

On Colour Spaces for Change Detection and Shadow Suppression

Philipp Blauensteiner, Horst Wildenauer, Allan Hanbury, and Martin Kampel

Pattern Recognition and Image Processing Group – TU Vienna
 Favoritenstr. 9-11, 1040 Vienna, Austria – {blau, wilde, hanbury, kampel}@prip.tuwien.ac.at

Abstract *In this paper, common colour models for background subtraction and problems related to their utilisation are discussed. A novel approach to represent chrominance information more suitable for robust background modelling and shadow suppression is proposed. Our method relies on the ability to represent colours in terms of a 3D-polar coordinate system having saturation independent of the brightness function; specifically, we build upon an Improved Hue, Luminance, and Saturation space (IHLS). The additional peculiarity of the approach is that we deal with the problem of unstable hue values at low saturation by modelling the hue-saturation relationship using saturation-weighted hue statistics. The effectiveness of the proposed method is shown in an experimental comparison with an approach based on normalised RGB.*

1 Introduction

The underlying step of visual surveillance applications like target tracking and scene understanding is the detection of moving objects. Background subtraction algorithms are commonly applied to detect these objects of interest by the use of statistical colour background models. Most of the present systems exploit the properties of the normalised RGB to achieve a certain degree of insensitivity with respect to changes in scene illumination.

Hong and Woontack [10] apply the normalised RGB space in their background segmentation system. McKenna et al. [14] use this colour space in addition to gradient information for their adaptive background subtraction. The AVITRACK project [3] utilises normalised RGB for change detection and adopts the shadow detection proposed by Horprasert et al. [11].

Beside normalised RGB, representations of the RGB colour space in terms of 3D-polar coordinates (hue, saturation, and brightness) are used for change detection and shadow suppression in surveillance applications. François and Medioni [5] suggest the application of HSV for background modelling for real-time video segmentation. In their work, a complex set of rules is introduced to reflect the relevance of observed and background colour information during change detection and model update. Cucchiara et al. [1, 2] propose a RGB-based background model which they transform to the HSV representation in order to utilise the properties of HSV chrominance information for shadow suppression.

Our approach differs from the aforementioned in the way that we build upon the IHLS colour space, which is more suitable for background subtraction. Additionally, we propose the application of saturation-weighted hue statistics [8] to deal with unstable hue values at weakly saturated colours. Also, a novel technique to efficiently classify changes in scene illumination (e.g., shadows), which explicitly models the relationship between saturation and luminance has been devised.

The remainder of this paper is organised as follows: In Section 2 the normalised RGB is reviewed and the Improved Hue, Luminance and Saturation (IHLS) colour space is outlined. Section 3 gives a short overview over circular colour statistics. Section 4 presents how these statistics can be applied in order to model the background in image sequences. The conducted experiments and their results are presented in Section 5. Section 6 concludes this paper.

2 Colour Spaces

In this section, the Normalised RGB and IHLS colour spaces used in this paper are described.

2.1 Normalised RGB

The Normalised RGB space aims to separate the chromatic components from the brightness component. The red, green and blue channel can be transformed to their normalised counterpart by using the formulae

$$l = R + G + B, \quad r = R/l, \quad g = G/l, \quad b = B/l \quad (1)$$

if $l \neq 0$ and $r = g = b = 0$ otherwise [4]. One of these normalised channels is redundant, since by definition

$$r + g + b = 1 \quad (2)$$

Therefore, the Normalised RGB space is sufficiently represented by two chromatic components (e.g. r and g) and a brightness component l .

The Normalised RGB space, however, suffers from a problem inherent to the normalisation: low intensities result in instable chromatic components, as Kender [12] describes.

2.2 IHLS Space

The Improved Hue, Luminance and Saturation (IHLS) colour space was introduced in [7]. It is obtained by placing an *achromatic axis* through all the grey ($R = G = B$) points in the RGB colour cube, and then specifying the coordinates of each point in terms of position on the achromatic axis (brightness), distance from the axis (saturation

s) and angle with respect to pure red (hue θ^H). The IHLS model is improved with respect to the similar colour spaces (HLS, HSI, HSV, etc.) by removing the normalisation of the saturation by the brightness. This has the following advantages:

- The saturation of achromatic pixels is always low.
- The saturation is independent of the brightness function used. One may therefore choose any function of R , G and B to calculate the brightness.

It is interesting that this normalisation of the saturation by the brightness, which results in the colour space having the shape of a cylinder instead of a cone or double-cone, is usually implicitly part of the transformation equations from RGB to a 3D-polar coordinate space. This is mentioned in one of the first papers on this type of transformation [16], but often in the literature the equations for a cylindrically-shaped space (i.e. with normalised saturation) are shown along with a diagram of a cone or double-cone (for example in [6, 9]).

The following formulae are used for the conversion from RGB to hue θ^H , luminance y and saturation s of the IHLS space:

$$\begin{aligned} s &= \max(R, G, B) - \min(R, G, B) \\ y &= 0.2125R + 0.7154G + 0.0721B \\ c_{r1} &= R - \frac{G+B}{2}, \quad c_{r2} = \frac{\sqrt{3}}{2}(B-G) \\ c_r &= \sqrt{c_{r1}^2 + c_{r2}^2} \\ \theta^H &= \begin{cases} \text{undefined} & \text{if } c = 0 \\ \arccos\left(\frac{c_1}{c}\right) & \text{if } c \neq 0 \wedge c_2 \leq 0 \\ 360^\circ - \arccos\left(\frac{c_1}{c}\right) & \text{if } c \neq 0 \wedge c_2 > 0 \end{cases} \end{aligned} \quad (3)$$

where c_{r1} and c_{r2} denote the chrominance coordinates and $c_r \in [0, 1]$ the chroma. The saturation assumes values in the range $[0, 1]$ independent of the hue angle (the maximum saturation values are shown by the circle on the chromatic plane in Figure 1). The chroma has the maximum values shown by the dotted hexagon in Figure 1.

The inverse transformation is easily derivable. When using this representation, it is important to remember that the hue is undefined if $s = 0$, and that it does not contain much useable information when s is low.

One may ask why the CIELAB space is not used, as it is standardised and also separates chrominance and luminance information. The problem with the conversion from the RGB to CIELAB space is that the coordinates of the white point are needed. As surveillance cameras are required to work at all hours of the day in many different lighting conditions, this calibration is not possible. Instead of introducing extra uncertainty into the process by a white point estimation, it was decided to use an alternative representation of the RGB coordinates which attempts to separate the luminance and chrominance information.

3 Colour statistics

In a 3D-polar coordinate space, standard (linear) statistical formulae can be utilised to calculate statistical descriptors

for brightness and saturation coordinates. The hue, however, is an angular value, and consequently the appropriate methods from circular statistics are to be used. Before we introduce the concept of saturation-weighted hue statistics, a brief review on the application of circular statistical descriptors to the hue channel is given.

3.1 Hue statistics

Let $\theta_i^H, i = 1, \dots, n$ be n observations sampled from a population of angular hue values. Then, the vector \mathbf{h}_i pointing from $\mathbf{O} = (0, 0)^T$ to the point on the circumference of the unit circle, corresponding to θ_i^H , is given by the Cartesian coordinates $(\cos \theta_i^H, \sin \theta_i^H)^T$ ¹

The mean direction $\bar{\theta}^H$ is defined to be the direction of the resultant of the unit vectors $\mathbf{h}_1, \dots, \mathbf{h}_n$ having directions θ_i^H . That is, we have

$$\bar{\theta}^H = \arctan2(\mathcal{S}, \mathcal{C}), \quad (4)$$

where

$$\mathcal{C} = \sum_{i=1}^n \cos \theta_i^H, \quad \mathcal{S} = \sum_{i=1}^n \sin \theta_i^H \quad (5)$$

and $\arctan2(y, x)$ is the four-quadrant inverse tangent function.

Therefore, since

$$\mathcal{R} = \sqrt{\mathcal{C}^2 + \mathcal{S}^2} \quad (6)$$

is the length of the resultant, its mean length is

$$\bar{\mathcal{R}} = \frac{\mathcal{R}}{n}. \quad (7)$$

The mean length of the resultant is an indicator of the dispersion of the observed data. If the n observed directions θ_i^H cluster tightly about the mean direction $\bar{\theta}^H$ then $\bar{\mathcal{R}}$ will approach 1. Conversely, if the angular values are widely dispersed $\bar{\mathcal{R}}$ will be close to 0. The *circular variance* is defined as

$$\mathcal{V} = 1 - \bar{\mathcal{R}} \quad (8)$$

While the circular variance differs from the linear statistical variance in being limited to the range $[0, 1]$, it is similar in the way that lower values represent less dispersed data. Further measures of circular data distribution are given in [13].

3.2 Saturation-weighted hue statistics

The use of statistics solely based on the hue, as described above, has the disadvantage of ignoring the relationship between the chrominance components hue and saturation. For weakly saturated colours the hue channel is unimportant and behaves unpredictably in the presence of colour changes induced by image noise (e.g. sensor- or compression noise). In fact, for colours with zero saturation the hue is undefined.

By projecting the RGB-space onto the chromatic plane, we arrive at a hexagon, as shown in Figure 1. This hexagon is slightly deformed into a circle for the calculation of the saturation so that the maximum saturation is 1 for all hue

¹Note that, when using the IHLS space (Eq. 4), no costly trigonometric functions are involved in the calculation of \mathbf{h}_i , since $\cos(\theta_i^H) = c_{r1}/c_r$ and $\sin(\theta_i^H) = -c_{r2}/c_r$.




Figure 1: The chromatic plane of the IHLS color space.

values. As one can see, the chromatic components may be represented by means of Cartesian coordinate vectors \mathbf{c}_i with direction and length given by hue and saturation respectively. Using this intuitive approach, we are able to introduce the aforementioned relationship into the hue statistics by weighting the unit hue vectors \mathbf{h}_i by their corresponding saturations s_i .

Now, let (θ_i^H, s_i) , $i = 1, \dots, n$ be n pairs of observations sampled from a population of hue values and associated saturation values. We proceed as described in Section 3.1, with the difference that instead of calculating the resultant of unit vectors, the vectors \mathbf{c}_i , which we will dub *chrominance vectors* throughout this paper, have length s_i .

That is, we weight the vector components in Eq. 5 by the saturation values s_i

$$\mathcal{C}_s = \sum_{i=1}^n s_i \cos \theta_i^H, \quad \mathcal{S}_s = \sum_{i=1}^n s_i \sin \theta_i^H, \quad (9)$$

and replace \mathcal{C} and \mathcal{S} in Eq. 6 by their weighted counterparts \mathcal{C}_s and \mathcal{S}_s getting

$$\mathcal{R}_s = \sqrt{\mathcal{C}_s^2 + \mathcal{S}_s^2}. \quad (10)$$


Although other formulations are possible (see, e.g. [8]) we choose the mean resultant length of the chrominance vectors to be

$$\bar{\mathcal{R}}_n = \frac{\sqrt{\mathcal{C}_s^2 + \mathcal{S}_s^2}}{n}. \quad (11)$$

Consequently, for the mean resultant chrominance vector we get

$$\bar{\mathbf{c}}_n = (\mathcal{C}_s/n, \mathcal{S}_s/n)^T. \quad (12)$$

Here, the length of the resultant is compared to the length obtained if all vectors had the same direction and maximum saturation. Hence, $\bar{\mathcal{R}}_n$ gives an indication of the saturations of the vectors which gave rise to the mean of the chrominance vector, as well as an indication of the angular dispersion of the vectors. Alternatively, $\bar{\mathbf{c}}_n$ can be interpreted as the mean of mean vectors, where the vectors have dispersions given by the saturations s_i .


 Figure 2: The combined resultant length R .

In order to test if a mean chrominance vector $\bar{\mathbf{c}}_n$ is similar to an observed chrominance vector, we use the Euclidean distance in the chromatic plane:

$$D = \sqrt{(\bar{\mathbf{c}}_n - s_o \mathbf{h}_o)^T (\bar{\mathbf{c}}_n - s_o \mathbf{h}_o)} \quad (13)$$

Here, \mathbf{h}_o and s_o denote the observed hue vector and saturation respectively. A rational justification for the use of the Euclidean distance follows. We have (see Fig. 2)

$$D^2 = \bar{\mathcal{R}}_n^2 + s_o^2 - 2\bar{\mathcal{R}}_n s_o \cos(\bar{\theta}_n^H - \theta_o^H) \quad (14)$$

and

$$\mathcal{R}^2 = \bar{\mathcal{R}}_n^2 + s_o^2 + 2\bar{\mathcal{R}}_n s_o \cos(\bar{\theta}_n^H - \theta_o^H). \quad (15)$$

So that if the difference of the observed angle θ_o^H and the mean angle $\bar{\theta}_n^H$ is zero then $\mathcal{R} = \bar{\mathcal{R}}_n + s_o$ and $D = |\bar{\mathcal{R}}_n - s_o|$. Conversely, for increasing angular differences, \mathcal{R} and D approach $|\bar{\mathcal{R}}_n - s_o|$ and $\bar{\mathcal{R}}_n + s_o$ respectively. As one can see, the use of the Euclidean distance provides a simple and efficient means of taking into account the dispersions as well as the directions when comparing $\bar{\mathbf{c}}_n$ and \mathbf{c}_o .

4 The IHLS Background Model

With the theoretical foundations laid out in Section 3.2 we proceed with devising a simple background subtraction algorithm based on the IHLS colour model and saturation-weighted hue statistics. Specifically, each background pixel is modelled by its mean luminance μ_y and associated standard deviation σ_y , together with the mean chrominance vector $\bar{\mathbf{c}}_n$ and the mean Euclidean distance σ_D between $\bar{\mathbf{c}}_n$ and the observed chrominance vectors (see Eq. 13).

On observing the luminance y , saturation s , and a Cartesian hue vector \mathbf{h} for each pixel in a newly acquired image, the pixel is classified as foreground if:

$$|(y - \mu_y)| > \alpha \sigma_y \quad \vee \quad \|\bar{\mathbf{c}}_n - \mathbf{c}\| > \alpha \sigma_D \quad (16)$$

where $\mathbf{c} = s\mathbf{h}$ and α is the foreground threshold, usually set between 2 and 3.5.

In order to decide whether a foreground detection was caused by a moving object or by its shadow cast on the static background, we exploit the chrominance information of the IHLS space. A foreground pixel is considered as shaded background if the following three conditions hold:

$$y < \mu_y \wedge |y - \mu_y| < \beta \mu_y, \quad (17)$$

$$s - \bar{\mathcal{R}}_n < \tau_{ds} \wedge |s\mu_y - \bar{\mathcal{R}}_n y| < \tau_{sl}, \quad (18)$$

$$\|\mathbf{h}\bar{\mathcal{R}}_n - \bar{\mathbf{c}}_n\| < \tau_h, \quad (19)$$

where $\bar{\mathcal{R}}_n = \|\bar{\mathbf{c}}_n\|$ (see Eq. 11).

These equations are designed to reflect the empirical observations that cast shadows cause a darkening of the background and usually lower the saturation of a pixel, while having only limited influence on its hue [2]. The first condition (Eq. 17) works on the luminance component, using a threshold β to take into account the strength of the predominant light source. Eq. 18 performs a test for a lowering in saturation, comparable to the approach proposed by Cucchiara et al. [1, 2]. However, our approach is different in two respects. Namely, the use of the mean resultant vector length $\bar{\mathcal{R}}_n$ as an estimate of the mean saturation μ_s , and the explicit modelling of the relationship between saturation and luminance in the IHLS space. If, assuming a linear camera response, a shadow scales light intensity by a certain factor; then a pixel's R, G, and B values scale by the same factor; which is also exploited in the normalised RGB space (see Eq. 1). I.e., also the luminance y and saturation s are linearly scaled by the same factor. Finally, in (Eq. 19) the observed hue vector \mathbf{h} is scaled to the same length as the mean chrominance vector $\bar{\mathbf{c}}_n$ and their deviation is tested using the Euclidean distance.

The choice of the thresholds τ_{ds} , τ_{sl} , and τ_h , for now, is done empirically with the assumption that the variation of chrominance caused by shading is rather small. The experiments presented in the following section confirmed this assumption, since in all cases it sufficed to choose very small positive values for all three thresholds. In fact, τ_h and τ_{ds} were set to 0.0015 and 0.05 respectively for all tests. τ_{sl} had to be varied in the range of [0.02, 0.03].

5 Experiments and Results

We compared our approach described in Section 4 with the NRGB using the *Colour Mean and Variance* approach to model the background [17]. A pixel is considered foreground if $|o_c - \mu_c| > \alpha \sigma_c$ for any channel $c \in \{r, g, l\}$, where o_c denotes the observed value, μ_c its mean, σ_c the standard deviation, and α the foreground threshold.

Both background models are maintained by means of exponential weighted averaging [17] using different learning rates for background and foreground pixels. Furthermore, it is ensured that the σ values do not fall below a predefined value in order to avoid oversensitivity whenever pixels do not change over a long period of time.

For the normalised RGB space, shadow suppression was implemented based on Horprasert's approach [11]: each

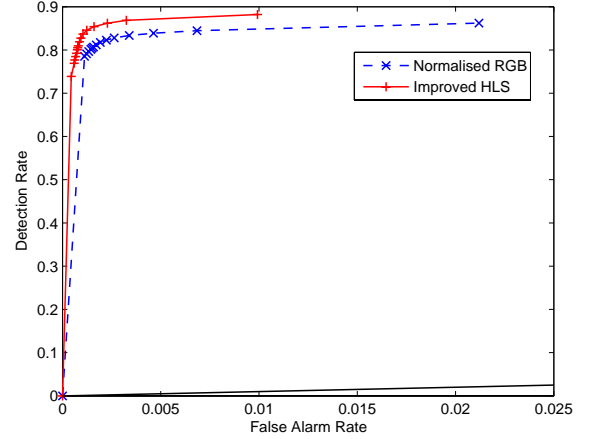


Figure 3: Receiver operating characteristic of normalised RGB and our approach evaluated on Test Sequence 1

foreground pixel is classified as shadow if:

$$o_l < \mu_l \wedge o_l > \beta \mu_l \wedge |o_r - \mu_r| + |o_g - \mu_g| < \tau_c \quad (20)$$

where β and τ_c denote thresholds for the maximum allowable change in the intensity and colour channels, so that a pixel is considered as shaded background.

As an example for an achromatic scene Test Sequence 1, recorded by an AXIS-211 network camera, shows a moving person in a stairway. For this sequence, ground truth was generated manually for 31 frames². The background model was initialised with 200 training frames. The same learning and update parameters were used for both background models. Before evaluation the segmented images are subject to morphological opening with a 3×3 structure element to eliminate salt noise. For the evaluation of the results, we defined the detection rate DR and the false alarm rate FR similar to Oberti et al. [15] as follows:

$$DR = \frac{TP}{FN + TP} \quad FR = \frac{FP}{N - (FN + TP)} \quad (21)$$

where TP denotes the number of true positives, FN the number of false negatives, FP the number of false positives, and N the total number of pixels in the image.

Test Sequence 2 was recorded with the same equipment mentioned above and shows moving objects in front of a coloured background. 50 frames were used for initialisation. Furthermore the approaches were tested on the *Laboratory* sequence of the ATON project³ (using the first 140 frames for initialisation) and on the camera 2 testing sequence of the PETS2001 Dataset 1 (using 100 frames for initialisation, starting with frame 2650).

5.1 Results

Figure 3 shows the receiver operating characteristic for both approaches evaluated on Test Sequence 1. The abscissa and

²Test Sequence 1 and Test Sequence 2 can be downloaded from the authors' homepage

³The video sequence may be found at <http://cvrr.ucsd.edu/aton/testbed/>.

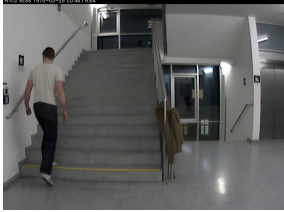









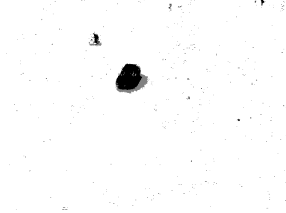
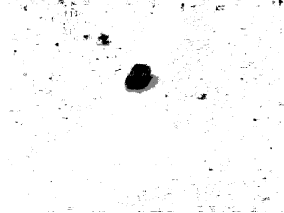
Sequence	Original Image	Our Approach	Normalised RGB
Test Sequence 1 Frame # 282			
Test Sequence 2 Frame # 70			
ATON Lab Frame # 685			
PETS2001 Dataset 1 Camera 2 Frame # 2982			

Table 1: Results of our approach in comparison with NRGB.

the ordinate denotes the false alarm rate and the detection rate, respectively. For this characteristic, the two approaches were tested for values of the foreground threshold α between 1.5 and 5.0. One can see that our approach yields better results in terms of low false alarm and high detection rate for the tested parameters.

Table 1 displays the results of the comparison of our approach and normalised RGB on the aforementioned sequences. For each sequence, parameters yielding the best balance between false alarm and detection rate are given in Table 2.

The example image of *Test Sequence 1* shows that the light beige t-shirt is better segmented in our approach, indicating a higher sensitivity w.r.t. to colour changes. In the results of the normalised RGB worse shadow detection performance as well more noise in dark background regions can be noticed. The comparison of the results in *Test Sequence 2* shows that the normalised RGB approach fails to classify the strong shadow in the lower left corner correctly. On the *Laboratory* sequence our approach yields a better segmentation of the right person as well as the inside of the opened locker. Finally, it can be noticed that for the *PETS2001* sequence, our approach yields a comparable segmentation result on the car; however, the present JPEG noise had greater impact on the normalised RGB approach.

6 Conclusion

We proposed the usage of the IHLS colour space for change detection and shadow suppression in visual surveillance tasks. In the proposed framework, we advocate the application of saturation-weighted hue statistics to deal with the problem of the unstable hue channel at weakly saturated colours.

Furthermore, we proposed a novel technique for detecting shadows exploiting the chrominance information and explicitly modelling the saturation-luminance relationship. The results on several challenging image sequences showed that our approach outperforms the normalised RGB colour space in terms of higher detection rate and less noise sensitivity.

One problem of our approach, however, is the fact that due to the use of saturation weighted hue statistics, it is impossible to tell whether a short chrominance vector in the background model is the result of unstable hue information or of a permanent low saturation. Although in the conducted experiments no impairments were evident, it is subject of further research in which cases this shortcoming poses a problem. Other fields of interest are the examination of alternatives to the Euclidean distance for the comparison of the chrominance vectors and an experimental in-depth investigation of the shadow classification.

		Sequence			
		Test Sequence 1	Test Sequence 2	Laboratory	PETS2001
NRGB	α	3.5	3.5	3.5	3.5
	β	0.5	0.9	0.5	0.5
	τ_c	0.02	0.1	0.02	0.05
Our Approach	α	3.5	3.5	3.5	3.5
	β	0.5	0.9	0.5	0.5
	τ_{ds}	0.05	0.05	0.05	0.05
	τ_{sl}	0.02	0.02	0.025	0.03
	τ_h	0.0015	0.0015	0.0015	0.0015
# Initialisation Frames		200	50	140	100
Foreground Learning Rate		0.9999			
Background Learning Rate		0.995			
Minimum σ -Value		0.01			

Table 2: Parameters used for evaluation of the approaches

Acknowledgement

This work was supported by the European Union, grant AVI-TRACK (AST3-CT-3002- 502818)⁴, the Austrian Science Foundation (FWF) under grant SESAME (P17189-N04), and the of the CABS project.

References

- [1] R. Cucchiara, C. Grana, M. Piccardi, and A. Prati. Detecting Moving Objects, Ghosts and Shadows in Video Streams. *IEEE Transactions on Pattern Recognition and Machine Intelligence*, 25(10):1337–1342, 2003.
- [2] R. Cucchiara, C. Grana, M. Piccardi, A. Prati, and S. Sirotti. Improving Shadow Suppression in Moving Object Detection with HSV Color Information. In *Intelligent Transport Systems*, pages 334–339. IEEE, 2001.
- [3] J. Ferryman, M. Borg, D. Thirde, F. Fusier, V. Valentin, F. Brémond, M. Thonnat, J. Aguilera, and M. Kampel. Automated Scene Understanding for Airport Aprons. In *Australian Joint Conference on Artificial Intelligence*, pages 593–503, Sydney, Australia, 2005.
- [4] Graham D. Finlayson, Bernt Schiele, and James L. Crowley. Comprehensive Colour Image Normalization. In *5th European Conference on Computer Vision*, pages 475–490, 1998.
- [5] Alexandre R. J. François and Gerard G. Medioni. Adaptive Color Background Modeling for Real-Time Segmentation of Video Streams. In *International Conference on Imaging Science, Systems, and Technology*, pages 227–232, 1999.
- [6] Rafael C. Gonzalez and Richard E. Woods. *Digital Image Processing*. Prentice-Hall, second edition, 2002.
- [7] Allan Hanbury. A 3D-polar coordinate colour representation well adapted to image analysis. In *Proceedings of the Scandinavian Conference on Image Analysis (SCIA)*, pages 804–811, 2003.
- [8] A. Hanbury and W. G. Kropatsch. Colour Statistics for Matching in Image Databases. In C Beleznaï and Th Schoegl, editors, *27th OEAGM Workshop*, Laxenburg, Austria, 2003.
- [9] D. Hearn and M. P. Baker. *Computer Graphics, C Version*. Prentice Hall, second edition, 1997.
- [10] D. Hong and W. Woo. A Background Subtraction for Vision-based User Interface. In *Fourth International Conference on Information, Communications and Signal Processing and Fourth IEEE Pacific-Rim Conference on Multimedia*, volume 1, pages 263–267, Singapore, 2003.
- [11] T. Horprasert, D. Harwood, and L.S. Davis. A Statistical Approach for Real-time Robust Background Subtraction and Shadow Detection. In *IEEE Conference on Computer Vision, FRAME-RATE Workshop*, 1999.
- [12] J. R Kender. Saturation, Hue and Normalized Colors: Calculation, Digitisation Effects, and Use. Technical report, Department of Computer Science, Carnegie Mellon University, 1976.
- [13] K. V. Mardia. *Statistics of Directional Data*. Academic Press, London, 1972.
- [14] S.J McKenna, S. Jabri, Z. Duric, H. Wechsler, and A. Rosenfeld. Tracking Groups of People. *Computer Vision and Image Understanding*, 80:42–56, 2000.
- [15] F. Oberti, A. Teschioni, and C.S. Regazzoni. Roc curves for performance evaluation of video sequences processing systems for surveillance applications. In *International Conference on Image Processing*, volume 2, pages 949– 953, Kobe, Japan, 1999. IEEE.
- [16] Alvy Ray Smith. Color gamut transform pairs. *Computer Graphics*, 12(3):12–19, 1978.
- [17] C. R. Wren, A. Azarbayejami, T. Darrel, and A. Pentland. Pfinder: Real-Time Tracking of the Human Body. *IEEE Transactions on Pattern Analysis and Machine Intelligence*, 19(7):780–785, 1997.

⁴However, this paper does not necessarily represent the opinion of the European Community, and the European Community is not responsible for any use which may be made of its contents.