

Increasing Illumination Invariance of SURF Feature Detector Through Color Constancy

Marcelo R. Petry^{1,2}, Antonio Paulo Moreira¹, and Luis Paulo Reis^{2,3}

¹ INESC TEC and Faculty of Engineering, University of Porto,
Rua Dr. Roberto Frias, s/n, 4200-465, Porto, Portugal
`{marcelo.petry, amoreira}@fe.up.pt`

² LIACC, University of Porto

³ DSI/School of Engineering, University of Minho,
Campus de Azurem, 4800-058, Guimares, Portugal
`lpreis@dsi.uminho.pt`

Abstract. Most of the original image feature detectors are not able to cope with large photometric variations, and their extensions that should improve detection eventually increase the computational cost and introduce more noise to the system. Here we extend the original SURF algorithm increasing its invariance to illumination changes. Our approach uses the local space average color descriptor as working space to detect invariant features. A theoretical analysis demonstrates the impact of distinct photometric variations on the response of blob-like features detected with the SURF algorithm. Experimental results demonstrate the effectiveness of the approach in several illumination conditions including the presence of two or more distinct light sources, variations in color, in offset and scale.

Keywords: Feature detection, SURF, Color Invariance

1 Introduction

Many computer vision tasks depend heavily on extraction and matching of corresponding points (features) over consecutive images. Applications of feature based algorithms include, but are not limited to, image classification, image segmentation, object recognition and camera calibration. In robotics, motion estimation methodologies like visual odometry and visual SLAM have been able to complement traditional navigation sensors (like global navigation satellite systems, encoders and inertial measurement units), offering smooth (30Hz) and locally highly accurate localization information.

In general, this problem is tackled by searching for image regions whose low-level characteristics (i.e size, shape, luminance, color, texture, binocular disparity) significantly differs from the background. As important as the feature distinctiveness, is its ability to be repeatedly identified in consecutive images. However, image RGB values are significantly influenced by variations in scene illuminant. Such variations introduce undesirable effects and negatively affect the

performance of computer vision methods. For this reason, one of the most fundamental tasks of visual systems is to distinguish the changes due to underlying imaged surfaces from those changes due to the effects of the scene illumination. In order to increase the probability of image features to be re-detected in subsequent images, it is extremely important for them to be robust to noise and invariant with regard to geometric (changes in scale, translation, rotation, affine/projective transformation) and photometric variations (illumination direction, intensity, color and highlights).

The Speeded Up Robust Feature (SURF) [1] is a widely used feature detector in robotics motion estimation due to its low computation time. The algorithm provides features that are invariant to image scale and rotation, but only partially invariant to changes in viewpoint and illumination. In this paper the problem of SURF illumination invariance is addressed.

The remainder of the paper is organized as follows. The second section reports the related works under this topic. The third section briefly presents the theory regarding image formation and color invariance. In the fourth section we perform a mathematical analysis over the SURF feature detection function to identify its weakness regarding photometric variations. In the fifth section we present an approach to improve SURF illumination invariance by exploiting the local space average color descriptor as working space for feature detection. Next it is presented the experimental work and results comparing the repeatability rate of the proposed approach with the original SURF implementation. Finally some conclusions and directions for future work conclude the paper.

2 Related Work

Originally, most of the feature detectors and descriptors were designed to cope only with the image luminance. Later, in order to take advantage of illumination invariance properties of other color spaces, some researchers proposed extensions for the original algorithms. In [2], Ancuti and Bekaert proposed an extension to the SIFT descriptor (SIFT-CCH) that combines the SIFT approach with the color co-occurrence histograms (CCH) computed from the Nrgb color space. Their algorithm performs the same as SIFT in the detection step, but introduces one dimension to the descriptor. Thus, features are described by a two element vector that combines the SIFT and the CCH descriptor vectors. The main problem of such an approach is the increase in the computational effort during the feature matching due to the extra 128 elements added to the descriptor vector. The color-SURF proposed by Fan et al.[3] was maybe the first to approach suggesting the use of colors in SURF descriptors. Through a methodology similar to the SIFT-CCH, the authors propose the addition of a new dimension to the descriptor vector. This extra information corresponds to the color histogram computed from the YUV color space, and adds a 64-element vector for each feature descriptor. For this reason, just like in the SIFT-CCH, the extra elements in the descriptor vector increase the computational effort necessary during the matching step.

In [4], Abdel-Hakim and Farag uses the invariant property H (related to hue) of the Gaussian color model as working space. Thus, instead of using gray gradients to track SIFT features, they use the gradients of the color invariant to detect and describe features. Although the authors used the H invariant instead of the C invariant, the approach is called CSIFT in a reference to the introduction of color in the SIFT operator. [5], also use invariants derived from the Gaussian color model to reduce the photometric effects in SIFT descriptions. They compare the individual performance of four invariants with the original SIFT approach, with the CSIFT approach [4] and with the HSV-SIFT approach [6]. Their evaluation suggests that the C-invariant, which can be intuitively seen as the normalized opponent color space, outperforms the original SIFT description and all the other approaches. In reference to the results of the C-invariant, the combination of this invariant with the SIFT operator is called C-SIFT. Sande et al. [7] presents an evaluation of the different approaches that attempt to provide photometric invariance to SIFT like descriptors.

3 Image Theory

The geometric distribution of the body reflection is sometimes assumed to reflect light evenly in all directions. Therefore, the luminance in such isotropic surfaces, also known as Lambertian surfaces, is the same regardless of the viewing angle. Assuming that a scene contain surfaces which exhibits Lambertian reflectance properties, its resulting image I can be modelled in terms of the surface reflectance $S(\lambda, x_{obj})$ and the light spectral power distribution $E(\lambda, x_{obj})$ falling onto an infinitesimal small patch on the sensor array.

$$I(x_i) = \int E(\lambda, x_{obj})S(\lambda, x_{obj})p(\lambda)d\lambda . \quad (1)$$

Where $p(\lambda)$ is the camera spectral sensitivity of wavelength λ , x_{obj} is the object location in the world coordinate frame and x_i is its location in the image coordinate frame. Although each sensor responds to a range of wavelengths, the sensor is often assumed to respond to the light of a single wavelength. Thus, one can approximate the sensor response characteristics by Dirac's delta functions. Through the former assumption, it is possible to simplify the Equation (1) and express the intensity $I_k(x_i)$ measured by the sensor $k \in \{R, G, B\}$ in the position x_i as:

$$I_k(x_i) = E_k(x_{obj})S_k(x_{obj}) . \quad (2)$$

3.1 Diagonal Model

One of the most difficult problems when working with colors is that the object's apparent color varies unpredictably with variations in the intensity and temperature of the light source. A well-known example occur in outdoor environments

with daylight variations, the color shift between sunny and cloudy days is simply not well modeled as Gaussian noise in RGB [8].

One of the most used models to describe those kind of variations is the von-Kries model, or Diagonal Model (DM), which corresponds to a diagonal transformation of the color space. According to Diagonal model, it is possible to map an observed image I_o taken under an unknown illuminant to a corresponding image I_c under a canonical illuminant through a proper transformation in order to render images color constant. Finlayson et al. [9] note that the DM model present shortcomings when mapping near saturated colors, and propose an extension that includes the "diffuse" light term by adding an offset. Such model is known as the Diagonal-Offset Model, and is given by:

$$\begin{bmatrix} R_c \\ G_c \\ B_c \end{bmatrix} = \begin{bmatrix} a & 0 & 0 \\ 0 & b & 0 \\ 0 & 0 & c \end{bmatrix} \begin{bmatrix} R_o \\ G_o \\ B_o \end{bmatrix} + \begin{bmatrix} o_1 \\ o_2 \\ o_3 \end{bmatrix}. \quad (3)$$

Using the Diagonal-offset model, illumination variations can be classified according to the values of the scalar and offset into five distinct categories [7]. In the light intensity change (LIC), the three RGB components of a given image varies equally by a constant factor, such that $a = b = c$ and $o_1 = o_2 = o_3 = 0$. Hence, when a function is invariant to light intensity changes, it is scale-invariant with respect to light intensity. In the light intensity shift (LIS), a constant shift affects equally all the RGB channels of a given image, such that $a = b = c = 1$ and $o_1 = o_2 = o_3 \neq 0$. The light intensity change and shift (LICS) is a combination of the two above mentioned categories, and also affect all three RGB channels equally, in such a way that $a = b = c$ and $o_1 = o_2 = o_3 \neq 0$. Thus, when a function is invariant to light intensity changes and to light intensity shift, it is known as scale-invariant and shift-invariant with respect to light intensity. The two remaining categories do not assume that RGB channels are equally affected by variations in the light source. The light color change (LCC) corresponds to the Diagonal model, and assumes that $a \neq b \neq c$ and $o_1 = o_2 = o_3 = 0$. Since images are able to vary differently in each channel, this category can model changes in the illuminant color temperature and light scattering. The last, light color change and shift (LCCS), corresponds to the full Diagonal-offset model and takes into consideration independent scales $a \neq b \neq c$ and offsets $o_1 \neq o_2 \neq o_3$ for each image channel.

3.2 Color Constancy

The ability to perceive color as constant under changing conditions of illumination is known as color constancy, and is a natural ability of human observers. The problem of computing a color constant descriptor based only on data measured by the retinal receptors is actually underdetermined, as both $E(\lambda, X_{obj})$ and $p_k(\lambda)$ are unknown. Therefore, one need to impose some assumptions regarding the imaging conditions. The most simple and general approaches to color constancy (i.e. White Patch [10] and the Gray World [11]) make use of a single

statistic of the scene to estimate the illuminant, which is assumed to be uniform in the region of interest. Approaches like Gammut Mapping, on the other hand, make use of assumptions of the surface reflectance properties of the objects.

A more recent method is based on the Local Space Average color (LSAC), which can be defined as a computational model of how the human visual system performs averaging of image pixels [12]. The model proposed by Ebner makes two important assumptions. The first is that the essential processing required to compute a color constant descriptor in human observers is located in the visual area V4 of the extrastriate visual cortex [13]. The second is that gap junctions behave like resistors. Thus, Ebner models the gap junctions between neurons in V4 as a resistive grid, which can be used to compute Local Space Average color, and then color constant descriptors. Each neuron of this resistive grid computes the local space average color by iterating update equations indefinitely for all three bands. According to Ebner, the iterative computation of Local Space Average Color produces results which are similar to the convolution of the input image with a Gaussian kernel.

The Local Space Average Color alone is just a biologically inspired theory that tries to explain how the brain averages image pixels. However, when combined with the Gray World hypothesis, LSAC can provide means to derive color invariant descriptors. The advantage of Ebner's work is that if we consider the Gray World assumption in a local perspective, it is possible to estimate the color of the illuminant at each image pixel. For a more detailed theoretical formulation we may refer to [12]. Given the local space average color a_k , one can derive a local color invariant descriptor O_k through:

$$O_k(x, y) = \frac{I_k(x, y)}{2a_k(x, y)} \approx \frac{S_k(x, y)E_k(x, y)}{E_k(x, y)} \approx S_k(x, y) . \quad (4)$$

4 Analysis of Photometric Variations in SURF

In order to understand the effects of the light source variation in SURF responses consider an observed single channel image I_o with pixel intensity $I_o(x, y)$ at a given point $X = (x, y)$. Through the central difference method it is possible to express the second derivatives of $I_o(x, y)$ as:

$$\frac{\partial^2 I_o(x, y)}{\partial x^2} = I_o(x + 1, y) - 2I_o(x, y) + I_o(x - 1, y) . \quad (5)$$

Now, consider that I_o has a corresponding image I_u , taken under unknown illuminant. Assuming the Diagonal-offset model these two images are related by a linear transformation determined by a scalar constant α and an offset β . Therefore, the pixel intensity $I_u(x, y)$ of the image I_u at the same point $X = (x, y)$ can be modeled as:

$$I_u(x, y) = \alpha I_o(x, y) + \beta . \quad (6)$$

Thus, it is possible to conclude that the second derivatives of $I_u(x, y)$ with respect to x is:

$$\frac{\partial^2 I_u(x, y)}{\partial x^2} = \alpha \frac{\partial^2 I_o(x, y)}{\partial x^2} . \quad (7)$$

The same applies to the second derivatives with respect to y and XY . When computing the derivatives, the diffuse term β is canceled out, causing no impact on the final result. However, by varying the illumination with a scalar α the second derivatives vary proportionally with the scalar. Feature localization is a three-step process that starts disregarding points in which blob-response is lower than a fixed threshold value. Thus, if the detector responses vary with the illumination, a given feature that is detected in a bright image may not be detected in a corresponding image with lower illumination levels. SURF detector response $R_u(x, y)$ of a given pixel $I_u(x, y)$ is given by the determinant of the Hessian matrix:

$$R_u(x, y) = \frac{\partial^2 I_u(x, y)}{\partial x^2} \frac{\partial^2 I_u(x, y)}{\partial y^2} - \left(\frac{\partial^2 I_u(x, y)}{\partial xy} \right)^2 . \quad (8)$$

Replacing (7) into (8), the filter response R_u can be expressed in terms of the SURF response $R_o(x, y)$ of the I_o :

$$R_u(x, y) = \alpha \frac{\partial^2 I_o(x, y)}{\partial x^2} \alpha \frac{\partial^2 I_o(x, y)}{\partial y^2} - \left(\alpha \frac{\partial^2 I_o(x, y)}{\partial xy} \right)^2 = \alpha^2 R_o . \quad (9)$$

The degree of the scalar (α^2) in (9) provides the theoretical explanation to why even small variations in the scene illuminant cause significant variations in the magnitude of the detector response.

5 The Proposed LSAC SURF Approach

Among color constancy methods, gamut mapping is referred in literature as one of the most successful algorithms [14]. It has demonstrated good results in different datasets of several works. The method is though computationally quite complex. Its implementation requires the computation two convex hulls, which is a difficult problem when using finite precision arithmetic. Another drawback is that the algorithm requires an image data set with known light sources to estimate the canonical gamut (learning phase) that will be used to compute the transformation matrix, and thus estimate the illuminant (testing phase). In practice, such methodology is not viable for robotic vision systems since robots are not constrained to one specific scenario, but subjected to multiple and dynamic environments.

Low level color constant algorithms, on the other hand, are less complex, faster and only slightly outperformed by the gamut mapping [15]. These characteristics make them perfect candidates for improving robotic vision systems.

One limitation of the Gray World assumption is that it is only valid in images with sufficient amount of color variations. Only when the variations in color are random and independent, the average value of the R, G, and B components of the image would converge to a common gray value. This assumption is, however, held very well in several real world scenarios, where it is usually true that there are a lot of different color variations.

Another limitation of most color constancy algorithms is that they are modeled with the assumption that the scene is uniformly illuminated. Since in practice multiple illuminants are present in the scene, the illumination is not uniform, and thus the premise is not fully verified. For instance, some daylight may be falling through a window while an artificial illuminant may be switched on inside the room. In fact, that may be the main advantage of the descriptors derived from the Local Space Average color methodology. Since LSAC estimates the illuminant locally for each point of the scene, its descriptors are better prepared to deliver color constancy in real world images.

Most color invariant feature detectors proposed combines the original detector with some sort of color space mapping. Our approach to achieve photometric invariant feature responses (LSAC SURF) consists on taking advantage of the invariant properties of the LSAC descriptor, using it as working space to perform SURF feature extraction. The inclusion of this pre-processing step adds a small computational load, but may provide a significant increase in feature detection robustness.

The size of the window that LSAC is computed plays a important role in the robustness of the feature detection. Empirical observation demonstrated that feature repeatability tends to perform better when LSAC is computed over small neighborhoods. In fact, due to the multiple illumination sources the values of α and β tends to vary significantly in distant image pixels, which makes the assumption that $E_k(x, y) \approx 2a_k(x, y)$ to be valid only for a small regions.

When a pixel reaches saturation, it does not present the same variation as its neighbors, causing non linear variations in the response of the feature detector and decreasing the probability to be correctly matched in subsequent images. Therefore, features which pixel intensities are close to saturation are not good candidates for matching. However, such features can not simply be ignored since under certain illumination variations their pixel intensity can move away from saturation, and make them good candidates for matching in subsequent images. For this reason, each detected feature is classified into hard and soft features according to their pixel intensities. If the pixel intensity of a distinct image region is lower than an upper threshold and higher than a lower threshold the feature is classified as hard feature, on the contrary, the feature is classified as soft feature. The choice of the proper upper and lower threshold values might be determined according to the expected variation in the scene illumination.

Since hard features are more likely to be found in subsequent images, we can reduce the search space and match only the current hard features with the subsequent set of features. In this context, soft features are used only to

support matching of previous hard features, while hard features are used in the computation of sensitive visual tasks.

6 Results

This section presents the experimental validation of the proposed method. First, the influence of several illumination conditions on the performance of the proposed method is studied using controlled indoor images, Fig. 1. Next, the proposed method is applied on a data set of real images, Fig. 3. In the following experiments the optimum parameters of SURF [1] were assigned to both algorithms. The LSAC was approximated with a Gaussian kernel of size 15×15 , with standard deviation $\sigma = 4$. The lower and upper thresholds used for feature classification into hard and soft were set to 20 and 220 respectively.

6.1 Performance Measurement

To evaluate the performance of our approach we adopted the repeatability criterion similar to the proposed by Schmid et al. [16]. The repeatability rate evaluates the ratio between the number of point-to-point correspondences that can be established for detected points in all images of the same scene $C(I_1, \dots, I_n)$ and the total number of features detected in the current image m_i . Therefore, the higher the repeatability, the more likely features are to be matched and the better the matching results tend to be.

$$R_i = \frac{C(I_1, \dots, I_n)}{m_i} . \quad (10)$$

Where i denotes the image under analysis and n the number of images of the same scene. The repeatability rate of our approach was compared with the repeatability rate of the SURF algorithm implemented in the OpenCV library.

6.2 Controlled Image Set

First experiments are performed on images available in the Amsterdam Library of Object Images (ALOI) [17]. ALOI provides several image collections, like the light color change (ALCC) collection. ALCC is a collection of images in which the color of the illumination source was varied from yellow to white, according to the voltage v_0 of the lamps (where $v_0 = 12i/255$ volts and $i \in \{110, 120, 130, 140, 150, 160, 170, 180, 190, 210, 230\}$). Since among the ALOI collections only the ALCC collection have a direct correspondence to the illumination variations modeled through the Diagonal offset model, three new controlled collections were artificially created: LIC, LIS, and LCC collections.

To create the controlled collections we selected all images from the ALCC collection with color temperature of 2750k, and performed the proper transformations. Thus, all collections contain a set of 9.000 images of 1.000 objects designed

to evaluate the effects of specific variations in the scene illuminant. The LIS collection was created by shifting all the color channels equally by an offset $\beta \in \{-20, -15, -10, -5, 0, 5, 10, 15, 20\}$. The LIC collection was created scaling all the RGB channels equally by a factor $\alpha \in \{1/2.0, 1/1.5, 1/1.2, 1/1.1, 1, 1.1, 1.2, 1.5, 2.0\}$. Finally, the LCC collection, which mimics the effect of a light variation from bluish to white, was created by scaling both the Red and Green channels by a factor $\alpha \in \{0.2, 0.3, 0.4, 0.5, 0.6, 0.7, 0.8, 0.9, 1.0\}$, while keeping $\alpha \in \{1.0\}$ for the Blue channel.

The repeatability results in the LIS collection, Fig. 2a, confirm the theoretical analysis and demonstrate that offset variations indeed do not affect the performance of the detection algorithms. The results demonstrate a good performance for both algorithms, in which mean repeatability remained above 90%.

Variations of scalar order, on the other hand, greatly impact SURF repeatability performance. Fig. 2b demonstrates the low repeatability rate of SURF algorithm in the LIC collection. LSAC SURF demonstrated a much higher and constant mean repeatability rate, presenting a significant improvement in the mean repeatability for all values of α . Note that a higher mean repeatability rate occurs for the smaller values of α , in both SURF and LSAC SURF, due to the tendency to find a smaller number of features in darker images.

Results of the mean repeatability rate of the ALCC collection, Fig. 2c, demonstrate that the mean repeatability rate was not significantly affected by the illumination color variation, presenting a mean repeatability rate above 90% for both algorithms. This result can be justified by the weight of the color components in the grayscale conversion. When varying the light source from yellow to white, only the blue component of the RGB model varies. Since the weight of the blue component (0.114) is considerably lower than the weight of the red (0.299) and green components (0.587), the variation in this color channel does not cause a sufficient large photometric variation to impact the grayscale image used in the feature detection. However, when varying the color of the light source from bluish to white (LCC collection) the mean repeatability rate of LSAC SURF significantly outperformed SURF. Fig. 2d demonstrates once again the the low




Fig. 1: Samples of the controlled indoor image set. From top row to bottom: samples of one object in the LIS, LIC, ALCC and LCC collections

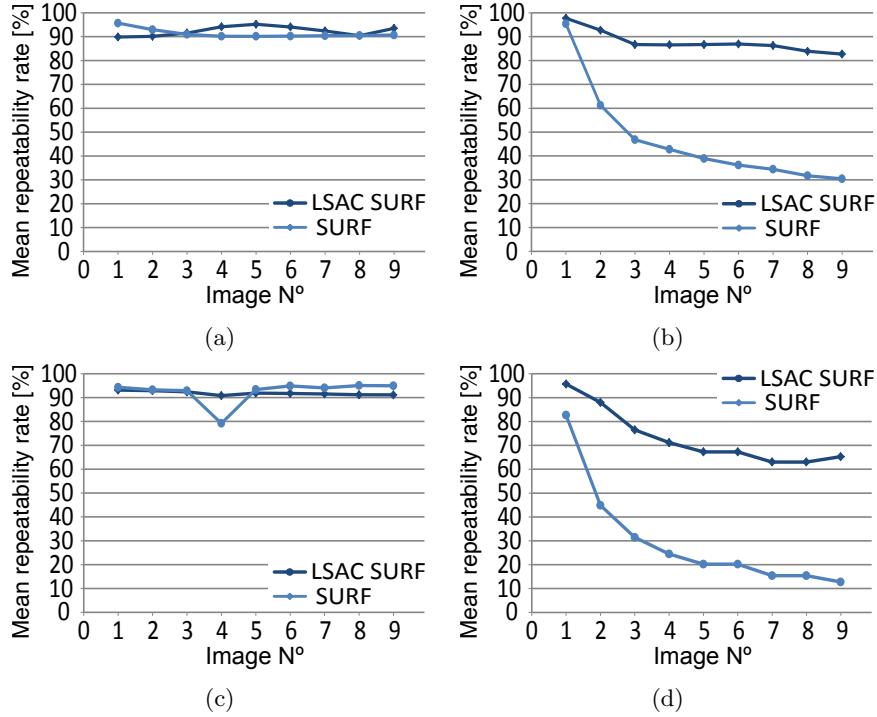


Fig. 2: Mean repeatability of (a)LIS, (b)LIC, (c)ALCC and (d)LCC collections.

repeatability rate of SURF algorithm and its tendency to decrease with higher values of α , while LSAC SURF presented a much higher and constant mean repeatability rate.

6.3 Real World Image Set

Here, the proposed method is tested on a data set of real images. This data set consists of twelve indoor images of a robotic soccer field. Images were taken with the camera mounted in a fixed position, while the scene illumination was varied through several combinations of individually regulated ceiling lightings. This dataset offers a challenging environment for robust feature detection since it contains non-uniform illumination due to multiple sources (different bulb lamps



Fig. 3: Dataset of real images: robotic soccer field




Fig. 4: Repeatability of the Robotic Soccer Field collection

in the ceiling and natural illumination from the windows), as well as variations in the color of the illuminant, shading, shadows, specularities, and interreflections.

The theoretical improvement offered by the LSAC SURF in the robustness of feature detection can be experimentally verified through the Fig. 4. This result demonstrates the low repeatability rate of SURF algorithm in the Robotic Soccer Field collection, which remained around 29% for all images. LSAC SURF, on the other hand, demonstrated a much higher repeatability rate, not lower than 65%, presenting a significant improvement in the repeatability score in all illumination conditions.

7 Conclusion

In this paper, the LSAC SURF approach was introduced as an extension of SURF algorithm. The methodology proposed has shown to be able to improve feature detection repeatability rates in scenarios where the uniform light-source assumption is too restrictive. The theoretical analysis demonstrated which variations in the illuminant affects images derivatives and SURF responses. We demonstrated that SURF response is proportional to the square of the scalar variation of the illuminant.

The experimental results validate the theoretical invariance property of the proposed approach. We have shown that LSAC SURF detection can be as accurate as the original SURF detection when the light source is (approximately) uniform. Furthermore, when the illumination conditions vary significantly (presence of two or more distinct light sources, variations in color, in scale, etc.) in an image, the proposed methodology is able to overcome the performance of the existing algorithm considerably. Future works will concern with conducting experimental tests in real robot localization datasets, comparing the drift of the proposed methodology with the original SURF algorithm.

Acknowledgments

This work was funded by the ERDF European Regional Development Fund through the COMPETE Programme and by National Funds through FCT - Portuguese Foundation for Science and Technology within project IntellWheels, RIPD/ADA/109636/2009. The first author thanks FCT for his PhD Grant SFRH/BD/60727/2009.

References

1. Bay, H., Ess, A., Tuytelaars, T., Gool, L.V.: Speeded-up robust features (surf). *Computer Vision and Image Understanding* **110**(3) (June 2008) 346–359
2. Ancuti, C., Bekaert, P.: Sift-cch: Increasing the sift distinctness by color co-occurrence histograms. In: *International Symposium on Image and Signal Processing and Analysis, Istanbul, Turkey (27-29 September 2007)* 130–135
3. Fan, P., Men, A.D., Chen, M.Y., Yang, B.: Color-surf: A surf descriptor with local kernel color histograms. In: *IEEE International Conference on Network Infrastructure and Digital Content, Beijing, China (6-8 November 2009)* 726–730
4. Abdel-Hakim, A.E., Farag, A.A.: Csift: A sift descriptor with color invariant characteristics. In: *IEEE Computer Society Conference on Computer Vision and Pattern Recognition - CVPR. Volume 2., New York, NY, USA, IEEE Computer Society (17-22 June 2006)* 1978–1983
5. Burghouts, G.J., Geusebroek, J.M.: Performance evaluation of local colour invariants. *Computer Vision and Image Understanding* **113** (January 2009) 48–62
6. Bosch, A., Zisserman, A., Muoz, X.: Scene classification via pls. In Leonardis, A., Bischof, H., Pinz, A., eds.: *Computer Vision ECCV 2006. Volume 3954 of Lecture Notes in Computer Science. Springer Berlin / Heidelberg (2006)* 517–530
7. van de Sande, K.E.A., Gevers, T., Snoek, C.G.M.: Evaluating color descriptors for object and scene recognition. *IEEE Transactions on Pattern Analysis and Machine Intelligence* **32**(9) (September 2010) 1582–1596
8. Buluswar, S.D., Draper, B.A.: Color recognition in outdoor images. In: *International Conference on Computer Vision, Bombay, India (04 - 07 January 1998)* 171–177
9. Finlayson, G.D., Hordley, S.D., Xu, R.: Convex programming colour constancy with a diagonal-offset model. In: *IEEE International Conference on Image Processing - ICIP. Volume 3., Genova, Italy (11-14 September 2005)* III-948–51
10. van de Weijer, J., Gevers, T., Gijssen, A.: Edge-based color constancy. *Image Processing, IEEE Transactions on* **16**(9) (2007) 2207–2214
11. Buchsbaum, G.: A spatial processor model for object color-perception. *Journal of the Franklin Institute-Engineering and Applied Mathematics* **310**(1) (July 1980) 1–26
12. Ebner, M.: How does the brain arrive at a color constant descriptor? In Mele, F., Ramella, G., Santillo, S., Ventriglia, F., eds.: *Advances in Brain, Vision, and Artificial Intelligence. Volume 4729 of Lecture Notes in Computer Science. Springer Berlin / Heidelberg (2007)* 84–93
13. Zeki, S., Marini, L.: Three cortical stages of colour processing in the human brain. *Brain* **121** (Sep 1998) 1669–1685
14. Gijssen, A., Gevers, T., van de Weijer, J.: Generalized gamut mapping using image derivative structures for color constancy. *International Journal of Computer Vision* **86**(2-3) (Jan 2010) 127–139
15. Barnard, K., Cardei, V., Funt, B.: A comparison of computational color constancy algorithms - part i: Methodology and experiments with synthesized data. *IEEE Transactions on Image Processing* **11**(9) (September 2002) 972–984
16. Schmid, C., Mohr, R., Bauckhage, C.: Evaluation of interest point detectors. *International Journal of Computer Vision* **37**(2) (June 2000) 151–172
17. Geusebroek, J.M., Burghouts, G.J., Smeulders, A.W.M.: The amsterdam library of object images. *International Journal of Computer Vision* **61**(1) (January 2005) 103–112