Classifying Color Edges in Video Into Shadow-Geometry, Highlight, or Material Transitions

Theo Gevers, *Member, IEEE,* and Harro Stokman

***Abstract*—We aim at using color information to classify the physical nature of edges in video. To achieve physics-based edge classification, we first propose a novel approach to color edge detection by automatic noise-adaptive thresholding derived from sensor noise analysis. Then, we present a taxonomy on color edge types. As a result, a parameter-free edge classifier is obtained labeling color transitions into one of the following types: 1) shadow-geometry, 2) highlight edges, and 3) material edges. The proposed method is empirically verified on images showing complex real world scenes.**

***Index Terms*—Adaptive edge thresholding, color edge detection, color invariance, edge classification, noise propagation, video.**

I. Introduction

THE COLOR of objects vary with changes in illuminant color and viewing conditions. As a consequence, color boundaries are generated by a large variety of imaging variables such as shadows, highlights, illumination and material changes. Therefore, classifying the physical nature of edges is useful for a large number of applications such as video segmentation, video indexing and content recognition, where corresponding edge types (e.g. material transitions) from distinct image frames are selected for image matching while discounting other "accidental" edge types (e.g. shadows and highlight transitions). In this paper, we consider the problem of discriminating different edge types based on local surface reflectance properties.

Edge detection in intensity images is well established in [1] and [2], for example. In general, to achieve color edge detection, these intensity-based edge detection techniques are extended by taking the sum or Euclidean distance from the individual gradient maps. Further, to obtain robustness against illumination effects, Tsang and Tsang [3] show that edge detection in the hue color space is effective in suppressing specular reflection. However, no edge classification scheme is provided. Another approach is given by Zhang and Bergholm [4] to classify edges into diffuse and sharp edges. The idea is that illumination phenomena, such as indoor shadows and reflections on glossy surfaces tend to cause gradual transitions, whereas edges between distinct objects tend to be sharp. A similar but more elaborated approach is given by Stander [5] for detecting moving shadows. However, the method is based on a complex geometry model restricted to the detection of cast shadows. Furthermore, all of the above mentioned classification techniques are based on color edge detection, which is, in general, dependent on the appropriate setting of threshold values to determine the edge maps. This threshold is found, in general, by trial-and-error. For general video segmentation and content recognition, manual settings of thresholds should be avoided. Therefore, an automatic way for threshold value selection is required.

In this paper, we aim at automatically classifying the physical nature of edges in images using color and reflectance information. To achieve this, we first propose a novel framework to compute edges by automatic gradient thresholding. Then, we present a taxonomy on edge types based upon the sensitivity of edges with respect to different imaging variables. Finally, a parameter-free edge classifier is provided labeling color transitions into one of the following types: 1) shadow-geometry edges, 2) highlight edges, and 3) material edges. The proposed method is empirically verified on video sequences recorded from complex real world scenes.

The paper is organized as follows. In Section II, the basics on reflection are discussed first. Further, different color models are presented and a taxonomy on color invariance is given. In Section III, computational methods are proposed to get to color invariant gradients. Next, in Section IV, error estimation and propagation is discussed. Finally, in Section V, the color edge classification scheme is proposed.

II. Photometric Invariance

In Gevers and Smeulders [6], different color models are proposed which show some degree of invariance for the purpose of object recognition. In this paper, we use the different color models for the purpose of color edge classification in video sequences. Therefore, in this section, we first reconsider the basic reflection definitions in Section II-A. Then, we discuss the different color models in Section II-B and their sensitivity with respect to the imaging conditions. We conclude this section with a taxonomy on photometric invariance.

A. Basic Reflection Definitions

The reflection from inhomogeneous dielectric materials under white or spectrally smooth illumination is given by [7]

$$\omega_k = G_B(\vec{n},\vec{s})E\int_\lambda B(\lambda)F_k(\lambda)d\lambda + G_S(\vec{n},\vec{s},\vec{v})ESF \quad (1)$$

Manuscript received July 25, 2000; revised May 23, 2002. The associate editor coordingating the review of this paper and approving it for publication was Dr. Chuang Gu.

The authors are with the Informatics Institute, University of Amsterdam, 1098 SJ, Amsterdam, The Netherlands (gevers@science.uva.nl).

Digital Object Identifier 10.1109/TMM.2003.811620

for $\omega_k \in \{R, G, B\}$ giving the red, green and blue sensor response of an infinitesimal matte surface patch under the assumption of a white or spectrally smooth light source. Spectral sensitivities are given by $F_R(\lambda)$, $F_G(\lambda)$ and $F_B(\lambda)$ respectively, where λ denotes the wavelength. $B(\lambda)$ is the surface albedo. Further, E denotes the white light source and S is the Fresnel reflectance. These are constant over the wavelengths assuming white or spectrally smooth illumination (i.e., approximately equal/smooth energy density for all wavelengths within the visible spectrum) and the neutral interface reflection (NIR) model (i.e. $S(\lambda)$ has a constant value independent of the wavelength). Consequently, we have $E(\lambda) = E$ and $S(\lambda) = S$. Further, $\vec{n}$ is the surface patch normal, $\vec{s}$ is the direction of the illumination source, and $\vec{v}$ is the direction of the viewer. Finally, geometric terms G_B and G_S denote the geometric dependencies on the body and surface reflection component.

B. Color Models

We focus on normalized color c_1c_2 defined by [6]:

$$c_1(R,G,B) = \arctan\left(\frac{R}{B}\right) \tag{2}$$

$$c_2(R,G,B) = \arctan\left(\frac{G}{B}\right). \tag{3}$$

Further, we consider the two-dimensional opponent color space (leaving out the intensity component here) defined by

$$o_1(R,G,B) = \frac{(R-G)}{2} \tag{4}$$

$$o_2(R,G,B) = \frac{B}{2} - \frac{(R+G)}{4}. \tag{5}$$

The opponent color space is well-known and has its fundamentals in human perception.

Consider the body reflection term of (1):

$$\beta_k(\vec{x}) = G_B(\vec{x},\vec{n},\vec{s})E(\vec{x})\int_\lambda B(\vec{x},\lambda)F_k(\lambda)d\lambda \tag{6}$$

giving the kth sensor response of an infinitesimal matte surface patch under the assumption of a white light source. Again, for a color camera we have $k = \{R, G, B\}$.

The body reflection component describes the way light interacts with a dull surface. The light spectrum E falls on a surface B. The geometric and photometric properties of the body reflection depends on many factors. If we assume a random distribution of the pigments, the light exits in random directions from the body. In this simple case, the distribution of exiting light can be described by Lambert's law. Lambertian reflection models dull, matte surfaces which appear equally bright regardless from angle they are viewed. They reflect light with equal intensity in all directions. As a consequence, a uniformly colored surface which is curved (i.e. varying surface orientation) gives rise to a broad variance of RGB values. The same argument holds for intensity I.

In contrast, the c_1c_2 color model is a photometric invariant for dull objects, cf. (6) and (2) and (3):

$$c_1(\beta_R,\beta_G,\beta_B) = \arctan\left(\frac{G_B(\vec{n},\vec{s})EK_R(\lambda)}{G_B(\vec{n},\vec{s})EK_B(\lambda)}\right) = \arctan\left(\frac{K_R(\lambda)}{K_B(\lambda)}\right) \tag{7}$$

$$c_2(\beta_R,\beta_G,\beta_B) = \arctan\left(\frac{G_B(\vec{n},\vec{s})EK_G(\lambda)}{G_B(\vec{n},\vec{s})EK_B(\lambda)}\right) = \arctan\left(\frac{K_G(\lambda)}{K_B(\lambda)}\right) \tag{8}$$

where

$$K_C(\lambda) = \int_\lambda B(\lambda)F_C(\lambda)d\lambda \text{ for } C \in \{R,G,B\} \tag{9}$$

is the compact formulation depending on the sensors and surface albedo only. Note that the dependency on illumination, object pose, camera position, and object shape is factored out, i.e., c_1c_2 is only dependent on the sensors and the surface albedo.

For shiny surfaces, o_1o_2 is independent of highlights as follows from substituting (1) in (4) and (5) (see (10), shown at the bottom of the page). Equal argument also holds for o_2. Note that o_1o_2 is still dependent on $G_B(\vec{n},\vec{s})$ and E, and consequently being sensitive to object geometry and shading.

In conclusion, c_1c_2 varies with a change in material and highlights, o_1o_2 with a change in material and object-geometry, and RGB varies with a change in material, highlights and object-geometry.

III. Photometric Invariant Gradients

A number of well established techniques for edge detection in ordinary (one-band) images is available, [1], [2]. However, for color edge detection, the computation of the gradient magnitude is more complex due to the multivalued nature of a color image. In this paper, we compute the distance in color space by the Euclidean metric over the various channels.

Further, the color channels of an image are differentiated in the x and y direction using the Prewitt filter giving the gradient as $((\partial c_i)/(\partial x), (\partial c_i)/(\partial y))$. Here, c_i is the notation for a particular color channel. The Prewitt operator is chosen merely for its simplicity.

Then, the modulus of the gradient ∇F of the color planes is obtained by taking the Euclidean distance

$$\nabla F = \sqrt{\sum_{i=1}^{N}\left[\left(\frac{\partial c_i}{\partial x}\right)^2 + \left(\frac{\partial c_i}{\partial y}\right)^2\right]} \tag{11}$$

where N is the dimensionality of the color space.

$$\begin{aligned} o_1(\omega_R,\omega_G,\omega_B) &= \frac{\left(\left(G_B(\vec{n},\vec{s})E\int_\lambda B(\lambda)F_R(\lambda)d\lambda + G_S(\vec{n},\vec{s},\vec{v})ESF\right) - \left(G_B(\vec{n},\vec{s})E\int_\lambda B(\lambda)F_G(\lambda)d\lambda + G_S(\vec{n},\vec{s},\vec{v})ESF\right)\right)}{2} \\ &= \frac{\left(G_B(\vec{n},\vec{s})E\int_\lambda B(\lambda)F_R(\lambda)d\lambda - G_B(\vec{n},\vec{s})E\int_\lambda B(\lambda)F_G(\lambda)d\lambda\right)}{2} \end{aligned} \tag{10}$$

Often false edges are introduced due to sensor noise. These false edges are usually eliminated by using a threshold value determining the minimum acceptable gradient modulus. In this paper, we aim at providing a computational framework to determine automatically this local threshold value. To achieve this, sensor noise characteristics on color transformation are studied in the next section.

	shape edges	shadow edges	highlight edges	material edges
∇C_{RGB}	+	+	+	+
∇C_{c1c2}	−	−	+	+
∇C_{o1o2}	+	+	−	+

Fig. 1. Taxonomy of color edges based upon the sensitivity of the different color edge models with respect to the imaging conditions. − denotes invariant and + denotes sensitivity of the color edge model to the imaging condition.

IV. Error Propagation

In this paper, we assume that the noise is normally distributed, because the most frequently occurring noise is additive Gaussian noise. It is widely used to model thermal noise and is the limiting behavior of photon counting noise and film grain noise.

Then, let the result of a number of measurements of a random quantity u be given by

$$\hat{u} = u_{\text{best}} \pm \sigma_u \tag{12}$$

where u_{best} is the average value which is the best estimate for the quantity u and σ_u the standard deviation denoting the uncertainty or error in the measurement of u. Suppose that $u, \cdots, w$ are measured with corresponding uncertainties $\sigma_u, \cdots, \sigma_w$, and that the measured values are used to compute the function $q(u, \cdots, w)$. If the uncertainties in $u, \cdots, w$ are independent, random and relatively small, then the standard deviation or the so-called predicted uncertainty in $\hat{q}$ is given by [8]

$$\sigma_q = \sqrt{\left(\frac{\partial q}{\partial u}\sigma_u\right)^2 + \cdots + \left(\frac{\partial q}{\partial w}\sigma_w\right)^2} \tag{13}$$

where $\partial q/\partial u$ and $\partial q/\partial w$ are the partial derivatives of q with respect to u and w. In any case, the uncertainty in q is never larger than the ordinary sum

$$\sigma_q \leq \left|\frac{\partial q}{\partial u}\right|\sigma_u + \cdots \left|\frac{\partial q}{\partial w}\right|\sigma_w \tag{14}$$

if and only if the uncertainties $\sigma_u, \cdots, \sigma_w$ are relatively small. Assuming normally distributed random quantities, a way to calculate the standard deviations σ_R, σ_G, and σ_B is to compute the mean and variance estimates derived from a homogeneously colored surface patch in an image under controlled imaging conditions. Although (13) and (14) are deduced for random errors, they have been used as universal formulas for all kinds of errors.

After calculating the noise variance, the uncertainty of c_1c_2 can be found by substitution of (2) and (3) into (13) as

$$\sigma_{c1} = \sqrt{\frac{R^2\sigma_B^2 + B\sigma_R^2}{(R^2+B^2)^2}} \tag{15}$$

$$\sigma_{c2} = \sqrt{\frac{G^2\sigma_B^2 + B\sigma_G^2}{(G^2+B^2)^2}} \tag{16}$$

where σ_R^2, σ_G^2 and σ_B^2 denote the sensor noise variance, and σ_{c_1} and σ_{c_2} represent the uncertainty (standard deviation) in the normalized red and green color components, respectively. From the analytical study of (15) and (16), it can be derived that normalized color becomes unstable around the black point $R = G = B = 0$.

Further, the uncertainty of the o_1o_2 opponent coordinates is given as

$$\sigma_{o_1} = \frac{1}{2}\sqrt{\sigma_G^2 + \sigma_R^2} \tag{17}$$

$$\sigma_{o_2} = \frac{1}{2}\sqrt{4\sigma_B^2 + \sigma_G^2 + \sigma_R^2} \tag{18}$$

which is relatively stable at all RGB points.

Further, to propagate the uncertainties from these color components through the gradient modulus, the uncertainties are determined using (14) because the transformed color components are dependent. Using (14), the propagation of uncertainty of the Prewitt filter can be implemented by filtering the uncertainty planes of the different color spaces with the absolute masks yielding the uncertainties in the gradient $\sigma_{\partial c/\partial x}$ and $\sigma_{\partial c/\partial y}$. Then, the uncertainty in the gradient modulus of (11) is determined using (14) as

$$\sigma_{\nabla F} \leq \frac{\sum_i \left[\left(\frac{\partial c_i}{\partial x}\right)\cdot\sigma_{\frac{\partial c_i}{\partial x}} + \left(\frac{\partial c_i}{\partial y}\right)\cdot\sigma_{\frac{\partial c_i}{\partial y}}\right]}{\sqrt{\sum_i \left[\left(\frac{\partial c_i}{\partial x}\right) + \left(\frac{\partial c_i}{\partial y}\right)\right]}} \tag{19}$$

where i is the dimensionality of the color space. In this way, the effect of measurement uncertainty due to thermal and photon noise is propagated through the color invariant gradient.

V. Automatic Edge Thresholding and Classification

In this section, techniques are presented to automatically select the gradient threshold value. The threshold value is computed locally in an adaptive way. In fact, the amount of uncertainty at image locations will steer the threshold value. Finally, a color edge taxonomy is presented on which the novel rule-based edge classifier is based on.

A. Parameter-Free Edge Thresholding

From (19), the uncertainty associated with the gradient modulus is known. Color edges are thresholded taking this uncertainty into account. Assuming that the noise is normally distributed, then distribution is well-approximated by the Gauss distribution [8]. For a Gaussian distribution 99% of the values fall within a 3σ margin. If a gradient modulus is detected which

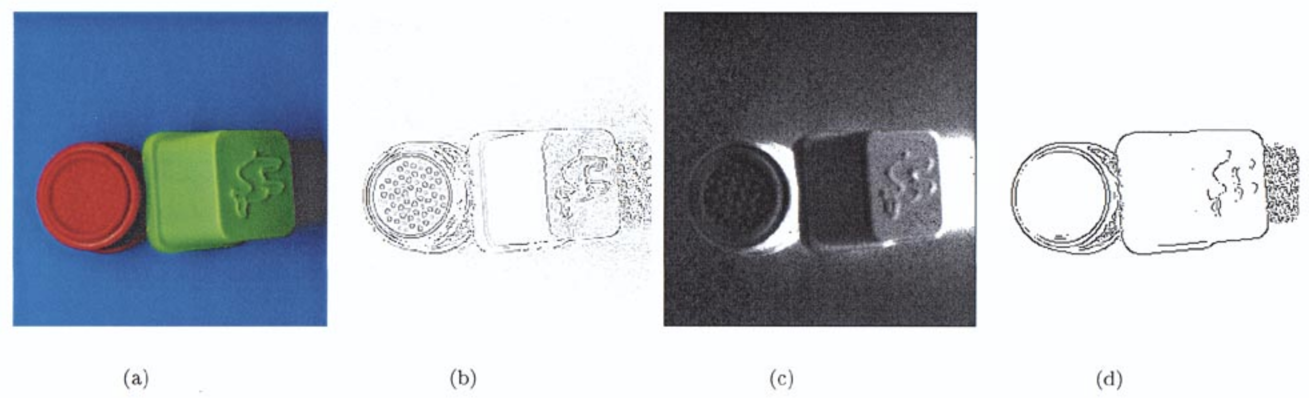

Fig. 2. Edges in b) normalized color space, and (c) the associated uncertainty. In (d), the result of manual global thresholding.

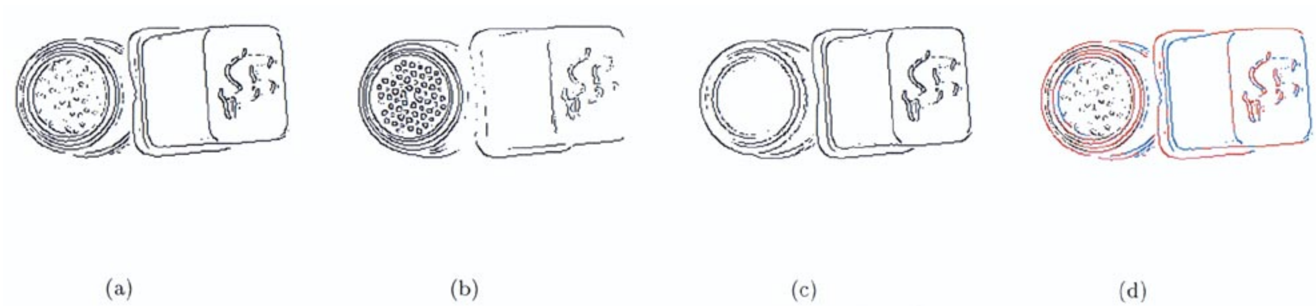

Fig. 3. Results of automatic local thresholding in (a) sensor, (b) normalized, and (c) opponent color space. (d) Result of color edge classification. Here, black edges are highlight edges, blue edges are geometry or shadow edges, red edges are material transitions.

exceeds $3\sigma_{\nabla F}$, we assume that there is a chance of 1% that this gradient modulus corresponds to no color transition:

$$\nabla C(x,y) = \begin{cases} 1 & \text{if } \nabla F(x,y) > 3\sigma_{\nabla F}(x,y) \\ 0 & \text{otherwise} \end{cases} \tag{20}$$

deriving a local threshold value.

The novelty of our approach is that the threshold value is automatically and locally adapted to the amount of uncertainty of the color invariant edge. For example, for c_1c_2 edges (unstable near the black point) at pixels with low intensity, the threshold value is automatically augmented. In this way, a local, noise-adaptive and automatic thresholding scheme is obtained.

B. Reflectance Based Edge Classification

In the previous sections, the effect of varying imaging circumstances have been analyzed first: c_1c_2 varies with a change in material and highlights, o_1o_2 with a change in material and object geometry, and RGB varies with a change in material, highlights and object geometry. Further, color invariant edges have been computed with their associated uncertainty.

As a consequence, we conclude that $\nabla\mathcal{C}_{RGB}$ [denoting the edge map in RGB-space with noise-adaptive thresholding corresponding to (20)] measures the presence of 1) shadow or geometry edges, 2) highlight edges, and 3) material edges. Further, $\nabla\mathcal{C}_{c_1c_2}$ (denoting the edge map in c_1c_2 normalized space) measures the presence of 2) highlight edges, and 3) material edges. Finally, $\nabla\mathcal{C}_{o_1o_2}$ measures the presence of 1) shadow or geometry edges and 3) material edges.

In this way, a taxonomy of color edge types is obtained (see Fig. 1). The color edge taxonomy is based upon the sensitivity of the color gradients with respect to the following imaging conditions: object geometry, shadows, highlights, and material.

Then the rule-based reflectance classifier is as follows:

IF $\nabla\mathcal{C}_{RGB} \neq 0$ AND $\nabla\mathcal{C}_{c_1c_2} = 0$
THEN classify as **shadow or geometry edge**
ELSE
IF $\nabla\mathcal{C}_{c_1c_2} \neq 0$ AND $\nabla\mathcal{C}_{o_1o_2} = 0$
THEN classify as **highlight edge**
ELSE
classify as **material edge**

only computed at color edge maxima using non-maxima suppression. Note again that the color edges are computed by (20). In this way, color edges and threshold values are automatically calculated in a parameter-free setting.

VI. Experiments

In this paper, experiments are conducted on still images (for illustration purposes) and video sequences recorded from complex scenes. To this end, in Section VI-A, we focus on color images taken from simple objects. In Section VI-B, image sequences are considered which are of average quality commonly used in digital consumer photography. In Section VI-C, video sequences are taken into account.

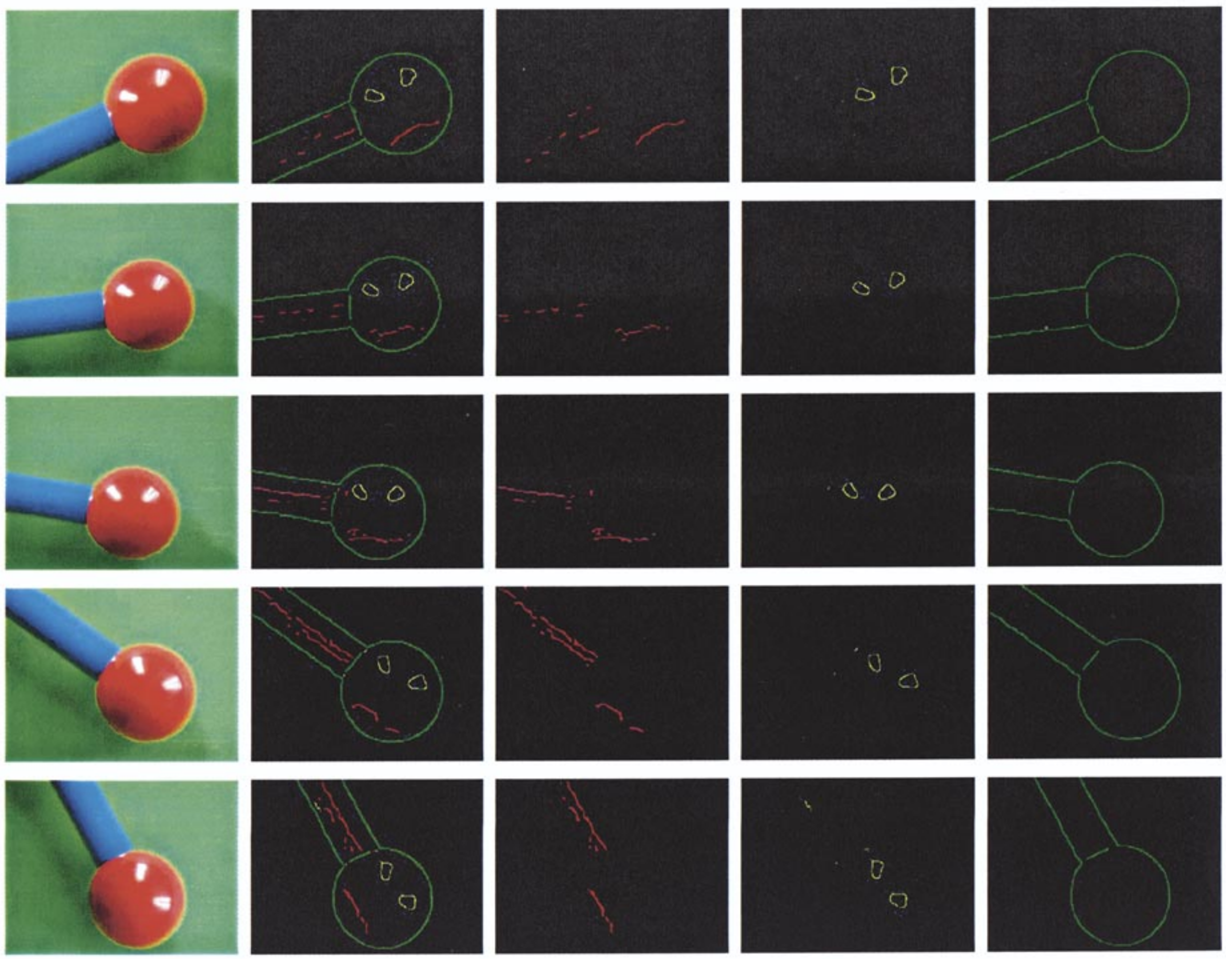

(a) (b) (c) (d) (e)

Fig. 4. Images from a sequence showing a toy against a background. (a) Original color image. (b) Classified edges. (c) Shadow and geometry edges. (d) Highlight edges. (e) Material edges.

A. Automatic Edge Classification in Still Images

For this experiment, an image is taken of two homogeneously colored plastic toys (red and green) against a blue paper background. The experiments are performed using a Sony 3CCD color camera XC-003P, Matrox Corona Frame-grabber, and four Osram 18 W "Lumilux deLuxe daylight" fluorescent light sources.

The image is shown in Fig. 2(a). The red object shows two circles enclosing a number of small specularities. A homogeneous, almost black, shadow region is visible at the right side of the green toy. The edge map, computed in the normalized color space, is shown in Fig. 2(b). As expected, the normalized color is sensitive to highlights. This results in the specularities at the red object. Further, it is experimentally established that the normalized color space is highly unstable at dark colors (i.e., low intensity). The uncertainty map of c_1c_2 edges is shown in Fig. 2(c). Note that uncertain values are depicted in black. In this way, regions with high intensity, in the original image, correspond to dark regions in the uncertainty map $\sigma_{\nabla_{c_1c_2}}$. In Fig. 2(d), the edge map is shown by thresholding the normalized color gradient by using a global threshold value. The most optimal threshold value has been selected by visual inspection. As color invariant instabilities are not stationary over the image (i.e., at each image location a different threshold is required), the specularities on the red cup having gradient moduli below the threshold value are shown up, whereas the noise edges in the dark region have gradient moduli values exceeding the threshold value. The experiment shows the inappropriateness of the use of a global (manual) threshold due to the local instabilities of color invariant edges.

The result of the newly proposed noise-adaptive thresholding scheme is shown in Fig. 3. In Fig. 3(a), the gradient is shown computed for the RGB color space. The image shows that many false edges are correctly suppressed while edges caused by material, geometry and specularity transitions are retained. Fig. 3(b) shows the result for normalized color space. Here, noise-adaptive thresholding correctly discards the edges present in the edge map of Fig. 2(b), while retaining the highlight edges on the bottom of the red cup. Fig. 3(c) shows the result of automatic thresholding for the opponent color space. As expected, the color space is invariant for highlights, which consequently do not show up in the edge map. Note that the opponent color space still depends on shadows and object geometry. Finally, in Fig. 3(d), the result of automatic edge labeling is shown. Black edges correspond to highlight edges, red edges to material

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

Fig. 5. Frames from a video showing a person against a textured background playing ping-pong. From left to right: (a) Original color frame. (b) Classified edges. (c) Material edges.(d) Shadow and geometry edges.

transitions, and blue edges to geometry changes. It is shown that the proposed method successfully classifies color edges in still images without the need for parameter settings.

B. Automatic Edge Classification in Digital Camera Sequences

In this section, we consider recordings from average quality which are commonly in use in digital consumer photography. Therefore, a sequence of five images have been recorded, shown in Fig. 4(a), by the Nikon Coolpix 950, a commercial digital camera of average quality. Further, the ScanDisk CompactFlash memory card has been used for image storage. The images have size 267 $\times$ 200 pixels. The digitization was done in 8 bits per color. The object is recorded against a cardboard background. Two light sources of average day-light color are used to illuminate the objects in the scene. There was no attempt to individual control the focus or the illumination for the objects in the scene. Images show a considerable amount of noise, shadows, shading, specularities and interreflections. As a result, recordings are of as average quality, corresponding to consumer photography.

From the edge classification result, shown in Fig. 4(b) it is observed that edges are introduced due to abrupt surface orientation and two highlights, caused by the two illumination sources. These highlights are well detected and classified, see Fig. 4(b) and Fig. 4(d). The highlight-edge map is fairly independent of shadows and abrupt surface orientation changes. Only a few erroneous edges are generated by small inter-reflections. Inter-reflections occur when an object receives the reflected light from other objects. Similarly, good performance is shown for the material transitions as depicted in Fig. 4(b) and 4(e).

C. Automatic Edge Classification in Video

In Fig. 5(a), six frames are shown from a standard video often used as a test sequence in the literature. It shows a person against a textured background playing ping-pong. The size of the image is 260 $\times$ 135. The images are of low quality. The frames are clearly contaminated by shadows, shading and inter-reflections. Note that each individual object-parts (i.e., t-shirt, wall, and table) is painted homogeneously with a distinct color. Further, that the wall is highly textured. The results of the proposed reflectance based edge classifier are shown in Fig. 5(b). As no highlights are present in the scene, the edge classifier discriminates edges in the color image to be one of the following types: 1) material edges shown in Fig. 5(c), and 2) shadow or geometry edges shown in Fig. 5(d). As one can see, the cast shadow of the person on the wall is well detected and classified. Also the geometry transitions on the t-shirt have been classified successfully. Material edges, as shown in Fig. 5(c), are well-defined ignoring radiometrical effects. Only interreflections and smoothly changing shading disturb the edge map slightly. From the observed results, it is concluded that the edge classifier discriminates the various edge types satisfactory. Only minor errors are caused when intensity change smoothly over a wide image range due to the local behavior of the edge classifier.

VII. Conclusion

Color information has been used to classify the physical nature of a color edge. A novel framework has been proposed for color edge detection and automatic noise-adaptive thresholding. The framework is derived from sensor noise analysis and propagation. Further, a parameter-free color edge classifier has been proposed labeling color transitions into the following types: 1) shadow, geometry or shading edges, 2) highlight edges, and 3) material edges. From the theoretical and experimental results, it is concluded that the proposed method successfully classifies color edges in video without the need for parameter settings.

Acknowledgment

The authors are grateful to A. W. M. Smeulders for advice and suggestions.

References

[1] J. Canny, "A computational approach to edge detection," *IEEE Trans. Pattern Anal. Machine Intell.*, vol. PAMI-8, no. 6, pp. 679–698, 1986.

[2] D. Ziou and S. Tabbone, "Edge detection techniques—an overview," *Pattern Recognit. Image Anal.*, vol. 8, no. 4, pp. 537–559, 1998.

[3] W. H. Tsang and P. W. M. Tsang, "Suppression of false edge detection due to specular reflection in color images," *Pattern Recognit. Lett.*, vol. 18, pp. 165–171, 1997.

[4] W. Zhang and F. Bergholm, "Multi-scale blur estimation and edge type classification for scene analysis," *Int. J. Comput. Vis.*, vol. 24, no. 3, pp. 219–250, 1997.

[5] J. Stander, R. Mech, and J. Ostermann, "Moving cast shadows detection for object segmentation," *IEEE Trans. Multimedia*, vol. 1, pp. 65–76, Mar. 1999.

[6] Th. Gevers and A. W. M. Smeulders, "Color based object recognition," *Pattern Recognit.*, vol. 32, pp. 453–464, 1999.

[7] S. A. Shafer, "Using color to separate reflection components," *Color Res. Applicat.*, vol. 10, no. 4, pp. 210–218, 1985.

[8] J. R. Taylor, *An Introduction to Error Analysis*. Mill Valley, CA: University Science, 1982.

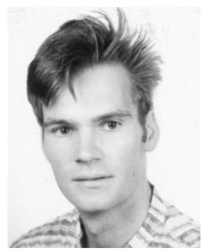

Theo Gevers (M'02) is an Associate Professor of computer science at the University of Amsterdam, The Netherlands. His main research interests are in the fundamentals of image database system design, image retrieval by content, theoretical foundation of geometric and photometric invariants and color image processing.

Harro Stokman, photograph and biography not available at the time of publication.