Edge and Corner Detection by Photometric Quasi-Invariants

Joost van de Weijer, *Member*, *IEEE*,
Theo Gevers, *Member*, *IEEE*, and
Jan-Mark Geusebroek

Abstract—Feature detection is used in many computer vision applications such as image segmentation, object recognition, and image retrieval. For these applications, robustness with respect to shadows, shading, and specularities is desired. Features based on derivatives of photometric invariants, which we will call *full invariants*, provide the desired robustness. However, because computation of photometric invariants involves nonlinear transformations, these features are unstable and, therefore, impractical for many applications. We propose a new class of derivatives which we refer to as *quasi-invariants*. These quasi-invariants are derivatives which share with full photometric invariants the property that they are insensitive for certain photometric edges, such as shadows or specular edges, but without the inherent instabilities of full photometric invariants. Experiments show that the quasi-invariant derivatives are less sensitive to noise and introduce less edge displacement than full invariant derivatives. Moreover, quasi-invariants significantly outperform the full invariant derivatives in terms of discriminative power.

Index Terms—Edge and feature detection, invariants, color.

✦

1 INTRODUCTION

FEATURE detection, such as edge and corner detection, plays an important role in many computer vision applications such as image segmentation, object recognition, and image retrieval [5]. A large number of feature detectors is based on the differential structure of images [1], [6], [9]. However, in real-world applications there are various physical phenomena which trigger differential-based features, such as shadows, shading, specularities, and object reflectance changes. It is important to differentiate between the various physical causes of a feature.

An improvement in color understanding was the introduction of the dichromatic reflection model by Shafer [14]. The model separates the reflected light into body reflection (object color) and surface reflection (specularities). This separation results in the classification of physical events, such as shadows and highlights. This is suited for photometric invariant segmentation, object recognition, and retrieval [3], [8], [10]. However, these methods are based on the zeroth order structure of images and mostly involve the analysis of the RGB-values in color histograms. For the photometric invariant theory to be applicable to differential-based operations, other methods are needed.

The connection between differential-based features and photometric invariance theory is proposed by Geusebroek et al. [2]. This work provides a set of photometric invariant derivative filters which uses them for invariant edge detection. However, the nonlinear transformations used to compute photometric invariants have several drawbacks such as instabilities and loss of discriminative power. These drawbacks limit the applicability of operations based on derivatives of these invariants. Traditionally, the effect of instabilities is suppressed by ad hoc thresholding of the transformed values [7], [11]. A more elaborate approach is to apply error propagation through the various color spaces to compensate for the undesired effects of instabilities and nonlinearities of the different photometric invariant spaces [4]. However, this approach is based on a proper noise estimation system which is not always available in practice.

In this paper, we propose a new class of derivatives which we refer to as photometric quasi-invariants. These derivatives link derivative-based operations to the theory of photometric invariance. Quasi-invariants are derived from the dichromatic reflection model and are proven to differ from full photometric invariants by a scaling factor. These quasi-invariants do not have the inherent instabilities of full photometric invariants, and from theoretical and experimental results it is shown that quasi-invariants have better noise characteristics, discriminative power, and introduce less edge displacement than full photometric invariants. The lack of full photometric invariance limits the applicability of quasi-invariants to methods which are based on a single image, such as edge and corner detection. Quasi-invariants cannot be used for applications in which responses between multiple images are compared, such as invariant object recognition.

2 THE DICHROMATIC REFLECTION MODEL

In this section, the dichromatic reflection model is discussed [14]. The dichromatic model divides the reflection in the body (object color) and surface reflection (specularities or highlights) component for optically inhomogeneous materials. Assuming a known illuminant, $\mathbf{c}^i = (\alpha, \beta, \gamma)^T$, and neutral interface reflection, the RGB vector, $\mathbf{f} = (R, G, B)^T$, can be seen as a weighted summation of two vectors,

$$\mathbf{f} = e(m^b \mathbf{c}^b + m^i \mathbf{c}^i), \tag{1}$$

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

$$\mathbf{f} = e m^b \mathbf{c}^b, \tag{2}$$

which is the well-known Lambertian reflection. For more on the validity of the photometric assumptions see [2], [3], [14] and for calibration [4].

From the dichromatic reflection model, photometric invariants can be derived (e.g., normalized RGB, hue). These invariants have the disadvantage that they are unstable; normalized RGB is unstable near zero intensity and hue is undefined on the black-white axis. The instabilities can be avoided by analyzing the RGB values in the RGB-histogram [8], [10]. That proved to be rather difficult and slow since you need a meaningful segmentation to generate a meaningful histogram and a meaningful histogram to get a good segmentation.

Instead of looking at the zeroth order structure (the RGB-values), we focus in this paper on the first order structure of the image. A straightforward extension of the photometric invariance theory to first order filters can be obtained by taking the derivative of the invariants. However, these filters would inherit the undesired instabilities of the photometric invariants. Therefore, we propose an alternative way to arrive at photometric derivatives by analyzing the spatial derivative of the dichromatic reflection model.

The spatial derivative of the dichromatic reflection model (1) gives the photometric derivative structure of the image:

$$\mathbf{f}_x = e m^b \mathbf{c}_x^b + (e_x m^b + e m_x^b) \mathbf{c}^b + (e m_x^i + e_x m^i) \mathbf{c}^i. \tag{3}$$

Here, the subscript indicates spatial differentiation. Since we assume a known illuminant and neutral interface reflection, $\mathbf{c}^i$ is independent of x. The derivative in (3) is a summation of three weighted vectors, successively caused by body reflectance, shading-shadow, and specular change. Further, we assume that shadows are not significantly colored.

● *The authors are with the Intelligent Sensory Information Systems, Faculty of Science, University of Amsterdam, Kruislaan 403, 1098 SJ Amsterdam, The Netherlands. E-mail: {joostw, gevers}@science.uva.nl.*
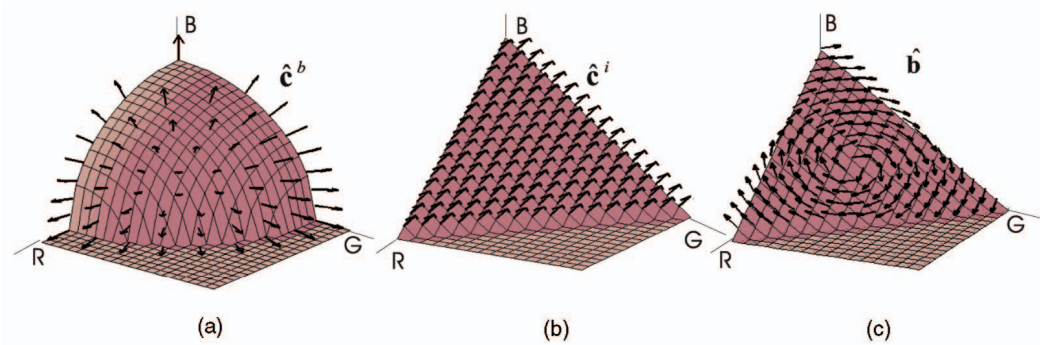
Fig. 1. (a) Shadow-shading direction $\hat{\mathbf{c}}^b$, (b) specular direction $\hat{\mathbf{c}}^i$, and (c) hue direction $\hat{\mathbf{b}}$.
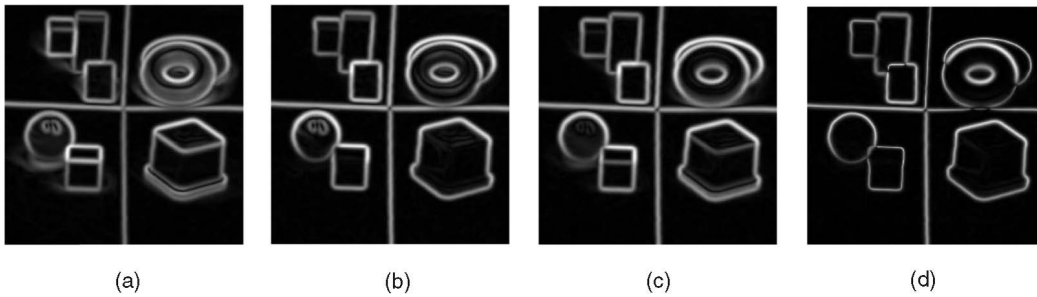

Fig. 2. Various derivatives applied to Fig. 4a: (a) color gradient ($\mathbf{f}_x$), (b) shadow-shading quasi-invariant ($\mathbf{S}_x^c$), (c) the specular quasi-invariant ($\mathbf{O}_x^c$), and (d) the specular-shadow-shading quasi-invariant ($\mathbf{H}_x^c$).

In fact, the direction of the shadow-shading changes (Fig. 1a) follows from (2). In the absence of interface reflection, the direction of $\mathbf{c}^b$ coincides with the direction of $\hat{\mathbf{f}} = \frac{1}{\sqrt{R^2+G^2+B^2}}(R,G,B)^T$. The hat is used to denote unit vectors. The shadow-shading direction is the multiplication of two scalars denoting two different physical phenomena. First, $e_x m^b$ indicates a change in intensity which corresponds to a shadow edge. And em_x^b is a change in the geometry coefficient which represents a shading edge.

Another direction is the specular direction c^i in which changes of the specular geometry coefficient m_x^i occur. In Fig. 1b, $\mathbf{c}^i$ is depicted for the case of a white light source for which $\hat{\mathbf{c}}^i = \frac{1}{\sqrt{3}}(1,1,1)^T$. The specular direction is multiplied by two factors. First, em_x^i is a change of geometric coefficient caused by changes in the angles between viewpoint, object, and light source. Second, the term $e_x m^i$ representing a shadow edge on top of a specular reflection.

Having the direction of two of the causes of an edge, we are able to construct a third direction which is perpendicular to these two vectors (Fig. 1c). This direction, named hue direction $\hat{\mathbf{b}}$, is computed by the outer product:

$$\hat{\mathbf{b}} = \frac{\hat{\mathbf{f}} \times \hat{\mathbf{c}}^i}{\left| \hat{\mathbf{f}} \times \hat{\mathbf{c}}^i \right|}. \qquad (4)$$

If $\hat{\mathbf{f}}$ and $\hat{\mathbf{c}}^i$ are parallel, we define $\hat{\mathbf{b}}$ to be the zero vector. Note that the hue direction is *not* equal to the direction in which changes of the body reflectance occur, $\hat{\mathbf{c}}_x^b$. It is perpendicular to the two other causes of an edge. Hence, changes in the hue direction can only be attributed to a body reflectance change.

In conclusion, changes in the reflection manifest themselves as edges in the image. There are three causes for an edge in an image: a hue change, a shadow-shading edge, or a specular change. We indicated three directions: the shadow-shading direction, the specular direction, and the hue direction. These directions are the same as the directions indicated by Klinker and Shafer [8] for the use of image segmentation. We use these directions for the construction of photometric invariant spatial derivatives.

3 PHOTOMETRIC VARIANTS AND QUASI-INVARIANTS

In this section, the goal is to propose a new set of photometric variants and quasi-invariants. To this end, the derivative of an image, $\mathbf{f}_x = (R_x, G_x, B_x)^T$, is projected on three directions found in the previous section. We will call these projections *variants*. For example, the projection of the derivative on the shadow-shading direction results in the shadow-shading variant. By removing the variance from the derivative of the image, we construct a complementary set of derivatives which we will call *quasi-invariants*.

The projection of the derivative on the shadow-shading direction is called the shadow-shading variant and is defined as

$$\mathbf{S}_x = \left(\mathbf{f}_x \cdot \hat{\mathbf{f}} \right) \hat{\mathbf{f}}. \qquad (5)$$

The dot indicates the vector inner product. The second $\hat{\mathbf{f}}$ indicates the direction of the variant. The shadow-shading variant is the part of the derivative which could be caused by shadow or shading. Due to correlation of the hue and specular direction with the shadow-shading direction, part of $\mathbf{S}_x$ might be caused by changes in hue or specular reflection.

What remains after subtraction of the variant is called the shadow-shading quasi-invariant, indicated by superscript c.

$$\mathbf{S}_x^c = \mathbf{f}_x - \mathbf{S}_x. \qquad (6)$$

The quasi-invariant $\mathbf{S}_x^c$ consists of that part of the derivative which is *not* caused by shadow-shading edges (Fig. 2b), hence, only contains specular and hue edges.

The same reasoning can be applied to the specular direction and results in the specular variant and the specular quasi-invariant

$$\begin{aligned} \mathbf{O}_x &= \left(\mathbf{f}_x \cdot \hat{\mathbf{c}}^i \right) \hat{\mathbf{c}}^i, \\ \mathbf{O}_x^c &= \mathbf{f}_x - \mathbf{O}_x. \end{aligned} \qquad (7)$$

The specular quasi-invariant is insensitive to highlight edges (Fig. 2c).

Finally, we can construct the shadow-shading-specular variant and quasi-invariant by projecting the derivative on the hue direction
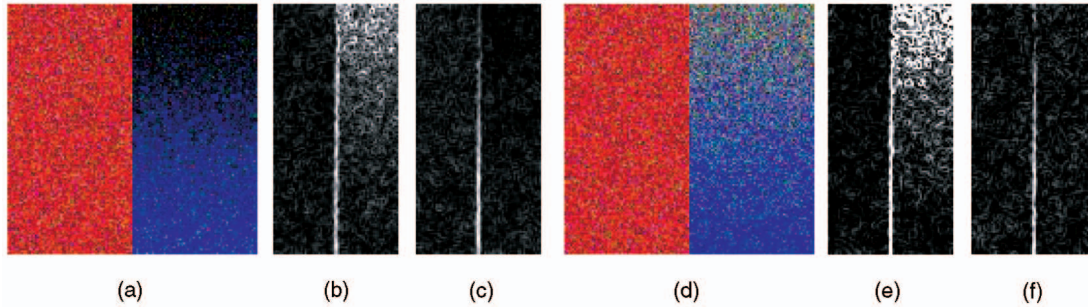
Fig. 3. (a) Red-blue edge, with a decreasing intensity of the blue patch going in the upward direction. Response of (b) normalized RGB derivative, and (c) shadow-shading quasi-invariant ($\mathbf{S}_x^c$). (d) Red-blue edge, with decreasing saturation going in the upward direction. Response of (e) hue derivative (h_x), and (f) specular-shadow-shading quasi-invariant ($\mathbf{H}_x^c$).

$$\mathbf{H}_x^c = \left(\mathbf{f}_x \cdot \hat{\mathbf{b}}\right)\hat{\mathbf{b}},$$
$$\mathbf{H}_x = \mathbf{f}_x - \mathbf{H}_x^c. \tag{8}$$

$\mathbf{H}_x^c$ does not contain specular or shadow-shading edges (Fig. 2d).

4 RELATIONS OF QUASI-INVARIANTS WITH FULL INVARIANTS

In this section, the resemblances and differences are analyzed between quasi-invariants and full invariants. A geometrical relation in RGB-space between the two is found by investigating underlying color space transformations. Conclusions with respect to stability are made. With stability it is meant that small changes in the RGB-cube do not cause large jumps in the invariant space. Further, we discuss the characteristics of quasi-invariants.

4.1 Spherical Color Space

An orthogonal transformation which has the shadow-shading direction as one of its components is the spherical coordinate transformation. Transforming the RGB-color space results in the spherical color space or $r\theta\varphi$-color space. The transformations are,

$$r = \sqrt{R^2 + G^2 + B^2} = |\mathbf{f}|$$
$$\theta = \arctan\left(\frac{G}{R}\right) \tag{9}$$
$$\varphi = \arcsin\left(\frac{\sqrt{R^2 + G^2}}{\sqrt{R^2 + G^2 + B^2}}\right).$$

Since r is pointing in the shadow-shading direction, its derivative corresponds to $\mathbf{S}_x$.

$$r_x = \frac{RR_x + GG_x + BB_x}{\sqrt{R^2 + G^2 + B^2}} = \mathbf{f}_x \cdot \hat{\mathbf{f}} = |\mathbf{S}_x|. \tag{10}$$

The quasi-invariant $\mathbf{S}_x^c$ is the derivative energy in the plane perpendicular to the shadow-shading direction. The derivative in the $\theta\varphi$-plane is given by

$$\left|\mathbf{S}_x^c\right| = \sqrt{(r\varphi_x)^2 + (r\sin\varphi\,\theta_x)^2}$$
$$= r\sqrt{(\varphi_x)^2 + (\sin\varphi\,\theta_x)^2}. \tag{11}$$

To conserve the metric of RGB-space, the angular derivatives are multiplied by their corresponding scale factors which follow from the spherical transformation. For matte surfaces, both θ and φ are independent of m^b (substitution of (2) in (9)). Hence, the part under the root is a shadow-shading invariant.

By means of the spherical coordinate transformation, a relation between the quasi-invariant and the full invariant is found. The difference between the quasi-invariant $\left|\mathbf{S}_x^c\right|$ and the full invariant $s_x = \sqrt{(\varphi_x)^2 + (\sin\varphi\,\theta_x)^2}$ is the multiplication with r which is the $L2$ norm

for the intensity (see (9)). In geometrical terms, the derivative vector which remains after subtraction of the part in the shadow-shading direction is *not* projected on the sphere to produce an invariant. This projection introduces the instability of the full shadow-shading invariants for low intensities,

$$\lim_{r \to 0} s_x \quad \text{does not exist}$$
$$\lim_{r \to 0} \left|\mathbf{S}_x^c\right| = 0. \tag{12}$$

The first limit follows from the nonexistence of the limit for both φ_x and θ_x at zero. The second limit can be concluded from $\lim_{r \to 0} r\varphi_x = 0$ and $\lim_{r \to 0} r\theta_x = 0$. Concluding, the multiplication of the full-invariant with $|\mathbf{f}|$ resolves the instability.

An example of the responses for the shadow-shading invariant and quasi-invariant is given in Fig. 3. In Fig. 3a, a synthetic image of a red-blue edge is depicted. The blue intensity decreases along the y-axis. Gaussian uncorrelated noise is added to the RGB channels. In Fig. 3b, the normalized RGB response is depicted and the instability for low intensities is clearly visible. For the shadow-shading quasi-invariant (Fig. 3c), no instability occurs and the response just diminishes for low intensities. Note that the instable region is particularly inconvenient because shadow-shading edges tend to produce low-intensity areas.

4.2 Opponent Color Space

The orthonormal transformation which accompanies the specular variant is known as the opponent color space. For a known illuminant $\mathbf{c}^i = (\alpha, \beta, \gamma)^T$, it is given by

$$o1 = \frac{\beta R - \alpha G}{\sqrt{\alpha^2 + \beta^2}}$$
$$o2 = \frac{\alpha\gamma R + \beta\gamma G - (\alpha^2 + \beta^2)B}{\sqrt{(\alpha^2 + \beta^2 + \gamma^2)(\alpha^2 + \beta^2)}} \tag{13}$$
$$o3 = \frac{\alpha R + \beta G + \gamma B}{\sqrt{\alpha^2 + \beta^2 + \gamma^2}}.$$

The relations with the variant and its complement are $|\mathbf{O}_x| = o3_x$ and $\left|\mathbf{O}_x^c\right| = \sqrt{o1_x^2 + o2_x^2}$.

4.3 The Hue Saturation Intensity Space

As discussed in Section 3, the shadow-shading-specular quasi-invariant is both perpendicular to the shadow-shading direction and the specular direction. An orthogonal transformation which satisfies this constraint is the hue-saturation-intensity transformation. It is actually a polar transformation on the opponent color axis $o1$ and $o2$.

$$h = \arctan\left(\frac{o1}{o2}\right)$$
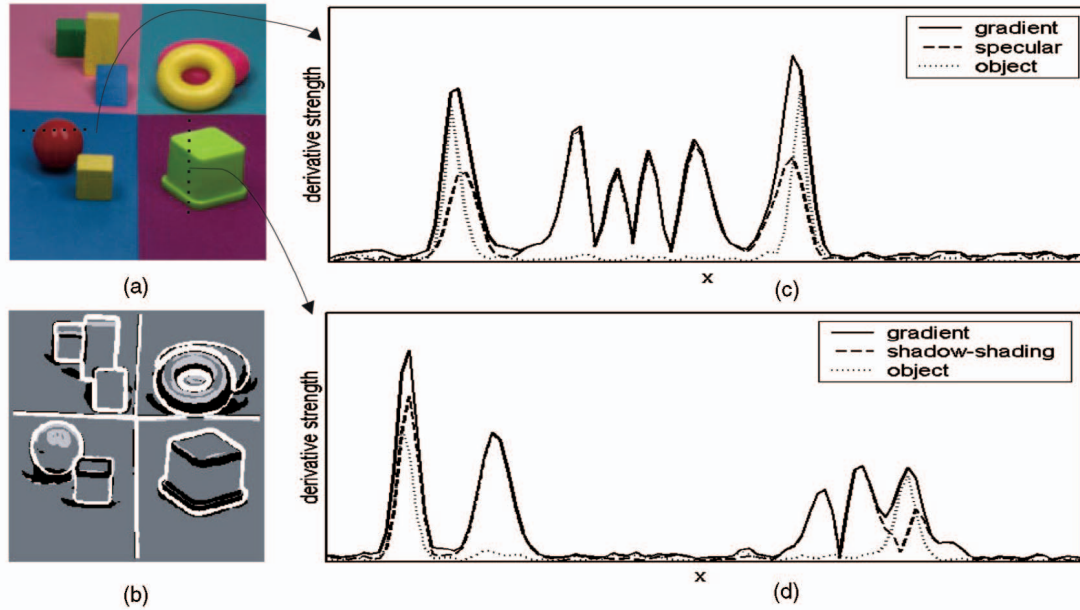$$s = \sqrt{o1^2 + o2^2} \tag{14}$$
$$i = o3.$$

Fig. 4. (a) Input image with two superimposed dotted lines which are plotted in the images (c) and (d). (b) Edge classification result, with white object edges, black shadow edges, and light gray specular edges. (c) and (d) The derivative strength along lines indicated in (a).

The changes of h occur in the hue direction and, hence, the derivative in the hue-direction is equal to the shadow-shading-specular quasi-invariant,

$$\left|\mathbf{H}_x^c\right| = s \cdot h_x. \tag{15}$$

The multiplication with the scale factor s follows from the fact that for polar transformations the angular derivative is multiplied by the radius.

The hue, h, is a well-known full shadow-shading-specular invariant. Equation (15) provides a link between the derivative of the full invariant h_x and the quasi-invariant $\left|\mathbf{H}_x^c\right|$. A drawback of hue is its undefinedness for points on the black-white axis, i.e., for small s. Therefore, the derivative of hue is unbounded. In Section 3, we derived the quasi-invariant as a linear projection of the spatial derivative. For these projections, it holds that $0 < \left|\mathbf{H}_x^c\right| < \left|\mathbf{f}_x\right|$ and, hence, the shadow-shading specular quasi-invariant is bounded. It should be mentioned that small changes round the gray axis, result in large changes of the direction or "color" of the derivative, e.g., from blue to red, in both the quasi-invariant and the full invariant. However, the advantage of the quasi-invariant is that the norm remains bounded for these cases. For example, in Fig. 3d a red-blue edge is depicted. The blue patch becomes more achromatic along the y-axis. The instability for gray values is clearly visible in Fig. 3e whereas, in Fig. 3f, the response of the quasi-invariant remains stable.

4.4 Characteristics of Quasi-Invariants

Full invariants are invariant with respect to a physical photometric parameter like for instance the geometric term m^b in the case of normalized RGB. Hence, the first order derivative response of such invariants does not contain any shadow-shading variation. Our approach determines the direction in the RGB-cube in which shadow-shading edges exhibit themselves. The derivative caused by other than shadow-shading edges can than be computed. Sharing with full invariants, the property that shadow-shading edges are ignored. However, the quasi-invariant is not invariant with respect to m_b. For the shadow-shading quasi-invariant, subtraction from (3) of the part in the shadow-shading direction $\mathbf{c}^b$ results in

$$\mathbf{f}_x = em^b\left(\mathbf{c}_x^b - \mathbf{c}_x^b \cdot \hat{\mathbf{c}}^b\right), \tag{16}$$

which is clearly not invariant for m^b and e. Also, in a similar way, the specular-shadow-shading quasi-invariant can be proven to be dependent on m^b and e.

The dependency of the quasi-invariants on m^b and e limits their applicability. They cannot be used for applications where edge responses are compared under different circumstances, such as content-based image retrieval. However, they can be used in applications which are based on a single frame, such as shadow-edge insensitive image segmentation, shadow-shading-specular independent corner detection, and edge classification.

A major advantage of the quasi-invariants is that their response to noise is independent of the signal. In the case of additive uniform noise, the noise in the quasi invariants is also additive and uniform since it is a linear projection of the derivative of the image. This means that the noise distortion is constant over the image. In Section 4, it was shown that the full invariants differ from the quasi-invariants by scaling with a signal depended factor (the intensity or saturation). And, hence, their noise response is also signal depended. Typically, the shadow-shading full invariant exhibits high noise distortion round low intensities while the shadow-shading-specular full invariant has high noise dependency for points around the achromatic axis. This is shown in Fig. 3. The uneven levels of noise throughout an image hinder further processing.

A second advantage of photometric variants and quasi-invariants is that they are expressed in the same units (i.e., being projections of the derivative they are in RGB-value per pixel). This allows for a quantitative comparison of their responses. An example

TABLE 1
The Displacement, Δ, and the Percentage of Missed Edges, ε, for Five Different Edge Detectors

std. noise →	5		20	
detector ↓	Δ	ε	Δ	ε
S_x^c	0.043	0.99%	0.43	10%
rgb	0.21	2.0%	1.1	18%
H_x^c	0.35	5.8%	0.98	20%
hue	0.85	9.8%	2.1	34%
$RGB - gradient$	0.003	.07%	0.08	2.0%

Gaussian noise of standard deviation 5 and 20 was added.
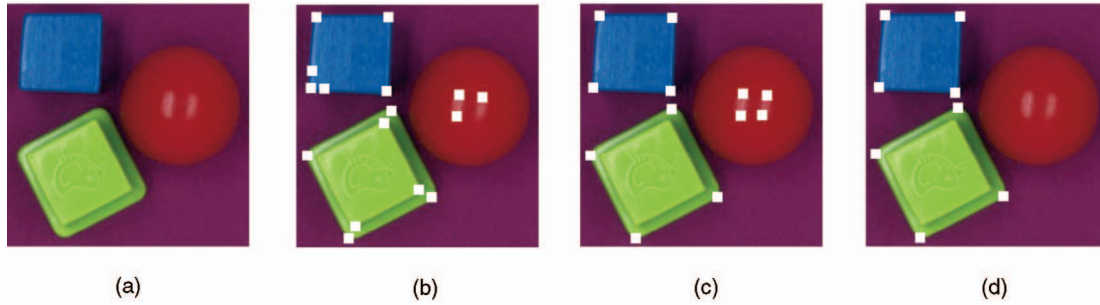
Fig. 5. (a) Input image and corner detector results based on (b) RGB gradient ($\mathbf{f}_x$), (c) shadow-shading quasi-invariant ($\mathbf{S}_x^c$), and (d) shadow-shading-specular quasi-invariant ($\mathbf{H}_x^c$).

is given in Fig. 4. Responses along two lines in the image are enlarged in Figs. 4c and 4d. The line in Fig. 4c crosses two object edges and several specular edges. It nicely shows that the specular-variant almost perfectly follows the total derivative energy for the specular edges in the middle of the line. In Fig. 4d, a line is depicted which crosses two object edges and three shadow-shading edges. Again, the shadow-shading variant follows the gradient for the three shading edges. A simple classification scheme results in Fig. 4b. Note that full-invariants cannot be compared quantitatively because they have different units.

5 EXPERIMENTS

We compare the performance of the quasi-invariants with the full invariants according to the following criteria 1) stability, 2) edge displacement, and 3) discriminative power. For the improved stability, a mathematical proof is given in chapter 5. Here, we will test the invariants on edge displacement and discriminative power.

Since the specular quasi-invariant is well-known and it does not counterpart a full invariant, its performance is not investigated here. The experiments were performed with normalized RGB, $c_1 c_2 c_3$, $l_1 l_2 l_3$, hue, Cw, and Hw [2], [3]. The results for the invariants $c_1 c_2 c_3$, $l_1 l_2 l_3$, Cw, and Hw were similar or worse than the results for normalized RGB and hue. Therefore, we have chosen normalized RGB and hue as exemplary for the set of invariants and compared them with the quasi-invariants. Implementation details of the quasi-invariants can be found in [15]. For the experiments, a white light source $\hat{\mathbf{c}}^i = \frac{1}{\sqrt{3}}(1,1,1)^T$ is used.

5.1 Edge Detection

First, we compare the edge detection performance of the quasi-invariants with the invariants from literature. These results can also be seen as an indication of the loss of discrimination due to invariance. Edge detection is performed between the 1,012 different colors from the PANTONE [12] color system. Every one of the 1,012 different RGB-values is combined with all other RGB-values, resulting in a total of $N = 1,012 * 1,011/2 = 511,566$ edges of $M = 25$ pixels length. The edge position is determined by computing the maximum response path of the derivative energy in a region of 20 pixels around the actual edge. This results in an edge estimation which is compared with the actual edge. We define two error measures. First, the average pixel displacement Δ,

$$\Delta = \frac{\sum\limits_{\{x_{i,j}; |x_{i,j}-x_0|>0.5\}} |x_{i,j} - x_0|}{N \cdot M}, \quad (17)$$

in which $x_{i,j}$ is the jth edge pixel of the ith edge. Because the actual edge is located between two pixels, displacements equal to .5 pixels are considered as a perfect match. The second error measure is the percentage of missed edges, ε. An edge was classified missed as the variation over one edge,

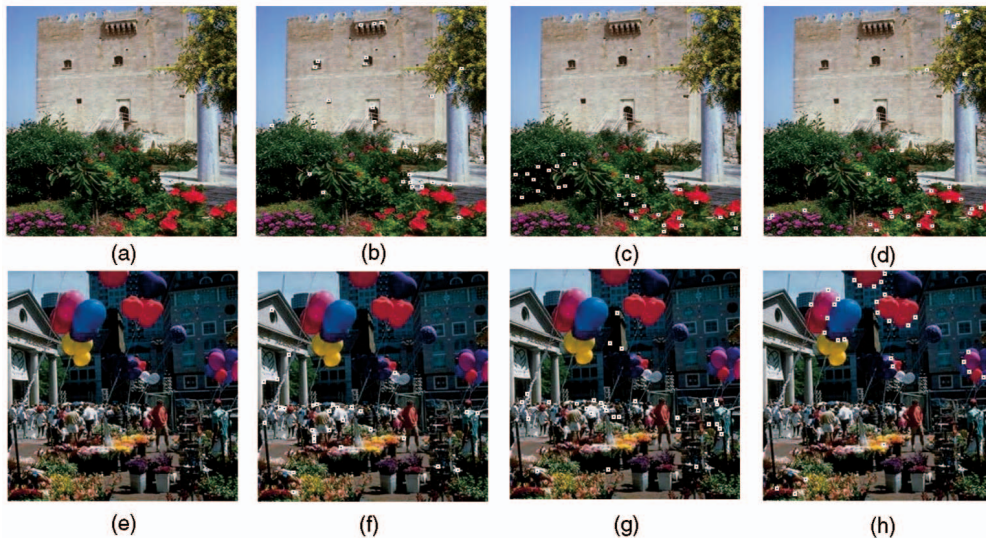$$\mathrm{var}(i) = \frac{1}{M}\sum_{j=1}^{M}\left| x_{i,j} - \frac{1}{M}\sum_k x_{i,k}\right|, \quad (18)$$

Fig. 6. (a) and (e) Input images. Corner detection based on (b) RGB gradient ($\mathbf{f}_x$), (c) normalized RGB, (d) shadow-shading quasi-invariant ($\mathbf{S}_x^c$), (f) RGB gradient ($\mathbf{f}_x$), (g) hue full invariant (h_x), and (h) shadow-shading quasi-invariant ($\mathbf{H}_x^c$).

is larger than 1 pixel. For the Gaussian derivative, a scale $\sigma = 1$ is chosen. The experiments were performed with uncorrelated Gaussian noise of standard deviation 5 and 20.

The results are depicted in Table 1. For both cases, the shadow-shading and shadow-shading-specular edges, the quasi-invariants substantially outperform the invariants. For comparison, the results without invariance based on the RGB gradient, $|\mathbf{f}_x|$, are inserted. Obviously, the RGB gradient has the best discriminative power. However, it will also find many edges which are caused due to scene incidental events.

To provide more insight in what kind of edges were still detected, we computed the average Euclidean RGB difference of the missed edges for the case with Gaussian noise with a standard deviation of 5. With $d_{ij} = |\mathbf{f}_i - \mathbf{f}_j|$ the Euclidean distance between patch i and j. For the RGB gradient-based method, we obtained an average distance of $\bar{d} = 4.6$, for the shadow-shading quasi-invariant $\bar{d} = 86$ and $\bar{d} = 109$ for the shadow-shading-specular invariant.

5.2 Photometric Invariant Corner Detection

Derivatives based on full photometric invariants are, due to their instability, unreliable input for geometrical operations such as photometric invariant corner detection, orientation estimation, curvature estimation, etc. Quasi-invariants, on the other hand are expected to be more stable in combination with geometrical operations. We used the following straightforward extension of the Harris corner detector [6] for color images

$$H\mathbf{f} = \overline{\mathbf{f}_x^T \mathbf{f}_x}\; \overline{\mathbf{f}_y^T \mathbf{f}_y} - \overline{\mathbf{f}_x^T \mathbf{f}_y}^2 - k(\overline{\mathbf{f}_x^T \mathbf{f}_x} + \overline{\mathbf{f}_y^T \mathbf{f}_y})^2. \qquad (19)$$

The overline indicates a Gaussian averaging window. The corner detection results are given in Fig. 5. The shadow-shading quasi-invariant detector does not find shadow-shading corners whereas the shadow-shading-specular quasi-invariant also ignores the specular corners.

In Fig. 6, the 30 most prominent Harris corners are detected for two real-world images (Corel gallery). The detected points can be used as interest points for object recognition [13]. Note that the images break several of the assumptions of the dichromatic reflection model (1). They do not have a known illuminant, nor are they taken with a linear acquisition system. The results for the full invariants are dominated by their instabilities. The shadow-shading full invariant is unstable in the low intensity areas and consequently finds most of the interest points in this area. The shadow-shading specular full invariant is unstable along the whole gray axis, which leads to false corners in gray areas. The RGB gradient method focusses on large RGB value changes which mostly coincide with light-dark transition which are rarely the most discriminative points. It is apparent that the quasi-invariants (Figs. 6d and 6h) suppress unwanted photometric variation and focus on body reflectance changes only.

6 CONCLUSIONS

In this paper, we proposed a set of quasi-invariant derivatives. These derivative filters are combined with derivative-based feature detectors to perform photometric invariant feature detection. Experiments show that they significantly outperform feature detection based on full invariants on both stability and discriminative power.

REFERENCES

[1] J. Canny, "A Computational Approach to Edge Detection," *IEEE Trans. Pattern Analysis and Machine Intelligence,* vol. 8, no. 6, pp. 679-698, 1986.

[2] J.M. Geusebroek, R. van den Boomgaard, A.W.M. Smeulders, and H. Geerts, "Color Invariance," *IEEE Trans. Pattern Analysis and Machine Intelligence,* vol. 23, no. 12, pp. 1338-1350, Dec. 2001.

[3] T. Gevers and A. Smeulders, "Color Based Object Recognition," *Pattern Recognition,* vol. 32, pp. 453-464, Mar. 1999.

[4] T. Gevers and H.M.G. Stokman, "Classification of Color Edges in Video into Shadow-Geometry, Highlight, or Material Transitions," *IEEE Trans. Multimedia,* vol. 5, no. 2, pp. 237-243, 2003.

[5] R.M. Haralick and L.G. Shapiro, *Computer and Robot Vision,* vol. II. Addison-Wesley, 1992.

[6] C. Harris and M. Stephens, "A Combined Corner and Edge Detector," *Proc. Fourth Alvey Vision Conf.,* vol. 15, pp. 147-151, 1988.

[7] G. Healey, "Segmenting Images Using Normalized Color," *IEEE Trans. Systems, Man, and Cybernetics,* vol. 22, pp. 64-73, 1992.

[8] G.J. Klinker and S.A. Shafer, "A Physical Approach to Color Image Understanding," *Int'l J. Computer Vision,* vol. 4, pp. 7-38, 1990.

[9] T. Lindeberg, "Feature Detection with Automatic Scale Selection," *Int'l J. Computer Vision,* vol. 30, no. 2, pp. 77-116, 1998.

[10] B.A. Maxwell and S.A. Shafer, "Physics-Based Segmentation of Complex Objects Using Multiple Hypothesis of Image Formation," *Computer Vision and Image Understanding,* vol. 65, pp. 265-295, 1997.

[11] Y. Ohta, T. Kanade, and T. Sakai, "Color Information for Region Segmentation," *Computer Graphics and Image Processing,* vol. 13, pp. 222-241, 1980.

[12] PANTONE, ed. 1992-1993, groupe BASF, Paris, France. PANTONE is a trademark of PATONE Inc.

[13] C. Schmid, R. Mohr, and C. Bauckhage, "Evaluation of Interest Point Detectors," *Int'l J. Computer Vision,* vol. 37, no. 2, pp. 151-172, 2000.

[14] S.A. Shafer, "Using Color to Seperate Reflection Components," *COLOR Research and Application,* vol. 10, no. 4, pp. 210-218, Winter 1985.

[15] J. van de Weijer, T. Gevers, and J.M. Geusebroek, "Color Edge Detection by Photometric Quasi-Invariants," *Proc. Int'l Conf. Computer Vision,* pp. 1520-1526, 2003.

▷ **For more information on this or any other computing topic, please visit our Digital Library at** www.computer.org/publications/dlib.