mathbf{H}_x^c$ can be computed by projection the object direction and contains the derivative signal which cannot be explained by shadow-shading changes nor by specularities.

$$\begin{array}{l} \mathbf{H}_x^c = \left(\mathbf{f}_x^T \cdot \hat{\mathbf{b}}\right) \hat{\mathbf{b}} \\ \mathbf{H}_x = \mathbf{f}_x - \mathbf{H}_x^c \end{array} . \tag{14}$$

By projecting the local spatial derivative on three photometric axis in the RGB cube we have derived photometric quasi-invariants. These can be combined with the structure tensor eq. 4 for photometric quasi-invariant feature detection. As discussed in section we would like features to be independent of orthonormal changes of color coordinates. For this to be true a rotation of the color coordinates should result in a rotation of the quasi-invariant

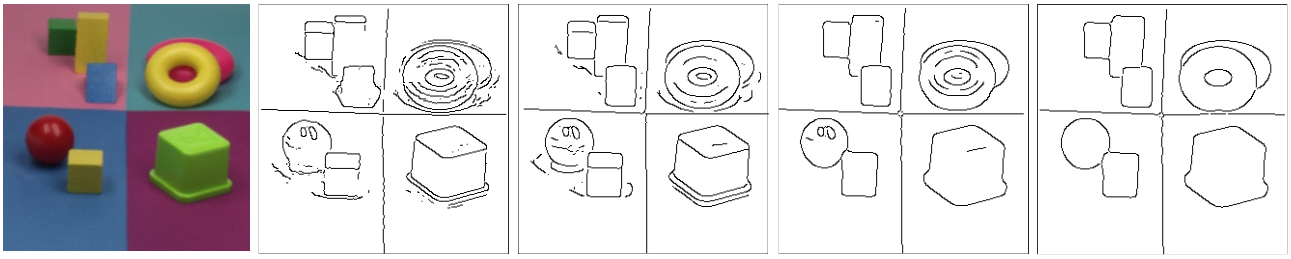

(a) (b) (c) (d) (e)

Figure 2: a) input image with Canny edge detection results on successively b) the luminance image c) the RGB image d) the shadow-shading quasi-invariant e) the specular shadow-shading quasi-invariant.

derivatives. For the shadow-shading quasi-variant $\mathbf{S}_x$ this can be proven by

$$\begin{aligned}\left((\mathbf{R}\mathbf{f}_x)^T\mathbf{R}\hat{\mathbf{f}}\right)\left(\mathbf{R}\hat{\mathbf{f}}\right) &= \left(\mathbf{f}_x^T\mathbf{R}^T\mathbf{R}\hat{\mathbf{f}}\right)\left(\mathbf{R}\hat{\mathbf{f}}\right)\\ &= \mathbf{R}\left(\mathbf{f}_x^T\hat{\mathbf{f}}\right)\hat{\mathbf{f}} = \mathbf{R}\mathbf{S}_x\end{aligned} \quad (15)$$

For the other photometric variants and quasi-invariants similar proofs can be given.

Experiments

In this section experiments of the tensor based features of section 3 in combination with photometric theory of section 4 are given. During these experiments we assume white illumination, and hence $\hat{\mathbf{c}}^i = \frac{1}{\sqrt{3}}(1,1,1)^T$.

Color Canny Edge Detection

In fig.2 the results of color Canny edge detection for several photometric quasi-invariants is shown. The algorithm finds the edges in images, based on the differentail structure within a local window. The algorithm consists of the following steps

1. Compute the spatial derivatives, $\mathbf{f}_x$ and combine them if desired into a quasi-invariant (eq. 12, 13, 14).
2. Compute the maximum eigenvalue (eq.4) and its orientation (eq.5). In this step the vectorial derivatives are projected on a scalar eigenvalue.
3. Suppress the pixels which are not a local maximum of λ_1 in the promiment orientation(, also known as non-maximum suppression).

The results show the advantage of applying Canny on the color image given in fig2a, above the luminance image fig2b. Also the removal of spurious edges by photometric invariance is demonstrated. In fig2d the edge detection is robust to shadow and shading changes and only detects material and specular edges. In fig2e only the material edges are depicted.

Color Features for Tracking

In [4] Shi and Tomasi propose a features for tracking. The method selects points for which $\lambda_2 > threshold$ (see eq.4), or points for which the derivative energy perpendicular to the prominent direction is above a threshold. This typically selects corners in images. In fig.3 results based on the photometric derivatives are depicted. Again the luminance based method misses some important features. The results in fig.3c,d show that scene incidental points can be removed by applying photometric invariant features.

Local Color Symmetry

Here we apply the circle detector disccussed in section 3. The experiment is performed on an image with lego-blocks. Because we know that the color within the blocks remains the same, the circle detection is done on the shadow-shading specular variant, $\mathbf{H}_x$ (eq. 14). The shadow-shading specular variant contains all the derivative energy except for the energy which can only be caused by a material edge. With the shadow-shading variant the circular energy λ_1 and the starlike energy λ_2 are computed (Eq. 9). Dividing the circular energy by the total energy yields a good descriptor of local circularity (see fig.9)

$$C = \frac{\lambda_1}{\lambda_1 + \lambda_2}. \quad (16)$$

In fig.4c the maxima of C are superimposed on the input image. Most of the circles are detected.

Conclusions

In this we paper we proposed a set of tensor based color image features. The tensor basis of these features ensures that opposing vectors in different channels do not cancel out, but instead reinforce each other. Experiments show that the features outperform their luminance counterparts, and that the combination with photometric invariance theory allows for the removal of undesired features.

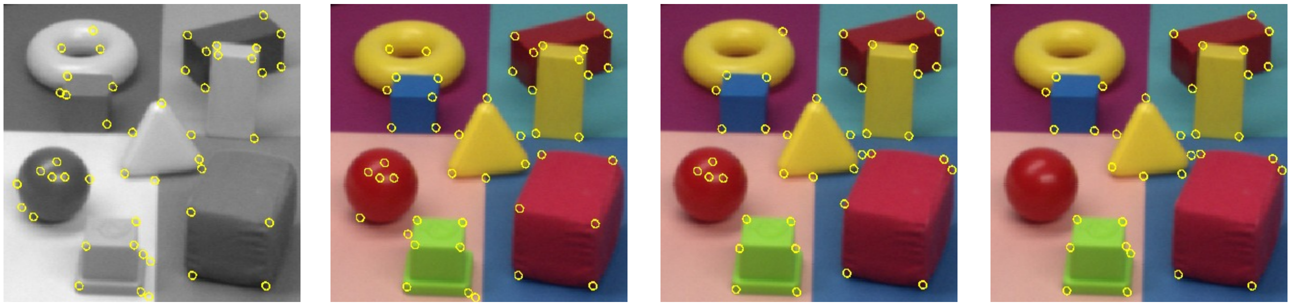

(a) (b) (c) (d)

Figure 3: *Features selection based on λ_2 successively applied on a) the grey image b) the RGB image c) the shadow-shading quasi-invariant d) the specular shadow-shading quasi-invariant.*

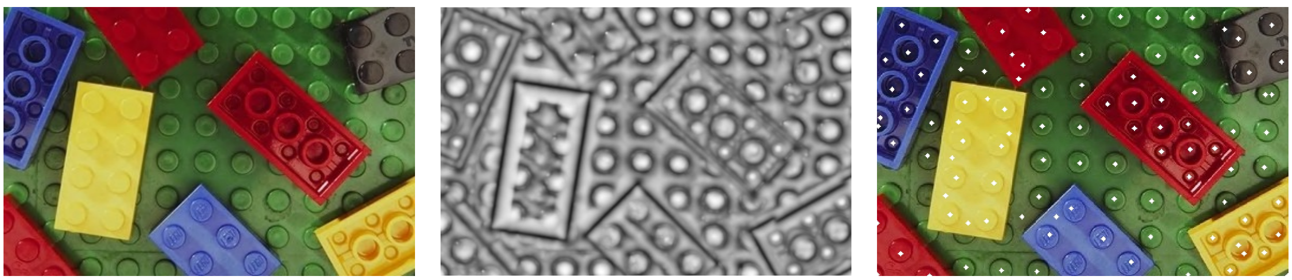

(a) (b) (c)

Figure 4: *a) input image (from Corel-database[16]) b) the circularity coefficient C c) the detected circles.*

References

1. R.M. Haralick and L.G. Shapiro, Computer and Robot Vision, volume II, Addison-Wesley (1992).
2. C. Schmid, R. Mohr and C. Bauckhage, Evaluation of interest point detectors, International Journal of Computer Vision, 37(2), 151 (2000).
3. N. Sebe, Q. Tian, E. Loupias, M.S. Lew and T.S. Huang, Evaluation of salient point techniques, Image and Vision Computing, 21(13-14), 1087 (2003).
4. J. Shi and C. Tomasi, Good features to track, in IEEE conference on Computer Vision and Pattern Recognition (1994).
5. Silvano Di Zenzo, Note: A note on the gradient of a multi-image, Computer Vision, Graphics, and Image Processing, 33(1), 116 (1986).
6. G. Sapiro and D. Ringach, Anisotropic diffusion of multivalued images with applications to color filtering, IEEE Trans. Image Processing, 5(11), 1582 (1996).
7. M. Kass and A. Witkin, Analyzing oriented patterns, Computer Vision, Graphics, and Image Processing, 37, 362 (1987).
8. J. Bigun, Pattern recognition in images by symmetry and coordinate transformations, Computer Vision and Image Understanding, 68(3), 290 (1997).
9. J. van de Weijer, Th. Gevers and J.M. Geusebroek, Color edge detection by photometric quasi-invariants, in Int'l Conf. Computer Vision, Nice, France, pp. 1520–1526 (2003).
10. S.A. Shafer, Using color to seperate reflection components, COLOR research and application, 10(4), 210 (1985).
11. Th. Gevers, Color image invariant segmentation and retrieval, Ph.D. thesis, University of Amsterdam (1997).
12. G.J. Klinker and S.A. Shafer, A physical approach to color image understanding, Int. Journal of Computer Vision, 4, 7 (1990).
13. J.M. Geusebroek, R. van den Boomgaard, A.W.M. Smeulders and H. Geerts, Color invariance, IEEE Trans. Pattern Analysis Machine Intell., 23(12), 1338 (2001).
14. O. Hansen and J. Bigun, Local symmetry modeling in multi-dimensional images, pattern Recognition Letters, 13, 253 (1992).
15. J. van de Weijer, L.J. van Vliet, P.W. Verbeek and M. van Ginkel, Curvature estimation in oriented patterns using curvilinear models applied to gradient vector fields, IEEE Trans. Pattern Analysis and Machine Intelligence, 23(9), 1035 (2001).
16. Corel Gallery, www.corel.com.

Biography

Joost van de Weijer received the M.Sc. degree in applied physics at the Delft University of Technology, The Netherlands, in 1998. Since 1999, he is a Ph.D. student in the ISIS group at the University of Amsterdam. His current research interests include filter theory, color image filtering, saliency detection and photometric invariance.

Theo Gevers is an Assoiciate Professor of Computer Science at the University of Amsterdam, The Netherlands. His main research interests are in the fundamentals of image database system design, image retrieval by content, theoretical foundations of geometric and photometric invariants and color image processing.