# Advances in Daylight Statistical Colour Modelling

Daniel Alexander
*GRASP Lab., U. of Penn., 3401 Walnut St., Philadelphia, PA19104.*
*daniel2@grip.cis.upenn.edu*

## Abstract

*In this paper, parametric statistical modelling of distributions of colour camera data is discussed. A review is provided with some analysis of the properties of some common models, which are generally based on an assumption of independence of the chromaticity and intensity components of colour data. Results of an empirical comparison of the performance of various models are also reviewed. These results indicate that such models are not appropriate for situations other than highly controlled environments. In particular, they perform poorly for daylight imagery. Here, a modification to existing statistical colour models is proposed and the resultant new models are assessed using the same methodology as for the previous results. This simple modification, which is based on the inclusion of an ambient term in the underlying physical model, is shown to have a major impact on the performance of the models in less constrained daylight environments.*

## 1 Introduction

This paper addresses the problem of modelling distributions of colour data corresponding to particular objects of interest in scenes. The problem can be defined as follows: given a sample distribution of data known to correspond to an object of interest, construct a statistical model reflecting the parent distribution from which that sample was drawn. The parent distribution is the asymptotic distribution of data corresponding to the object of interest that arises as the number of observations tends to infinity.

There are many applications of statistical colour models (SCMs). Common examples include, object or feature recognition, [1], segmentation, [2,3], and tracking, [5-8]. The application motivating this investigation is the tracking of unmarked roads or lanes, viewed in daylight, through sequences of images obtained from a camera mounted on the front of an autonomous land vehicle. No assumption about the distribution of data comprising the background of the scene is made, as this may arbitrarily change during a single tracking task. Thus, no sample data from the background of the image is used in the construction of the SCMs.

There are many ways to construct a statistical model from a sample distribution of data, [9]. This work concentrates solely on parametric statistical models. Parametric models are suitable for problems of this type, as only a small amount of sample data is provided and separate models must be constructed for each individual tracking task. When large amounts of data are available a priori non-parametric techniques, such as neural networks or nearest neighbour rules, may be more appropriate.

**Figure 1 An image containing an unmarked road with the seed region outlined.**

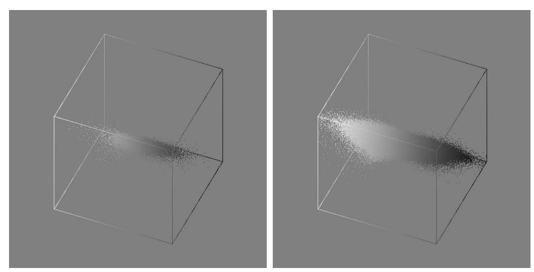

**Figure 2 RGB cubes showing the occurrence histograms of data in (left) the seed region of the track in Figure 1, and (right) the region of interest as a whole throughout the sequence.**

Figure 1 shows a typical image from the application described above. This image is the first of a sequence taken moving down the track in the image. A seed region is marked on the image, which provides a sample distribution of data from which to construct a SCM. Once the SCM is constructed, its parameters are fixed throughout the tracking task. The appearance of the object of interest does not change significantly through each sequence of images used here. Figure 2 shows occurrence histograms of data contained in the region of interest. The histogram on the left shows the sample distribution of data contained in the seed region of Figure 1. The one on the right shows the distribution of data contained in the entire region of interest, extracted by hand, throughout the sequence. This figure gives an idea of the shape of the parent distribution, which is to be captured by a SCM.

In the next section a review of existing SCMs is provided. Common models are described and some analysis of their properties is given. In section 3, a modification to these SCMs is described, which makes them more applicable in daylight situations. Experiments and results are presented in section 4, in which this modification is shown to have a major impact on performance. Finally, conclusions are drawn in section 5.

## 2 Existing SCMs

Several attempts have been made to find generic parametric models for distributions of colour data in the past. These existing SCMs are reviewed in this section.

### 2.1 Trivariate Gaussian model

A common statistical model is provided by the multivariate Gaussian distribution. The parameters of the distribution, the mean and covariance, are straightforward to estimate from a sample distribution of data by maximum likelihood estimation, [9]. Figure 3(left) shows a fixed likelihood boundary of a Gaussian model of the region of interest in Figure 1, constructed from the sample distribution shown in Figure 2. The fixed likelihood boundary has the ellipsoidal shape characteristic of the multivariate Gaussian distribution.

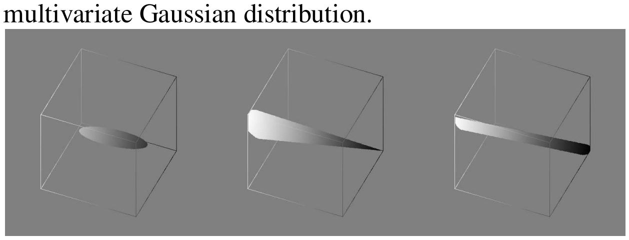

**Figure 3 Fixed likelihood boundaries in the RGB cube for SCMs, constructed from the sample distribution of Figure 2. Left: Gaussian SCM. Centre: Gaussian chromaticity, uniform intensity SCM. Right: Zhu and Yuille SCM.**

Trivariate Gaussian SCMs have been used in the past in, for example, [4]. Crisman, [8], and Thorpe, et al, [7], model clusters of colour data with Gaussian distributions; in both cases, however, multiple clusters are allowed. This model is included here to provide a benchmark against which to measure the performance of more sophisticated SCMs.

### 2.2 Chromaticity Models

A common way to construct a SCM is to assume a unichromatic model of reflection, [1-6], which affords the assumption that the intensity and chromaticity of the data are independently distributed. Such a reflection model assumes that there is one dominant component to the light reflected from the surface of the object of interest so that the spectral power density (SPD) of the reflected light is fixed up to a scalar factor depending on the geometry of the scene. Significant contributions from other reflection components, e.g., specular reflection, are assumed to be sparse, so that they can be considered noise. The SPD of the reflected light is therefore constrained to lie on a line passing through the origin of SPD space. The direction of this line is invariant with respect to the orientation of the surface and so characterises its reflection properties. More complete models of reflection, such as Shafer's dichromatic reflection model, have two (or more) reflection components, surface and body reflection, so that the SPD is constrained only to lie in a plane passing through the origin.

If the transformation made by the camera from the infinite dimensional space of SPDs to the 3D RGB cube is assumed to be linear, then distributions of RGB colour data have the same shape as the SPD distributions from which they were obtained. Thus data from a single surface patch is constrained to lie on a line passing through the origin of the RGB cube. The exact position on that line, i.e., the distance from the origin, or *intensity*, depends on the geometry of the scene, but the direction of the ray from the origin to the data point is fixed. The direction characterises the *chromaticity* component of the data.

Chromaticity based SCMs generally consist of a Gaussian model of normalised RGB data, which represent just the chromaticity component of the data. RGB vectors with the same chromaticity, i.e., that lie on the same line through the origin, are given the same likelihood, which implies that the model for the intensity component is uniform. These models therefore have conical contours in the RGB colour cube with elliptical cross-section. Figure 3(centre) shows a fixed likelihood boundary of such a model for the object of interest in Figure 1.

It is not clear, however, how a Gaussian model of directions should be constructed. In this section various approaches to that problem are reviewed.

First note that there are many colour spaces in which the chromaticity and intensity components are explicitly separated, for examples, see [10]. Many common colour spaces, such as the HSV space, are non-linear transformations of the original RGB space. Distributions of data in the RGB cube have different shape in these spaces, which can complicate the form of an appropriate SCM. Such spaces are generally intended for intuitive human navigation through the colour space rather than for the facilitation of automatic analysis of the data. All the models described here are constructed in the RGB cube or subspaces thereof.

One way to represent the direction of a line passing through the origin of a 3D space is by the point of intersection with a fixed plane, for example the unit plane, within that space. Ivins and Porrill, [4], use 2D Gaussian distributions in the normalised colour space (NCS). The NCS is equivalent to the unit plane in the RGB cube and normalised colour co-ordinates are obtained by taking the $L_1$ norm of an RGB vector to yield the intersection of the ray joining the origin to the data point with the unit plane.

The unit plane is a reasonable choice of fixed plane in which to represent directions, as the ray from the origin to every non-zero point in the RGB colour cube must intersect this plane. Matas, et al, [1], used similar 2D Gaussian models in the CIE chromaticity plane, which is equivalent to the plane B=1 in the RGB cube. This choice of fixed plane is less robust, however, as the chromaticity of data for which B=0 is not represented.

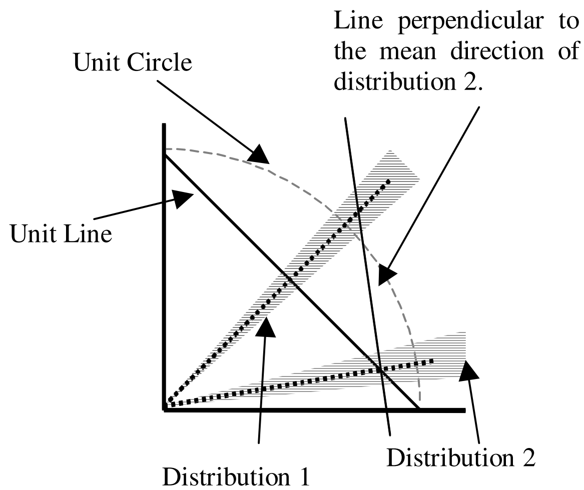


**Figure 4 Domains of representation for directional data.**

For simply characterising the chromatic properties of single items of colour data, the use of the unit plane is sufficient, as the $L_1$ norm is isomorphic to the set of positive directions from the origin, but there are problems when trying to characterise distributions. When the mean of a distribution of directions is close to one of the co-ordinate axes, the distribution of points of intersection with the unit plane is spread wider than a similar distribution whose mean is aligned with the normal to the plane. This effect is illustrated in 2D in Figure 4, where two directional distributions with similar variance are shown. The range over which intersections with the unit line occur, for a fixed range of likelihoods, is greater for the distribution that is not aligned with the normal to that line. A Gaussian model of distribution 2, in Figure 4, on the unit line will have greater variance than that of distribution 1, and will therefore provide a less tight model of the true distribution.

A consistently tight model can be ensured, as illustrated in Figure 4, by choosing a plane whose normal is aligned with the mean direction of the distribution. The distribution of directions can then be described in that plane by modelling the distribution of points at which the rays from the data points to the origin intersect the plane.

It is straightforward to construct a conical model, such as that shown in Figure 3(centre), using this approach, as described in [6]. Zhu and Yuille, [3], also used this chromaticity representation, but, in their model, lines *parallel* to the plane are lines of constant likelihood rather than all lines passing through the origin of the RGB cube. This gives their model a *cylindrical* shape, with elliptical cross-section. Such a model for the region of interest in Figure 1, is depicted in Figure 3(right). Zhu and Yuille argue that models constructed in this way are more robust to noisy, low intensity data for which the chromaticity is poorly defined, [3].

A more natural and consistent representation for the direction of a ray from the origin is by an $L_2$ unit vector, which is a point on the unit sphere. If a distribution has Gaussian profile on the unit sphere, then the projection of that distribution onto the unit plane is no longer Gaussian. The change of shape becomes more drastic if the mean direction of the distribution is not aligned with the normal to the unit plane. The distribution on the plane becomes asymmetric and skewed. This effect is illustrated well by Healey, [2], but can be seen by considering Figure 4 in which the unit circle is shown.

Healey, [2], devised a SCM based on the $L_2$ normalisation of RGB data for representation of chromaticity. First, a trivariate Gaussian model of the sample data is constructed from the seed region, exactly as described in section 2.1. The intensity component of the SCM is then made uniform by $L_2$ normalising the parameters of the Gaussian model. The mean of the model is found by normalising the mean vector of the trivariate Gaussian model and the covariance matrix is found by normalising the trivariate covariance matrix by the squared magnitude of the mean, see [2].

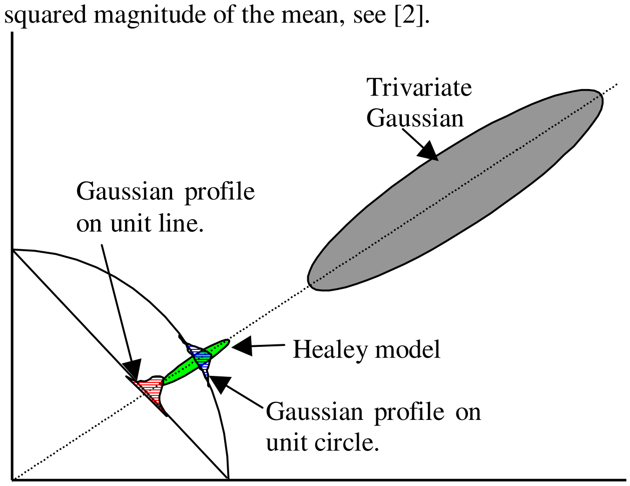


**Figure 5 Shows in 2D how Healey's SCM is related to the trivariate Gaussian model.**

The resultant model is still a trivariate Gaussian distribution and so is not strictly a distribution of points on the unit sphere. Although the mean is a point on the sphere, the covariance matrix still provides contours of the model with considerable depth perpendicular to the sphere in the RGB cube. Figure 5 shows how Healey's model relates to a trivariate Gaussian model of the full RGB data and Gaussian profiles on the unit plane and sphere. Healey's model has similar shape to the trivariate Gaussian model, whereas the others are flattened onto the chromaticity spaces in which they exist. Healey's model covers a much lesser portion of the sphere to within a particular fixed likelihood boundary as a large portion of its overall variance applies in a direction perpendicular to the sphere in the RGB cube.

With Healey's model, to assign likelihood to new data, the RGB vector is normalised to the unit sphere and the likelihood of the resultant 3D point is calculated from the normalised trivariate Gaussian model. The intensity component is thus factored out and so Healey SCMs have conical contours in the RGB cube like that shown in Figure 3(centre).

Healey's model is an ad-hoc approach to modelling distributions of spherical data. As pointed out by Alexander and Buxton, [5,6], there are established spherical analogues of the planar Gaussian distribution, such as the Bingham distribution, which can be used for modelling distributions of chromaticity data.

Alexander and Buxton, [5,6], compared the performance of several candidate SCMs, including all those described above, by computing the ROC (receiver operating characteristic) curve of a tracking algorithm driven by each SCM over two separate test sets of data. The first set of data consisted of images of matte items of clothing viewed under a single illuminant, so that the conditions of the unichromatic reflection model were well met. Over this data set, they found that conical shaped models performed much better than both the cylindrical, Zhu and Yuille, model and the Gaussian model. Within the conical models, SCMs based on the chromaticity plane oriented normal to the mean of the distribution were shown to produce better performance than fixed planar models. However, the best performance was obtained from the models based on a spherical chromaticity representation. Healey's model in fact produced better performance than the Bingham model. This result was attributed to the fact that Healey's model is more robust to outlying dark pixels in the sample distribution by virtue of its smaller range on the sphere, see Figure 5.

The second set of data consisted of unmarked roads and lanes like that shown in Figure 1. For these daylight images, although the performance ranking of the directional models was similar, it was found that none of them could match the performance of the Gaussian model. As the objects of interest are rigid in this data set, the distribution of intensities is quite constrained and so Alexander and Buxton combined a Gaussian intensity model with the chromaticity models. This improved performance somewhat, but the performance of the trivariate Gaussian SCM could still not be matched.

These results served as a reminder of the domain of applicability of directional SCMs. Although they work well in controlled environments where the assumptions of the unichromatic reflection model are well met, they do not generalise well to less constrained environments.

## 3 SCMs for Daylight

This section describes a simple modification to directional SCMs, which makes them more applicable to daylight environments, and other environments with complex lighting conditions.

In daylight there are two lighting components, a directional source from the sun and a hemispherical source from the sky. Shafer, [11], suggested the inclusion of an ambient term in the dichromatic reflection model to represent the contribution of light reflected from the sky. The effect of this term is to shift the SPD of reflected light away from the origin of SPD space. The SPD of reflected light is thus constrained to lie in a plane containing a fixed point described by the SPD of light reflected by the object of interest from the sky. This model implies that whatever direction the object is viewed from, the same amount of light is reflected, both body and surface reflection, from the hemispherical source. Koenderink and van Doorn, [12], pointed out that this ambient model is unrealistic because of local and global vignetting effects. Consider a surface patch at the bottom of a hole in a surface. As the hole gets deeper, a smaller fraction of the sky is visible from the bottom and so the amount of light incident on the surface patch decreases.

However, in many situations, particularly when the scene geometry is simple, the ambient model can be a reasonable approximation. For the objects of interest here, which are roughly planar and the orientation with respect to the sky varies very little, the ambient lighting term can be expected to provide a reasonable model.

In order to incorporate an ambient component into a directional model the axis of the distribution must be offset from a line through the origin. For the cylindrical Zhu and Yuille model, this is all that is required. It is interesting to note that in Alexander and Buxton's findings, the trivariate Gaussian SCM, significantly outperforms a Zhu and Yuille SCM combined with a Gaussian model for the intensity component. These two SCMs are almost identical apart from the fact that the major axis of the Gaussian model is not constrained to lie along a line that passes through the origin of the RGB cube, whereas the axis of the Zhu and Yuille model is. The added orientability of the Gaussian model allows it to find a more suitable major axis.

In order to construct the new, offset conical and cylindrical SCMs the axis of the distributions must first be estimated. Here a trivariate Gaussian model of the sample distribution is first computed and its eigenvectors are determined. The eigenvector that passes closest to the origin is used as the axis of the offset directional models. To obtain an offset Zhu and Yuille model, which incorporates a term for an ambient lighting model, a 2D Gaussian model of the data, projected onto a plane perpendicular to this offset axis, is constructed.

For the conical models, a position for the vertex of the cone along the offset axis must also be determined. Maximum likelihood estimation of the optimal vertex position can only be performed numerically. This is computationally intensive process and, furthermore, proves not to provide good estimates of the position (compared to empirically determined values). Geometric techniques for this estimation are currently under investigation, but in the experiments presented here, this value was set manually as described in the next section.

The offset axis typically misses the origin by about 10 grey levels and so illustrations of the offset conical model appear similar to before. For the conical models, the optimal position of the vertex is generally outside the RGB cube and well within the negative octant of $\Re^3$, the SCMs thus have almost cylindrical contours in the RGB cube. Fixed likelihood boundaries for both types of model are shown in Figure 6. Even if the fixed ambient lighting term corresponds to a single RGB vector, variation in the surface albedo means that a distribution of values will be observed when the object is viewed under ambient illumination alone. This places a distribution on the vertex of the conical distribution, which is captured in a single cone by moving the vertex further along the negative direction of the axis.

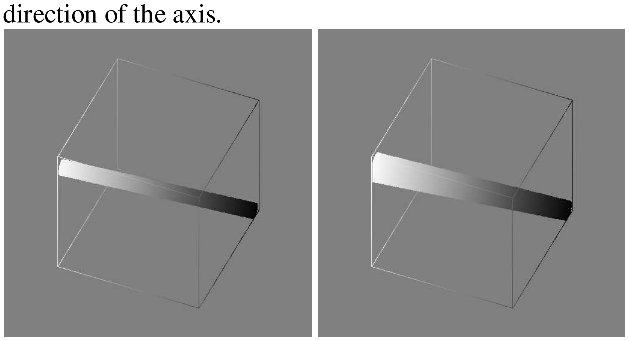

**Figure 6 Offset Zhu and Yuille (left) and Healey (right) SCMs of the track in Figure 1.**

## 4 Experiments and Results

Experiments have been performed on Alexander and Buxton's tracks data set, which is described in full in [6]. Their performance characterisation methodology, also detailed in [6], has been closely followed and the same tracking system has been used to assess the performance of the new offset directional SCMs.

The data set consists of six sequences of images, comprising about 200 images in total, all containing unmarked roads and lanes. Performance is assessed by comparison of the algorithm output with hand generated ground truth data, see [6]. For each of a number of settings of a parameter that controls how conservative the region boundaries extracted by the tracking algorithm are, counts of true and false positive pixel classifications are made. Thus, a ROC curve for the algorithm is generated for each SCM over the entire date set, which characterises the performance, see [6] for details. Two sequences are used as a training set to fix the parameters of the tracking algorithm for each SCM separately. The remaining four sequences, which comprise the test set, were used to produce a final ROC for each SCM.

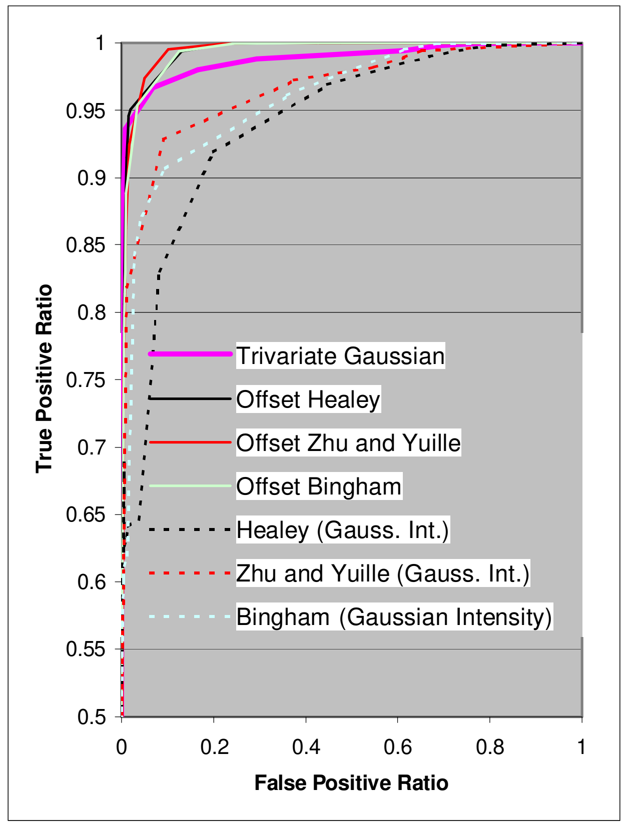


**Figure 7 Upper half of ROC curves for each SCM computed over the test set of data.**

Experiments have been performed using the following candidate SCMs: offset Zhu and Yuille model; offset Bingham model; and offset Healey model. For the latter two, conical models, the optimum position for the vertex of the distribution was determined by maximising the performance over the training data. It is then fixed at this value for each sequence in the test set. The intensity distribution was not constrained in any of these models, i.e., a uniform intensity model was implicitly adopted. The ROC curves computed over the test set of data for each of these models are shown in Figure 7. Figure 7 also shows the ROC curves computed by Alexander and Buxton for the trivariate Gaussian model; and Bingham Healey and Zhu and Yuille models, each combined with a

Gaussian intensity model, which was shown to enhance performance in each case.

A common measure of performance that is derived from ROC curves, [13], is the inverse cumulative normal distribution function, z, of the area under the curve, A. z(A) for each candidate SCM is tabulated in Table 1.

**Table 1 Performance measures for each candidate SCM.**

| SCM | z(A) |
|---|---|
| Healey offset | 2.48 |
| Zhu and Yuille offset | 2.46 |
| Bingham offset | 2.43 |
| Gaussian | 2.26 |
| Zhu and Yuille, Gauss. Int. | 1.86 |
| Bingham, Gauss. Int. | 1.75 |
| Bingham, Uniform Int. | 1.44 |

## 5 Conclusions

The current state-of-the-art in statistical colour modelling has been reviewed and some analysis of the properties of common models has been provided. Results, which suggest that the state-of-the-art is lacking in complex lighting environments such as daylight, have also been reviewed. A modification to commonly used directional models has been described, which is equivalent to the inclusion of an ambient lighting term in the underlying physical reflection model.

Results of a performance characterisation of these new models over a daylight image ensemble show that the inclusion of this term dramatically improves the performance of directional SCMs. These models, have been found to significantly outperform the default Gaussian model, which previous SCMs have been unable to do on this data. Note that if the offset Zhu and Yuille model were to be combined with a Gaussian intensity model, the resulting model would be identical to the trivariate Gaussian and so overall performance would decrease. Thus a uniform intensity model as been adopted with all the new models. There is extensive shadowing on some roads in the data set, which may cause the Gaussian intensity model to be less appropriate.

The major question to be addressed in future research on this subject is how to estimate the position of the vertex of conical models. The offset conical models show approximately the same level of performance as the offset, cylindrical model with the empirically selected and fixed position of the vertex. However, preliminary experimentation suggests that the actual optimal position of the vertex varies considerably between the different image sequences in the test set. If the optimal position could be determined accurately from the sample distribution of data, we may thus expect to observe further increases in overall performance. The difference in performance between the offset Bingham and Healey models is not considered significant, because of the fact that the vertex position may not be accurately estimated.

Other work for the future is to test the generalisation of these models to other data sets, for example, sets of data with less constrained object of interest shape, or with different types of illumination. For example, in overcast daylight, when there is less shading and shadowing better performance may be obtained from a constrained intensity model or even from the trivariate Gaussian SCM.

## References


[1] J. Matas, R. Marik, J. Kittler. "Generation, verification and localisation of object hypotheses based on colour". Proc. BMVC. 1993.

[2] G. Healey. "Segmenting Images Using Normalised Color". IEEE Transactions on Systems, Man, and Cybernetics. Vol. 22. No. 1. pp. 64-73. 1992.

[3] S.C. Zhu, A. Yuille. "Region Competition: Unifying Snakes, Region Growing, and Bayes/MDL for multiband Image Segmentation". IEEE Trans. PAMI. Vol. 18. No. 9. pp. 884-900. 1996.

[4] J. Ivins and J. Porrill. “Constrained Active Region Models for Fast Tracking in Colour Image Sequences”, CVIU. Vol. 72, No. 1, pp 54-71, 1998.

[5] D.C. Alexander, B.F. Buxton. "Modelling of Single Mode Distributions of Colour Data Using Directional Statistics". Proc. CVPR. pp. 319-324. 1997.

[6] D.C. Alexander, “Statistical Modelling of Colour Data and Model Selection for Region Tracking”, PhD Thesis, University of London, 1997.

[7] C. Thorpe, MH Herbert, T Kanade, SA Shafer. "Vision and Navigation for the Carnegie-Mellon Navlab". IEEE Trans. PAMI. Vol. 10. No. 3. pp. 362-373. 1988.

[8] J. Crisman. "Colour Region Tracking for Vehicle Guidance". in Active Vision edited by A. Blake and A. Yuille. pp. 107-122. MIT Press. 1992.

[9] R.O. Duda, P.E. Hart. "Pattern Classification and Scene Analysis". John Wiley and Sons, Inc. 1973.

[10] J. Foley, A. van Dam, S. Feiner, J. Hughes. "Computer Graphics". Addison Wesley. 1990.

[11] S.A. Shafer. "Using color to separate reflection components". Color: Research and Application. Vol. 10. No. 4. pp. 210-218. 1985.

[12] J.J. Koenderink, A.J. van Doorn. "Illuminance texture due to surface mesostructure". JOSA - A. Vol. 13. No. 3. pp. 452-463. 1996.

[13] J.A. Swets. "ROC analysis applied to the analysis of medical imaging techniques". Investigative Radiology. Vol. 14. pp. 109-121. 1979.