Boosting Saliency in Color Image Features

J. van de Weijer   Th. Gevers

Intelligent Sensory Information Systems
Faculty of Science, University of Amsterdam
Kruislaan 403, 1098 SJ Amsterdam, The Netherlands
{joostw, gevers }@science.uva.nl

Abstract

The aim of salient point detection is to find distinctive events in images. Salient features are generally determined from the local differential structure of images. They focus on the shape saliency of the local neighborhood. The majority of these detectors is luminance based which has the disadvantage that the distinctiveness of the local color information is completely ignored. To fully exploit the possibilities of color image salient point detection, color distinctiveness should be taken into account next to shape distinctiveness.

In this paper color distinctiveness is explicitly incorporated into the design of saliency detection. The algorithm, called color saliency boosting, is based on an analysis of the statistics of color image derivatives. Isosalient color derivatives can be closely approximated by ellipsoidal surfaces in color derivative space. Based on this remarkable statistical finding, isosalient derivatives are transformed by color boosting to have equal impact on the saliency. Color saliency boosting is designed as a generic method easily adaptable to existing feature detectors. Results show that substantial improvements in information content are acquired by targeting color salient features. Further, the generality of the method is illustrated by applying color boosting to multiple existing saliency methods.

1   Introduction

Indexing objects and object categories as a collection of salient points has been successfully applied to image matching, content-based retrieval, learning and recognition [2], [10], [11], [14], [16], [19]. Salient points are local features in the image which exhibit geometrical structure, such as T-junctions, corners, and symmetry points. The aim of salient point detection is to represent objects more concisely and robust to varying viewing conditions, such as changes due to zooming, rotation, and illumination changes. Applications based on salient points are generally composed of three phases: 1. a feature detection phase locating the fea-

tures. 2. an extraction phase in which local descriptions are extracted at the detected locations and 3. a matching phase in which the extracted descriptors are matched against a database of descriptors. In this paper, the focus is to improve the salient point detection phase.

Although the majority of image data is in color format nowadays, most salient point detectors are still luminance based. They typically focus on shape saliency rather than color saliency [9], [20]. They focus on corners without distinguishing low-salient black-and-white corners from high-salient red-green corners. Only recently color information has been incorporated in the detection phase. Montesinos et al. [12] propose an extension of the luminance Harris corner detector to color [4]. Heidemann [5] incorporates color into the generalized symmetry transform proposed by Reisfeld et al. [13]. Both methods achieve a performance gain for near isoluminant events. However, since the luminance axis remains the major axes of color variation in the RGB-cube, results do not differ greatly from luminance based feature detection. Itti et al. [6] use color contrast as a clue for salience. Their method is based on a zero-order signal (normalized red, green , blue, yellow). It is however not easily extendable to differential-based features.

For the evaluation of salient point detectors Schmid et al. [15] proposed two criteria: 1. *repeatability*, salient point detection should be stable under the varying viewing conditions, such as geometrical changes and photometric changes. 2. *distinctiveness*, salient points should focus on events with a low probability of occurrence. Most salient point detectors are designed according to these criteria. They focus on two dimensional structures, such as corners, which are stable and distinctive at the same time. Although color is believed to play an important role in attributing saliency [7], the explicit incorporation of color distinctiveness into the design of salient points detectors has, to our knowledge, not been done.

A remarkable phenomenon appears when studying the statistics of color image derivatives. In histograms of color derivatives, points of equal frequency form regular struc-

tures. The derivatives play two roles in salient point detection. Firstly, they are input to the saliency function, which based on local derivatives probes for salient structures. Secondly, they are part of the extracted local features, on which the distinctiveness of the salient point detector is based. This double role, together with the statistical finding described above, leads to the following question: How can we exploit the regularity of the distinctiveness of color image derivatives to improve salient feature detection ?

In this paper we aim to incorporate color distinctiveness into salient point detection. The extension should be general and hence be easy to incorporate in existing salient point detectors. For a color image, with values $\mathbf{f} = (R, G, B)^T$, salient points are the maxima of the saliency map, which compares the derivative vectors in a neighborhood fixed by scale σ,

$$s = H^\sigma (\mathbf{f}_x, \mathbf{f}_y) \tag{1}$$

where H is the saliency function and the subscript indicates differentiation with respect to the parameter. This type of saliency maps include [1], [4], [5], [9]. The impact of a derivative vector on the outcome of the local salience depends on its vector norm, $|\mathbf{f}_x|$. Hence, vectors with equal norm have an equal impact on the local saliency. Rather than deriving saliency from the vector norm, the challenge is to adapt the saliency function in order that vectors with equal color distinctiveness have equal impact on the saliency function.

2 Color Distinctiveness

The efficiency of salient point detection depends on the distinctiveness of the extracted salient points. At the salient points' positions, local neighborhoods are extracted and described by local image descriptors. The distinctiveness of the descriptor defines the conciseness of the representation and the discriminative power of the salient points. The distinctiveness is measured by its information content [15].

For luminance-based descriptors, the information content is measured by looking at the distinctiveness of the differential invariants described by the local 2-jet [8] at the detected points [14]. Montesinos et al. [12] argue that due to the extra information available in color images the color 1-jet is sufficient for local structure description. The color 1-jet descriptor is given by

$$\mathbf{v} = \begin{pmatrix} R & G & B & R_x & G_x & B_x & R_y & G_y & B_y \end{pmatrix}^T. \tag{2}$$

The information content of this color descriptor also provides an indication of the information content of more complex local color descriptors such as color differential invariant descriptors, since these complex descriptors are computed from the elements of Eq. 2.

From information theory it is known that the information content of an event is dependent on its frequency or probability. Rare events are more informative. The dependency of the information content on its probability is given by

$$I(\mathbf{v}) = -\log(p(\mathbf{v})) \tag{3}$$

where $p(\mathbf{v})$ is the probability of the descriptor $\mathbf{v}$. The information content of the descriptor, given by Eq. 2, is approximated by assuming independent probabilities of the zeroth order signal and the first order derivatives

$$p(\mathbf{v}) = p(\mathbf{f}) p(\mathbf{f}_x) p(\mathbf{f}_y) \tag{4}$$

To improve the information content of the salient point detector, defined by Eq. 1, the probability of the derivatives, $p(\mathbf{f}_x)$, should be small.

We can now restate the aim of this paper in a more precise manner. The aim is to find a transformation $g : \Re^3 \rightarrow \Re^3$ for which holds that

$$p(\mathbf{f}_x) = p\left(\mathbf{f}_x'\right) \leftrightarrow |g(\mathbf{f}_x)| = \left|g\left(\mathbf{f}_x'\right)\right|. \tag{5}$$

This implies that vectors with equal information content have equal impact on the saliency function. The transformation, attained by the function g, is called *color saliency boosting*. Similar equations hold for $p(\mathbf{f}_y)$. Once a color boosting function g has been found, the color boosted saliency can be computed with

$$s = H^\sigma (g(\mathbf{f}_x), g(\mathbf{f}_y)), \tag{6}$$

After color boosting the saliency map is based on the orientations and the information content of the derivatives. Gradient strength has been replaced by information content, thereby aiming for points with higher information content.

From Eq. 5 color boosting function g can be found by analyzing the probabilities of the derivatives. The channels of $\mathbf{f}_x$, $\{R_x, G_x, B_x\}$ are correlated due to the physics of the world. Photometric events in the world, such as shading, shadows, and reflection of the light source in specularities influence RGB values in a well defined manner [17]. Before investigating the statistics of color derivatives, the derivatives need to be transformed to a color space which is uncorrelated with respect to these photometric events.

3 Physics-Based Decorrelation

Here we describe three color coordinate transformations which partition RGB-space differently. The transformation are derived from photometric invariance theory [17]. Photometric invariance theory allows us to distinguish between various photometric causes for features in the image, such as shadows, shading, specularities and object reflectance

changes. The theory is based on the dichromatic reflection model introduced by Shafer [17]. Geusebroek et al. [3] extended the photometric reflection theory to differential-based operations. Recently, the quasi-invariant derivatives are proposed to improve noise characteristics [18]. Here we use the color transformations to decorrelate the spatial derivative, $\mathbf{f}_x$, into axes which are photometrically variant and photometrically invariant.

3.1 Spherical Color Spaces

The spherical color transformation is given by:

$$\left(\begin{array}{c} \theta \\ \varphi \\ r \end{array} \right) = \left(\begin{array}{c} \arctan(\frac{G}{R}) \\ \arcsin\left(\frac{\sqrt{R^2+G^2}}{\sqrt{R^2+G^2+B^2}} \right) \\ r = \sqrt{R^2 + G^2 + B^2} \end{array} \right). \qquad (7)$$

The spatial derivatives are transformed to the spherical coordinate system by:

$$S\left(\mathbf{f}_x \right) = \mathbf{f}_x^s = \left(\begin{array}{c} r \sin \varphi \, \theta_x \\ r\varphi_x \\ r_x \end{array} \right). \qquad (8)$$

The scale factors follow from the Jacobian of the transformation. They ensure that the norm of the derivative remains constant under transformation, hence $|\mathbf{f}_x| = |\mathbf{f}_x^s|$. In the spherical coordinate system the derivative vector is a summation of a shadow-shading variant part, $\mathbf{S}_x = (0, 0, r_x)^T$ and a shadow-shading quasi-invariant part, given by $\mathbf{S}_x^c = (r \, \sin \varphi \, \theta_x, r \, \varphi_x, 0)^T$.

3.2 Opponent Color Spaces

The opponent color space is given by:

$$\left(\begin{array}{c} o1 \\ o2 \\ o3 \end{array} \right) = \left(\begin{array}{c} \frac{R-G}{\sqrt{2}} \\ \frac{R+G-2B}{\sqrt{6}} \\ \frac{R+G+B}{\sqrt{3}} \end{array} \right). \qquad (9)$$

For this the following transformation of the derivatives follows:

$$O\left(\mathbf{f}_x \right) = \mathbf{f}_x^o = \left(\begin{array}{c} o1_x \\ o2_x \\ o3_x \end{array} \right). \qquad (10)$$

The opponent color space decorrelates the derivative with respect to specular changes. The derivative is divided into a specular variant part, $\mathbf{O}_x = (0, 0, o3_x)^T$, and a specular quasi-invariant part $\mathbf{O}_x^c = (o1_x, o2_x, 0)^T$.

3.3 Hue-Saturation-Intensity Color Spaces

The well known hue-saturation-intensity is given by

$$\left(\begin{array}{c} h \\ s \\ i \end{array} \right) = \left(\begin{array}{c} \arctan\left(\frac{o1}{o2} \right) \\ \sqrt{o1^2 + o2^2} \\ o3 \end{array} \right). \qquad (11)$$

The transformation into the hsi-space decorrelates the derivative with respect to specular, shadow and shading variations. It is given by

$$H\left(\mathbf{f}_x \right) = \mathbf{f}_x^h = \left(\begin{array}{c} s \, h_x \\ s_x \\ i_x \end{array} \right). \qquad (12)$$

The shadow-shading-specular variant is given by $\mathbf{H}_x = (0, 0, i_x)^T$ and the shadow-shading-specular quasi-invariant by $\mathbf{H}_x^c = (s \, h_x, s_x, 0)^T$.

Since the length of a vector is not changed by coordinate transformations, the norm of the derivative remains the same in all three representations $|\mathbf{f}_x| = |\mathbf{f}_x^c| = |\mathbf{f}_x^o| = |\mathbf{f}_x^h|$. For both the opponent color space and the hue-saturation-intensity color space, the photometrically variant direction is given by the $L1$ norm of the intensity. For the spherical coordinate system the photometric variant is equal to the $L2$ norm of the intensity.

In conclusion, three color spaces are discussed which decorrelate the color spaces with respect to various physical events. In the decorrelated color spaces often occurring physical variations, such as intensity changes, will only influence the photometric variant axes. In the next section the statistics of color image derivatives are examined in these physically decorrelated color spaces.

4 Statistics of Color Images

As discussed in Section 2 the information content of a descriptor depends on the probability of the derivatives, see Eq. 3 and Eq. 4. In this section we investigate the statistics of color derivatives in the decorrelated color spaces proposed in Section 3. From these statistics, we aim to find a mathematical description of surfaces of equal probability, so called isosalient surfaces. Since a description of these surfaces leads to the solution of Eq. 5.

The statistics of color images are computed for the Corel database, which consists of 40,000 images after the exclusion of black and white images. In Fig. 1 the distributions of the first order derivatives, $\mathbf{f}_x$, are given for the various color coordinate systems described in section 3 (HSI has been left out due to space considerations). The isosalient surfaces show a remarkably simple structure, approximately similar to an ellipsoid. For all three color spaces, the third coordinate coincides with the axis of maximum variation (i.e. the intensity). For the opponent and the spherical coordinate system, the first and second coordinate are rotated, with rotation matrix R^ϕ, so that the first coordinate coincides with the axis of minimum variation

$$\begin{array}{c} \left(r \sin \tilde{\varphi} \, \tilde{\theta}_x, r\tilde{\varphi}_x \right)^T = R^\phi \left(r \sin \varphi \theta_x, r\varphi_x \right)^T \\ \left(\tilde{o}1_x, \tilde{o}2_x \right)^T = R^\phi \left(o1_x, o2_x \right)^T . \end{array} \qquad (13)$$
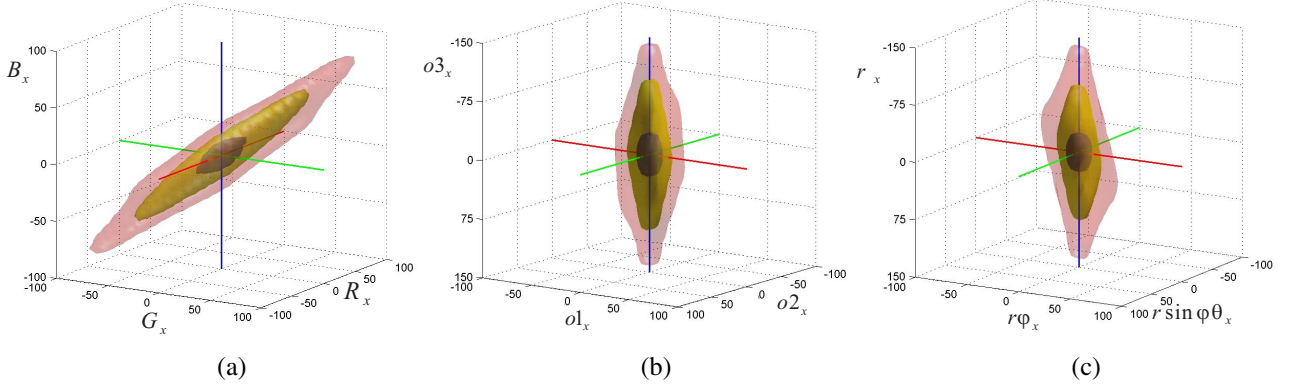
Figure 1: The histograms of the distribution of the transformed derivatives of the Corel image database in respectively the (a) RGB coordinates, (b) the opponent coordinates and (c) the spherical coordinates. The three planes correspond with the isosalient surfaces which contain (from dark to light) respectively 90%, 99%, 99.9%t of the total number of pixels.

The tilde indicates the color space transformation with the aligned axes. Similarly, the aligned transformations are given by $\tilde{S}(\mathbf{f}_x) = \mathbf{f}_x^{\tilde{s}}$ and $\tilde{O}(\mathbf{f}_x) = \mathbf{f}_x^{\tilde{o}}$.

After alignment of the axes isosalient surfaces of the derivative histograms can be approximated by ellipsoids

$$\left(\alpha h_x^1\right)^2 + \left(\beta h_x^2\right)^2 + \left(\gamma h_x^3\right)^2 = R^2 \qquad (14)$$

where $\mathbf{h}_x = h(\mathbf{f}_x) = \left(h_x^1, h_x^2, h_x^3\right)^T$ and h is one of the color transformations $\tilde{S}$, $\tilde{O}$, or H.

5 Boosting Color Saliency

We now return to our goal, that is to incorporate color distinctiveness into salient point detection. Or mathematically, to find the transformation for which vectors with equal information content have equal impact on the saliency function. In the previous section we saw that derivatives of equal saliency form an ellipsoid. Since Eq. 14 is equal to

$$\left(\alpha h_x^1\right)^2 + \left(\beta h_x^2\right)^2 + \left(\gamma h_x^3\right)^2 = |\mathbf{\Lambda} h(\mathbf{f}_x)|^2 \qquad (15)$$

the following holds

$$p(\mathbf{f}_x) = p\left(\mathbf{f}_x'\right) \leftrightarrow |\mathbf{\Lambda} h(\mathbf{f}_x)| = \left|\mathbf{\Lambda}^T h\left(\mathbf{f}_x'\right)\right|. \qquad (16)$$

where $\mathbf{\Lambda}$ is a 3x3 diagonal matrix with $\mathbf{\Lambda}_{11} = \alpha$, $\mathbf{\Lambda}_{22} = \beta$, and $\mathbf{\Lambda}_{33} = \gamma$. Λ is restricted to $\Lambda_{11}^2 + \Lambda_{22}^2 + \Lambda_{33}^2 = 1$. Eq. 16 satisfies the requirement which was imposed by Eq. 5, and hence the desired color saliency boosting function is

$$g(\mathbf{f}_x) = \mathbf{\Lambda} h(\mathbf{f}_x), \qquad (17)$$

where h is one of the color transformations $\tilde{S}$, $\tilde{O}$, or H. By a rotation of the color axes followed by a rescaling of the

axis, the oriented isosalient ellipsoids are transformed into spheres, and thus vectors of equal saliency are transformed into vectors of equal length.

5.1 Influence of Color Saliency Boosting on Repeatability

In the introduction two criteria for salient point detection were described, namely distinctiveness and repeatability. The color boosting algorithm is designed to focus on color distinctiveness, while adopting the geometrical characteristics of the operator to which it is applied. In this section we examine the influence of color boosting on the repeatability. By boosting the color saliency an anisotropic transformation is carried out. This will reduce the signal-to-noise ratio negatively. For isotropic uncorrelated noise, ε, the measured derivative $\hat{\mathbf{f}}_x$ can be written as

$$\hat{\mathbf{f}}_x = \mathbf{f}_x + \varepsilon \qquad (18)$$

and after color saliency boosting

$$g\left(\hat{\mathbf{f}}_x\right) = g(\mathbf{f}_x) + \mathbf{\Lambda}\varepsilon \qquad (19)$$

Note that isotropic noise remains unchanged under the orthogonal curvilinear transformations. Assume the worst case in which $\mathbf{f}_x$ only has signal in the photometric variant direction, then the noise can be written as

$$\frac{|g(\mathbf{f}_x)|}{|\mathbf{\Lambda}\varepsilon|} \approx \frac{\Lambda_{33}\,|\mathbf{f}_x|}{\Lambda_{11}\,|\varepsilon|}. \qquad (20)$$

Hence, the signal-to-noise ratio reduces by $\frac{\Lambda_{11}}{\Lambda_{33}}$, which will negatively influence repeatability to geometrical and photometrical changes.

	$\mathbf{f}_x$	$|\mathbf{f}_x|_1$	$\mathbf{f}_x^{\tilde{s}}$	$\tilde{\mathbf{S}}_x^c$	$\mathbf{f}_x^{\tilde{o}}$	$\tilde{\mathbf{O}}_x^c$	$\mathbf{f}_x^h$	$\mathbf{H}_x^c$
Λ_{11}	0.577	1	0.851	0.856	0.850	0.851	0.858	1
Λ_{22}	0.577	-	0.515	0.518	0.524	0.525	0.509	0
Λ_{33}	0.577	-	0.099	0	0.065	0	0.066	0

Table 1: The diagonal entries of Λ for the Corel data set computed for Gaussian derivatives with $\sigma = 1$.

6 Experiments and Illustrations

Color saliency boosting is tested on: information content and repeatability. The salient points based on color saliency boosting ($\mathbf{f}_x^{\tilde{s}}$, $\mathbf{f}_x^{\tilde{o}}$, $\mathbf{f}_x^h$) are compared to luminance ($|\mathbf{f}_x|_1$), RGB gradient ($\mathbf{f}_x$), and the quasi-invariant-based salient point detectors ($\tilde{\mathbf{S}}_x^c$, $\tilde{\mathbf{O}}_x^c$, $\mathbf{H}_x^c$). The generality of the approach is illustrated by applying color boosting to several existing feature detectors.

6.1 Initialization

Experiments are performed on a subset of 1000 randomly chosen images from the Corel data set. Before color saliency boosting can be applied, the Λ-parameters (Eq.15) have to be initialized by fitting ellipses to the histogram of the data set. First the axes of the opponent and the spherical transformation are aligned by Eq. 13. Next, the axes of the ellipsoid are derived by fitting the isosaliency surface which contains 99 percent of the pixels of the histogram of the Corel data set. Changing this parameter to 99.9 or 99.99 percent changes matrix Λ only slightly. The results for the various transformations are summarized in Table 1. The relation between the axes in the various color spaces clearly confirms the dominance of the luminance axis in the RGB-cube, since Λ_{33}, the multiplication-factor of the luminance axis, is much smaller than the color-axes multiplication factors, Λ_{11} and Λ_{22}.

We have chosen the Harris point detector [4] to test color boosting in experiment 6.2 and 6.3. It is computed with

$$H^\sigma(\mathbf{f}_x, \mathbf{f}_y) = \overline{\mathbf{f}_x \cdot \mathbf{f}_x\, \mathbf{f}_y \cdot \mathbf{f}_y} - \overline{\mathbf{f}_x \cdot \mathbf{f}_y}^2 - k\left(\overline{\mathbf{f}_x \cdot \mathbf{f}_x} + \overline{\mathbf{f}_y \cdot \mathbf{f}_y}\right)^2 \quad (21)$$

by substituting $\mathbf{f}_x$ and $\mathbf{f}_y$ by $g(\mathbf{f}_x)$ and $g(\mathbf{f}_y)$. The bar $\overline{\cdot}$ indicates convolution with a Gaussian filter and the dot indicates the inner product. We applied Gaussian derivatives of $\sigma = 1$ and computed the Harris detector response with a Gaussian of $\sigma = 3$.

6.2 Color Distinctiveness

Here we examine if color boosting improves the color distinctiveness of the Harris detector. In [15], the Harris detector has already been shown to outperform other detectors both on 'shape' distinctiveness and repeatability. The color

method	20 points			100 points		
	inf.	incr(%)	decr(%)	inf.	incr.	decr.
$\mathbf{f}_x$	20.4	-	-	20.0	-	-
$|\mathbf{f}_x|_1$	19.9	0	1.4	19.8	0	0.8
$\tilde{\mathbf{S}}_x^c$	22.2	45.5	10.1	20.4	9.1	17.7
$\mathbf{f}_x^{\tilde{s}}$	22.3	49.4	.6	20.8	13.1	1.3
$\tilde{\mathbf{O}}_x^c$	22.6	51.4	12.9	20.5	12.0	34.2
$\mathbf{f}_x^{\tilde{o}}$	23.2	62.6	0.0	21.4	21.5	0.9
$\mathbf{H}_x^c$	21.0	21.7	43.4	19.0	1.8	77.4
$\mathbf{f}_x^h$	23.0	57.2	0.3	21.3	16.7	1.1
rand.	14.4	0	99.8	14.4	0	100

Table 2: The information content of salient point detectors. Measured in 1. information content and 2. the percentage of images for which a substantial decrease (-5%) or increase ($+5\%$) of the information content occurs. The experiment is performed with both 20 or 100 salient points per image.

distinctiveness of salient point detectors is described by the information content of the descriptors extracted at the locations of the salient points. From the combination of Eq. 3 and Eq. 4, it follows that the total information is computed by summing the information of the zeroth and first order part,

$$I(\mathbf{v}) = I(\mathbf{f}) + I(\mathbf{f}_x) + I(\mathbf{f}_y). \quad (22)$$

The information content of the parts is computed from the normalized histograms with

$$I(\mathbf{f}) = -\sum_i p_i \log(p_i) \quad (23)$$

where p_i are the probabilities of the bins of the histogram of $\mathbf{f}$.

The results for 20 and 100 salient points per image are shown in Table 2. Next to the absolute information content we have also computed the relative information gain with respect to the information content of the color gradient based Harris detector. For this purpose the information content on a single image is defined as

$$I = -\sum_{j=1}^n \log(p(v_j)) \quad (24)$$

where $j = 1, 2, ...n$ and n is the number of salient points in the image. Here $p(v_j)$ is computed from the global histograms, which allows comparison of the results per image. The information content change is considered substantially for a 5 percent increase or decrease.

The highest information content is obtained with $\mathbf{f}_x^{\tilde{o}}$, which is the color saliency boosted version of the opponent
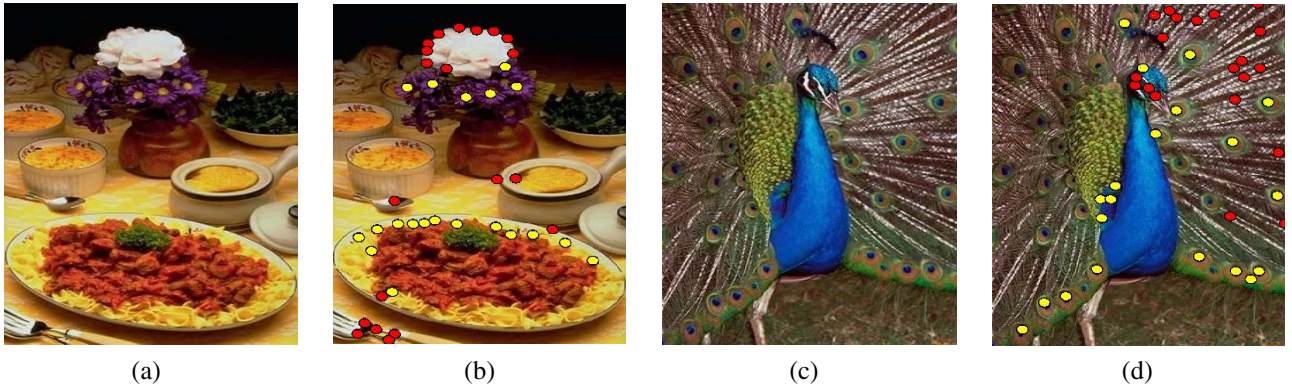
(a)   (b)   (c)   (d)

Figure 2: (a) and (c) Corel input images. (b) and (d) results of Harris detector (red dots) and the Harris detector with color boosting (yellow dots). The red dots mainly coincide with black and white events, while the yellow dots are focussed on colorful points.

derivatives. The boosting results in an increase of 7% (20. to 21.4 in the case of 100 points) to 14% (20.4 to 23.2 in the case of 20 points) of the information content compared to the color gradient based detector. On the images of the Corel set this resulted in a substantial increase on 22% to 63% of the images. The advantage of color boosting diminishes when increasing the number of salient points per image. This is caused by the limited number of color clues in many of the images, which is especially visible for the results of the photometric quasi-invariants, $\tilde{\mathbf{S}}_x^c$, $\tilde{\mathbf{O}}_x^c$, or $\mathbf{H}_x^c$. These detectors discard all intensity information, which in the case of 100 salient points per image results in many images with a substantial decrease in information content. Finally, it is noteworthy to see how small the difference is between luminance and RGB-based Harris detection. Since the intensity direction also dominates the RGB derivatives, using RGB-gradient instead of luminance-based Harris detection only results in a substantial increase in information content in 1% of the images.

In Fig. 2 results of the RGB-gradient based and color boosted Harris detector are depicted. From a color information point of view, the RGB-gradient based method does a poor job. Most of the salient points have a black and white local neighborhood, with a low color saliency. The salient points after color boosting focus on more distinctive points, such as the colorful flowers and the plumage of the peacock.

6.3 Repeatability

Repeatability measures the stability with respect to varying viewing conditions. As indicated in section 5.1 color saliency boosting reduces the signal-to-noise ratio. Repeatability with respect to geometrical changes, scaling, and affine transformations are considered a property of the detector and will not be considered here.

The loss of repeatability caused by color saliency boost-

method	20 points	100 points
$\mathbf{f}_x$	88	84
$|\mathbf{f}_x|_1$	88	83
$\tilde{\mathbf{S}}_x^c$	53	42
$\mathbf{f}_x^{\tilde{s}}$	62	54
$\tilde{\mathbf{O}}_x^c$	46	34
$\mathbf{f}_x^{\tilde{o}}$	51	41
$\mathbf{H}_x^c$	35	25
$\mathbf{f}_x^h$	52	42

Table 3: The percentage of Harris points which remain detected after adding Gaussian uncorrelated noise.

ing is examined by adding uniform, uncorrelated Gaussian noise of $\sigma = 10$. This yields a good indication of loss in signal-to-noise, which in its turn will influence results of repeatability under other variations, such as zooming, illumination changes, and geometrical changes. Repeatability is measured by comparing the Harris points detected in the noisy image to the points in the noise-free images. The results in Fig. 3a correspond to the expectation made by Eq. 20, namely the larger the difference between Λ_{11} and Λ_{33}, the poorer the repeatability.

6.4 Illustrations Generality

Color saliency boosting can be applied on all function which can be written as a function of the local derivatives. Here we apply it to two different feature detectors. First the focus point detector which was originally proposed by Reisfeld et al. [13] and recently extended to color by Heidemann [5]. The detector focuses on the center of locally symmetric
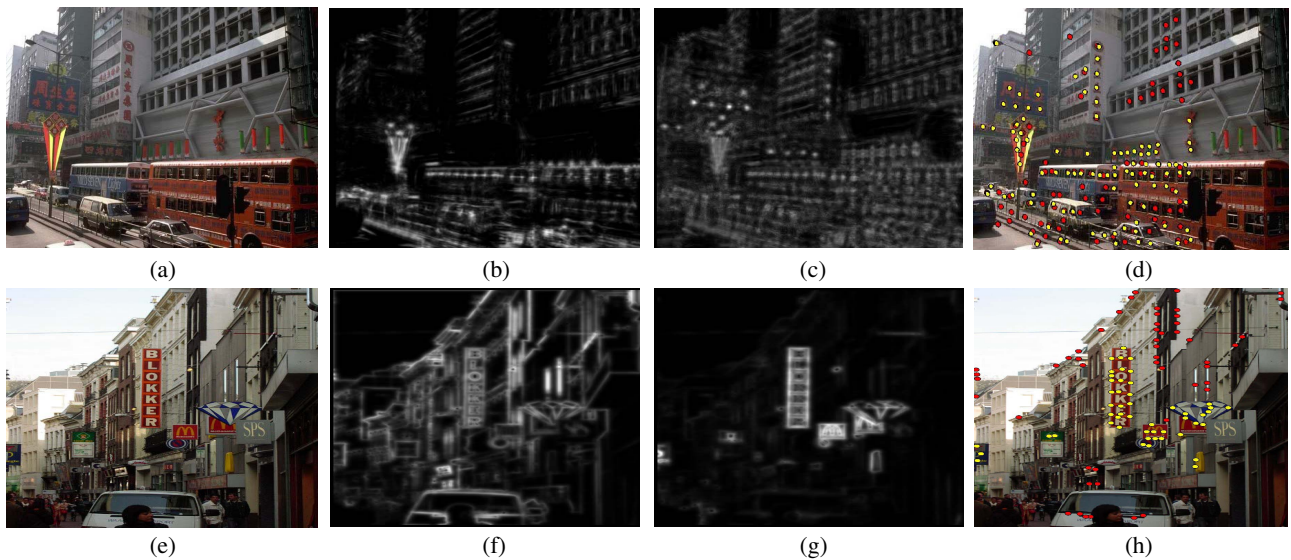
Figure 3: Respectively, the input image, RGB-gradient based saliency map, the color boosted saliency map and the results with red dots (lines) for the gradient-based method and yellow dots for the salient points after color saliency boosting. (a),(b),(c),(d) Results for the focus points, (e),(f),(g),(h) for the symmetry points.

structures. On the first row of Fig. 3 the results of the focus point detector are shown. Fig. 3b shows the saliency map as proposed in [5]. In Fig. 3c the result after saliency boosting is depicted. Although focus point detection is already an extension from luminance to color, black-and-white transition still dominate the result. Only after boosting the color saliency, the less interesting black-and-white structures in the image are ignored and most of the red Chinese signs are found. Similar difference in performance is obtained by applying color boosting to the linear symmetry detector proposed by Bigün [1]. This detector focuses on corner and junction like structures. The RGB gradient based method focuses mainly on black-and-white events while the more salient signboards are found only after color saliency boosting.

7 Conclusions

In this paper color distinctiveness is explicitly incorporated in the design of salient point detectors. The method, called color saliency boosting, can be incorporated into existing detectors which are mostly focused on shape distinctiveness. Saliency boosting is based upon an analysis of the statistics of color image derivatives. Isosalient derivatives form ellipsoids in the color derivative histograms. This fact is exploited to adapt derivatives in such a way that equal saliency implies equal impact on the saliency map. Experiments show that color saliency boosting substantially increases the information content of the detected points. A

substantial information content increase is obtained on up to 20% to 60% of the Corel images. Further, the generality of the method is illustrated by applying color boosting to various existing salient point detectors.

References

[1] J. Bigun. Pattern recognition in images by symmetry and co-ordinate transformations. *Computer Vision and Image Understanding*, 68(3):290–307, 1997.

[2] R. Fergus, P. Perona, and A. Zisserman. Object class recognition by unsupervised scale-invariant learning. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, volume 2, pages 264–271, June 2003.

[3] J.M. Geusebroek, R. van den Boomgaard, A.W.M. Smeulders, and H. Geerts. Color invariance. *IEEE Trans. Pattern Analysis Machine Intell.*, 23(12):1338–1350, 2001.

[4] C. Harris and M. Stephens. A combined corner and edge detector. In *Proc. 4th Alvey Vision Conf.*, volume 15, pages 147–151, 1988.

[5] G. Heidemann. Focus-of-attention from local color symmetries. *IEEE trans. on pattern analysis and machine intelligence*, 26(7):817–847, 2004.

[6] L. Itti, C. Koch, and E. Niebur. A model of saliency-based visual attention for rapid scene analysis. *IEEE Transactions on Pattern Analysis and Machine Intelligence*, 20(11):1254–1259, Nov 1998.

[7] L. Itti, C. Koch, and E. Niebur. Computation modeling of visual attention. *Nature Reviews Neuroscience*, 2(11):194–203, March 2001.

[8] J J Koenderink and A J van Doom. Representation of local geometry in the visual system. *Biol. Cybern.*, 55(6):367–375, 1987.

[9] M.S. Lee and G. Medioni. Grouping into regions, curves, and junctions. *Computer Vision Image Understanding*, 76(1):54–69, 1999.

[10] D.G. Lowe. Distinctive image features from scale-invariant keypoints. *Int. J. Computer Vision*, 60(2):91–110, 2004.

[11] K. Mikolajczyk and Cordelia Schmid. Scale and affine invariant interest point detectors. *Int. J. Computer Vision*, 60(1):62–86, 2004.

[12] P. Montesinos, V. Gouet, and R. Deriche. Differential Invariants for Color Images. In *Proceedings of 14^{th} International Conference on Pattern Recognition*, pages 838–840, Brisbane, Australia, 1998.

[13] D. Reisfeld, H. Wolfson, and Y. Yeshurun. Context free attentional operators: the generalized symmetry. *International Journal of Computer Vision*, 14:119–130, 1995.

[14] C. Schmid and R. Mohr. Local grayvalue invariants for image retrieval. *IEEE Transactions on Pattern Analysis and Machine Intelligence*, 19(5):530–534, May 1997.

[15] C. Schmid, R. Mohr, and C. Bauckhage. Evaluation of interest point detectors. *International Journal of Computer Vision*, 37(2):151–172, 2000.

[16] 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):1087–1095, 2003.

[17] S.A. Shafer. Using color to seperate reflection components. *COLOR research and application*, 10(4):210–218, Winter 1985.

[18] J. van de Weijer, Th. Gevers, and J.M. Geusebroek. Edge and corner detection by photometric quasi-invariants. *IEEE Trans. Pattern Analysis and Machine Intelligence*, 27(4):625–630, 2005.

[19] L. van Gool, T. Tuytelaars, and A. Turina. Local features for image retrieval. In *State-of-the-Art in Content-Based Image and Video Retrieval*, pages 21–41. Kluwer Academic Publishers, 2001.

[20] L.R. Williams and K.K. Thornber. A comparision of measures for detecting natural shapes in cluttered backgrounds. *International Journal of Computer Vision*, 34(2/3):81–96, 1999.