

## Attention-based color correction

Fred W. M. Stentiford<sup>\*a</sup>, Matt D. Walker<sup>b</sup>

<sup>a</sup> University College London, Adastral Park Campus, Martlesham Heath Ipswich, UK

<sup>b</sup> BT Research and Venturing, Adastral Park, Martlesham Heath, Ipswich, UK

### ABSTRACT

This paper proposes a new algorithm that extracts color correction parameters from pairs of images and enables the perceived illumination of one image to be imposed on the other. The algorithm does not rely upon prior assumptions regarding illumination constancy and operates between images that can be significantly different in content. The work derives from related research on visual attention and similarity in which the performance distributions of large numbers of randomly generated features reveal characteristics of images being analysed. A proposed color correction service to be offered over a mobile network is described.

### 1. INTRODUCTION

The color of illumination of a scene can have a dramatic affect on the performance of image retrieval systems. In addition different imaging devices will produce widely different responses. As there is normally no control over the camera characteristics, image preprocessing, the brightness or the color of the illumination, and the surface reflectances, this becomes a serious problem for object recognition and will lead to apparently identical images being assessed as different by the machine. It would not be acceptable for a photo taken in the late afternoon to reside in a different class to an identical photo taken in the morning.

Color constancy seeks a relationship between colors and surface illumination so that the recorded color can be mapped to the correct one [42]. In the absence of a precise definition of the 'correct' illumination in each case, most approaches to the color correction problem have to make assumptions about the statistics of the reflectances and the illuminants that will be encountered. The basic Retinex computation achieves a good measure of color constancy by mapping the maximum value of each channel to white, but sometimes forces scenes dominated by a single color to become gray [39]. Finlayson et al [34] used a diagonal model of illumination change and applied histogram equalization with some success but found that performance was lower on images with spatially varying illumination. The Color by Correlation approach [45] requires a set of training illuminants that encompasses the images in question and furthermore it must also overcome sources of mismatch between the model and the real world. Jackowski [43] corrects color distortions introduced by a camera by using known color charts to calibrate transformation functions that relate colors on the chart to those in the image. Reinhard et al [44] imposes one images' color characteristics on another by matching the average colors in Lab color space. The problem of comparing images possessing different compositions is addressed by comparing swatches taken from key parts of each image.

Human vision possesses an ability to interpret scenes without being troubled by the brightness or color of the illumination and should be an inspiration for new ideas in this field. Work on modeling the functionality of human vision is proving fruitful in the areas of visual attention [41], symmetry detection [40] and visual similarity [33].

Visual attention plays a major function in human vision both for our survival and our ability to spot patterns in our environment. Models of attention endeavour to predict which parts of an image or video will attract our attention and promise to find application in data compression, content base image retrieval, automated visual inspection and other sectors where human vision is currently an essential component. Measures of attention within images are closely related to measures of similarity between images because attention normally is drawn towards differences between foreground and background and this is necessarily determined by the large scale similarity of the background with *itself*. The

---

\* [f.stentiford@ucl.ac.uk](mailto:f.stentiford@ucl.ac.uk); <http://www.ee.ucl.ac.uk/~fstentif/>

identification of visual anomalies is therefore in many ways synonymous with detecting similarities elsewhere in an image and this suggests that attention mechanisms may be reused in some form to estimate similarity.

The work outlined in this paper is based on the notion that color constancy arises in human vision as a result of experiencing visual similarity rather than of some absolute definition of color correctness. This paper describes a new approach to illuminant estimation and correction as part of the need to obtain similarity estimates that are independent of the illumination.

### 1.1. Similarity measures

Similarity measures are central to most pattern recognition problems not least in computer vision and the need to access huge volumes of multimedia content now being broadcast and offered on the Internet. These problems have motivated considerable research into content based image retrieval [1,2,3,4,5] and many commercial and laboratory systems are described in the literature [6-13]. There are several approaches to similarity and pattern matching and much of this is covered in several survey papers [14-19]. Many approaches involve the use of pre-determined features such as edges, color, location, texture and functions dependent on pixel values e.g. [20]. Mikolajczyk et al [21] use edge models to obtain correspondences with similar objects. The advantages and disadvantages of using 3D histograms in which bins represent location are investigated by Ankerst et al [22].

The selection of features dependent upon the spatial arrangement of sets of points sampled from shapes is a strategy used by several authors to obtain some of the best results to date [23-25]. However, all the approaches use pre-determined point selection rules and metrics that can limit performance on unseen data. Viola et al [26] restrict themselves to a specific type of rectangle feature which works well in their face recognition application, but may not perform as well on data that is not suited to this feature.

Increasingly research is turning to models of perception in order to reflect the behaviour of the human visual system in measures of similarity. Mojsilovic et al [27] use perceptually important colors to construct a feature vector for similarity measurement, and overcome the problem of close colors occupying different quantization bins. Law et al [28] introduce a measure of saliency in their development of a feature selection and clustering algorithm. A feature is deemed irrelevant if its distribution is independent of class labels. Visual attention models by Itti [29] are used by Frintrop et al [30] to focus computational resource and recognise 3D objects. Shape contours are detected by Grigorescu et al [31] using a model of human visual surround suppression that identifies perceptually significant edges.

It is almost universal that established approaches to pattern analysis make use of a priori features to distinguish and recognise classes of data [Duda et al]. Unless the universe of data is completely understood it will always be possible to produce unseen data on which these systems will fail. The nature of visual data certainly is such that no single fixed set of features appears capable of encompassing all the complex relationships that exist between each item. In fact the strength of similarity between each pair of patterns in a class could be on a different basis in every case with no two patterns sharing the same features as any other pair. High dimensional feature spaces offer no solution to this problem.

### 1.2. Cognitive visual attention

Studies in neurobiology [32] are suggesting that human visual attention is enhanced through a process of competing interactions among neurons representing all of the stimuli present in the visual field. The competition results in the selection of a few points of attention and the suppression of irrelevant material. It means that people and animals are able to spot anomalies in a scene no part of which they have seen before and attention is drawn in general to the anomalous object in a scene.

Such a mechanism has been extended to apply to the comparison of two images in which attention is drawn to those parts that are in common rather than their absence as in the case of saliency detection in a single image [33]. Whereas saliency measures require no memory of data other than the image in question, cognitive attention makes use of other stored material in order to determine similarity with an unknown image.

The model of Cognitive Visual Attention (CVA) used in this paper relies upon the matching of large numbers of pairs of pixel groups (forks) taken from patterns A and B under comparison.

Let a location  $\mathbf{x}$  in a pattern correspond to a measurement  $\mathbf{a}$  where

$$\mathbf{x} = (x_1, x_2) \text{ and } \mathbf{a} = (a_1, a_2, a_3)$$

Define a function  $F$  such that  $\mathbf{a} = F(\mathbf{x})$ .

Select a fork of  $m$  random points  $S_A$  in pattern A where

$$S_A = \{x_1, x_2, x_3, \dots, x_m\}.$$

Similarly select a fork of  $m$  points  $S_B$  in pattern B where

$$S_B = \{y_1, y_2, y_3, \dots, y_m\} \text{ where} \\ x_i - y_i = \delta_j$$

The fork  $S_A$  matches the fork  $S_B$  if

$$|F(x_i) - F(y_i)| < \varepsilon \quad \forall i \text{ for some } \delta_j, j = 1, 2, \dots, N$$

An example of a  $m = 4$  fork fitting images A and B is shown in Figure 1. In general  $\varepsilon$  is not a constant and will be dependent upon the measurements under comparison i.e.

$$\varepsilon_j = f_j(F(x), F(y))$$

In effect up to  $N$  selections of the displacements  $\delta_j$  apply translations to  $S_A$  to seek a matching fork  $S_B$ .

The CVA similarity score  $C_{AB}$  is produced after generating and applying  $T$  forks  $S_A$  :

$$C_{AB} = \sum_{i=1}^T w_i \quad \text{where } w_i = \begin{cases} 1 & \text{if } S_A \text{ matches } S_B \\ 0 & \text{otherwise} \end{cases}$$

$C_{AB}$  is large when a high number of forks are found to match both patterns A and B and represents features that both patterns share. It is important to note that if  $C_{AC}$  also has a high value it does not necessarily follow that  $C_{BC}$  is large because patterns B and C may still have no features in common. The measure is not constrained by the triangle inequality and is able to model a greater degree of complexity than a more restricted metric.




Figure 1: Four pixel fork matching image A and image B with  $\delta_j = (3, -1)$

## 2. RELATIVE ILLUMINANT ESTIMATION

In this work we have applied random color shifts during the fork matching process as a means of obtaining illumination independence when calculating the CVA similarity measure. No restriction is placed on the choice of color shift (which effectively varies brightness) and will change for each new fork generated. In this case the fork  $S_A$  matches the fork  $S_B$  if

$$|F(x_i) - F(y_i) + \alpha_j| < \varepsilon \quad \forall i \text{ for some } \alpha_j, \delta_j, j = 1, 2, \dots, N$$

where  $\alpha_j = (\alpha_{j1}, \alpha_{j2}, \alpha_{j3})$  is a valid random displacement in the pixel color components in image A. Peaks  $\mu_{AB}$  in the distribution of  $\alpha_j$  for  $j$  corresponding to matches will reflect the overall difference in illumination between images A and B providing sufficient common geometry and fork matches are found.  $\mu_{AB}$  being an average does not reflect any variation in illumination across the image although by restricting diameter of forks an estimate of  $\mu_{AB}(y_i)$  as it varies over image B may be obtained. This approach has the advantage that no prior assumptions are made about the camera or the nature of the illumination that might impose anomalous results.

The color shift  $\mu_{AB}$  that stimulates the highest frequency of matches represents a measure of the relative illumination of image B with respect to image A. We apply the reverse color shift  $-\mu_{AB}$  to pixels in image B to obtain a transformation that approximates the illumination present in image A. In the results below  $\mu_{AB}$  was taken to be the average of the color shifts  $\alpha_{AB}$  that were applied when a match resulted:

$$\mu_{AB} = \frac{\sum_{S_A \text{ matches } S_B} \alpha_{AB}}{C_{AB}}$$

If the original reference image A is geometrically identical to the color corrected version B it becomes possible to measure the percentage color deviation  $e_{AB}$  between the two images on a pixel by pixel basis:

$$e_{AB} = \frac{\sum_{ij} \sum_{k=r,g,b} |x_{ijk} - y_{ijk}|}{\sum_{ijk} 2.55}$$

where  $x_{ijk}$  and  $y_{ijk}$  are corresponding pixel color component values (0-255) in each image.

### 3. RESULTS

The number of pixels ( $m$ ) to be included in each fork is a critical parameter. If  $m$  is too great the number of matches decreases leading to poor statistics and higher computation; if  $m$  is too small image structure may be overlooked leading to inaccurate color correction. The image color correction error  $e_{AB}$  was computed for images A and B in Figure 2 [47] for several values of  $m$  and averaged over 5 runs with  $T = N = 1000$ . The performances are shown in Table 1.

$m$	1	2	3	4	5	6	7	8
$e_{AB}$	7.86	6.00	4.28	3.99	4.02	4.00	4.00	4.01
Standard deviation	0.74	0.04	0.12	0.05	0.06	0.04	0.04	0.15

Table 1. Color correction errors for images A and B in Figure 2

The error rates in Table 1 indicate that little is to be gained by increasing  $m$  beyond 4 not least because computation is proportionate to  $m$ . In the case of  $m = 1$  the correction is heavily dependent upon the relative color values of the largest areas of co-colored pixels in each image and gives an approximate result only when the images are geometrically similar.

Subsequent images are all processed with  $T = 500$  randomly generated forks each possessing  $m = 4$  pixels with  $N = 500$ . The distribution of color shifts for red green and blue for matching forks between images A and B in Figure 2 is shown in Figure 3. In this case  $\mu_{AB} = (-28,42,-2)$  indicating that a reduction in red values and an increase in green encouraged a greater likelihood of a match. By the same token this same shift was a measure of the change in illumination between image A and Image B and applying a color shift of  $-\mu_{AB}$  produces the corrected version with an error value of 3.89%.

In Figure 4 image A is illuminated with yellow light and image B with blue light [35]. The difference in illumination is obtained with the same parameters as used in Figure 2 and gives  $\mu_{AB} = (-27,9,57)$ . The reverse shift in pixel colors removes the blue tinge and restores the yellow illumination to image B with an error value of 4.05%.

Figure 5 [36] shows an image (Image A) which has been manually restored from a color cast version (Image B). A color shift of  $-\mu_{AB} = (17,35,-33)$  produced the corrected version which compares favourably with the manually corrected version. In this case the original and color cast versions are not identical in composition and size.

As in Figure 5 it is not necessary for the images A and B to be identical in all respects except illumination. Figure 6 shows a cropped version of image A in Figure 2 together with the same image B as Figure 2. The resulting color shift  $\mu_{AB} = (-33,38,-9)$  yields an error of 4.16% only marginally greater than that obtained using geometrically identical images.

More significantly the images do not have to be of the same scene. Unlike Reinhard et al [46] we do not need to select areas within each image for specific comparison because the statistics of matching pairs of forks already take this into account. Matching forks which span important color adjacencies in different positions in each image contribute towards the estimate of the relative illumination regardless of the detailed composition of the images. Image A in Figure 7 is illuminated with sunlight [37]. Image B [38] has a color cast which is common in digital photography. This is removed with a color shift of  $-\mu_{AB} = (-23,-5,-15)$  and an error of 3.92% and may be compared with the manually treated image on the right.

The fact that relative illumination may be calculated from a sample of images exhibiting a preferred set of illuminations means that working images may be processed to appear under a range of lighting conditions without it being necessary to make special arrangements. Furthermore images may be tailored automatically to individual preferences by simply supplying examples of acceptable images and allowing the system to adjust the color balance accordingly.



Figure 2: Original image A, illuminated version B, and corrected version  $\mu_{AB} = (-33,39,-4)$ , error = 3.89%




Figure 3: Distributions of R,G,B color shifts for matching forks for images A and B in Figure 2.




Figure 4: Image A in yellow light, image B in blue light and corrected  $\mu_{AB} = (-28,6,56)$ , error = 4.05%.



Figure 5: Manually restored image A, image B with color cast, and corrected version  $\mu_{AB} = (-17,-35,33)$



Figure 6: Cropped image A in Figure 2, Image B in Figure 2, and corrected version  $\mu_{AB} = (-33,38,-9)$ , error = 4.16%



Figure 7: Image A, Image B with blue color cast, corrected version  $\mu_{AB} = (-23,-5,-15)$ , error = 3.92%, and manually corrected version.

#### 4. DISCUSSION

It is important to note that this approach identifies color shifts that maximise the similarity of two images. In this way it gives a good indication of a possible change in illumination between the two images, but it does not provide an absolute estimate for the illumination in either image. It is possible that with more information gleaned from more images under different illuminations good estimates of illuminants can be obtained

The color shifts  $\mu_{AB}$  extracted in this work are based on averages obtained from fork matches over the entire area of the images. They therefore do not take account of any variation of perceived illumination across the image but arrive at a compromise solution and therefore a possible source of error. Spatially differentiated color correction

could be obtained by first recording the positions and color shifts applied to matching fork pixels. A smoothed 2-dimensional profile of color shifts could then be used to correct images and obtain better conformance with the ideal result.

Shadows are sometimes a problem for object recognition and this approach offers some scope for minimizing their effects. The color and brightness of objects immersed in shadow have a relationship with those from the same object that are not immersed in shadow. It may be the differences are simply one of brightness but commonly the colors in shadow are shifted by other sources of light. It would be reasonable to expect that spatial color correction between similar images with and without shadows will reduce their prominence.

The computational requirements of this technique are independent of the size of the image, but increase as the pair of images under comparison become less similar. Dissimilar images will not yield many matching forks and the color correction statistics will become less reliable as the number of matches falls. Processing of geometrically similar images takes 12 sec on a 1Ghz machine with  $T = N = 1000$  using Visual Basic software. This rises to 25 sec for dissimilar images. Equivalent code in C++ takes under 100ms and related implementations using some parallelism on the Texas Instruments DM642 DSP platform indicate that processing can take place at video speeds.

## 5. REAL WORLD USAGE of ALGORITHM

The techniques used will be exposed to end-users in a real-world environment, utilising BT's Broadband Applications Testbed (BAT). The BAT enables new technologies to be rapidly deployed as new services, with the goal of understanding both the technical performance, and end-user acceptance and issues arising from the new service. In this case, end-users can upload photos taken from their camera-phones, and receive color-corrected images back to the same devices.

Figure 8 below shows the logical high-level architecture of the color algorithm as a web-service, available in a fixed-mobile convergence (FMC) scenario, where end-users can upload images from either their mobile phone using the cellular network (in this case MMS) and also from their PC within a broadband environment. In practice, the end-user tests will only utilise the mobile instantiation.

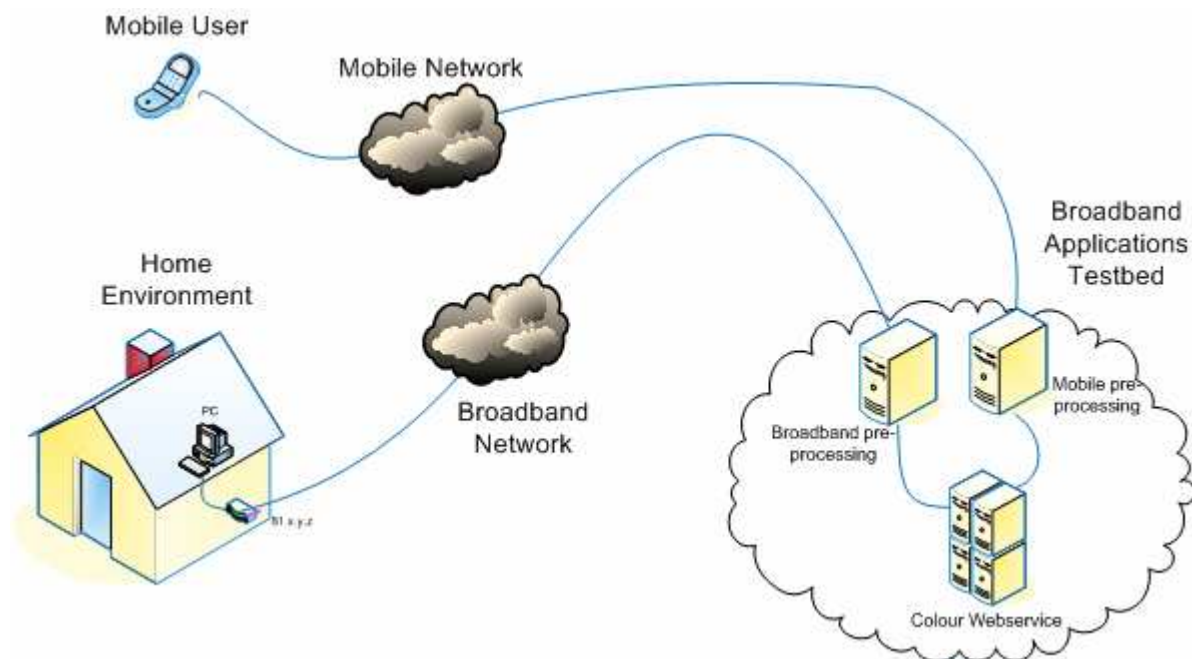


Figure 8: BT's Broadband Applications Testbed, showing high-level architecture for Color as a web-service.

For implementation from a mobile handset, users will send MMS messages containing both the image to correct, and reference images – the front-end mobile server will pass the two images to the web service running the color-correction algorithm, and will then send the corrected image back to the end-user via the MMS channel. Users will also have the option of selecting reference images from a centrally provided set.

Feedback from the triallists is expected to dictate how such a service could be improved – for example, by uploading the corrected images to a network-based file store that can be accessed from both PC and mobile handset.

Potential embodiments of a future service, could then be provided directly by a service provider through an online portal. One advantage of providing the implementation as a distributed web service however, means that the technology can be provided as a wholesale network service. This is beneficial for service providers in the retail space as they can leverage upon not only the technology, but computational resources provided by the wholesaler (if providing the service to large numbers of consumers). This approach leads to dramatically reduced overheads usually associated with installing costly infrastructure for computationally intense processes, and reducing the time to market.

## 6. CONCLUSIONS

This paper has presented a new way of adjusting the perceived illumination between pairs of images. The pairs of images do not have to be geometrically identical, as the reference image can be a portion of the image being corrected or can contain completely different content. More investigations are necessary to explore the performance on a greater range of image content and the extent to which the images can become structurally dissimilar before the correction becomes unreliable. The results, however, are encouraging and should find application in the optimization of camera settings and a color correction service to be offered over a mobile network..

This research has been conducted with the support of British Telecom and within the framework of the European Commission funded Network of Excellence “Multimedia Understanding through Semantics, Computation and Learning” (MUSCLE) [46].

## REFERENCES

1. Y. Rui, T. S.Huang, and S. F. Chang, “Image retrieval: current techniques, promising directions and open issues,” *J. Visual Communication and Image Representation*, Vol. 10, No. 1, pp. 39 – 62, 1999.
2. A. W. M. Smeulders, M. Worring, S. Santini, A. Gupta and R. Jain, “Content-based image retrieval at the end of the early years,” *IEEE Trans. Pattern Anal. Mach. Intell.*, Vol. 22, No. 12, pp. 1349-1379, 2000.
3. C. C. Ventners and M. Cooper, “A review of content-based image retrieval systems,” *JISC Technology Applications Programme*, <http://www.jtap.ac.uk/reports/htm/jtap-054.html>.
4. R. C. Veltkamp and M. Tanase, “Content-based image retrieval systems: a survey,” <http://www.aa-lab.cs.uu.nl/cbirsurvey/cbir-survey/>, March, 2001.
5. E. Izquierdo, V. Mezaris, E. Triantafyllou, and L-Q Xu, “State of the art in content-based analysis, indexing and retrieval,” *IST Project IST-2001-32795*, Network of Excellence in Content-Based Semantic Scene Analysis and Information Retrieval, <http://www.iti.gr/SCHEMA/library/index.html>, Sept, 2002.
6. M. Flickner, H. Sawhney, W. Niblack, J. Ashley, Q. Huang, B. Dom, M. Gorkani, J. Hafner, D. Lee, D. Petkovic, D. Steele and P. Yanker, “Query by image and video content: the QBIC system,” *IEEE Computer*, 1995.
7. Pentland, A., Pickard, R.W., and Sclaroff, S. Photobook: content-based manipulation of image databases. *Proceedings of SPIE Storage and Retrieval of Images and Video Databases*, Vol. 2185, San Jose, CA, USA, 1994.
8. Carson, C., Belongie, S., Greenspan, H., and Malik, J. Blobworld: image segmentation using expectation-maximisation and its application to image querying. *IEEE Trans. Pattern Anal. Mach. Intell.* 24(8) (2002) 1026-1038.

9. Smith, J.R., and Chang, S-F., VisualSEEK: a fully automated content-based image query system. *Proceedings of the ACM International Conference on Multimedia*, Boston MA, USA, 1996, pp 87-98.
10. Ma, W-Y., and Manjunath, B.S. NeTra: a toolbox for navigating large image databases. *Multimedia Systems*, 7 (1999), 184-198.
11. Gupta, A. and Jain, R. Visual information retrieval. *Communications of the ACM*, 40(5) (1997) 70-79.
12. Dowe, J. Content Based Retrieval in Multimedia Imaging. *Proceedings of SPIE Storage and Retrieval for Image and Video Databases*, 1993, pp 164-167.
13. Bach, J., Fuller, C., Gupta, A., Hampapur, A., Gorowitz, B., Humphrey, R., Jain, R., and Shu, C. Virage image search engine: an open framework for image management. *Proceedings of SPIE Storage and Retrieval for Image and Video Databases IV*, San Jose, CA, USA, 1996, pp. 76-87.
14. Veltkamp, R.C. Shape matching: similarity measures and algorithms. *Technical Report UU-CS-2001-03*, Utrecht, January 2001.
15. Santini, S., and Jain, R. Similarity measures. *IEEE Trans. Pattern Anal. Mach. Intell.* 21(9) (1999) 871- 883.
16. Tangelder, J.W.H., and Veltkamp, R.C. A survey of content based 3D shape retrieval methods. *Proceedings of Shape Modelling International Conference*, 2004, pp. 145-156.
17. Loncaric, S. A survey of shape analysis techniques. *Pattern Recognition*, 31(8) (1998) 983-1001.
18. Cardone, A., Gupta, S.K., and Karnik, M. A survey of shape similarity assessment algorithms for product design and manufacturing applications. *J. of Computing and Information Science in Engineering*, 3(2) (2003) 109-118.
19. Zhang, D., and Lu, G. Review of shape representation and description techniques. *Pattern Recognition*, 37 (2004) 1-19.
20. Manmatha, R., Ravela, S., and Chitti, Y. On computing local and global similarity in images. *Proceedings of SPIE Human and Electronic Imaging III*, 1998.
21. Mikolajczyk, K., Zisserman, A., and Schmid, C. Shape recognition with edge-based features. *Proceedings of the British Machine Vision Conference*, Norwich, UK, 2003.
22. Ankerst, M., Kastenmuller, G., Kriegel, H-P., and Seidl, T. 3D histograms for similarity search and classification in spatial databases. *Proceedings of the International Symposium on Spatial Databases*, Hong Kong, 1999.
23. Belongie, S., Malik, J., and Puzicha, J. Shape matching and object recognition using shape contexts. *IEEE Trans. Pattern Anal. Mach. Intell.*, 24(4) (2002) 509-522.
24. Grigorescu, C., and Petkov, N. Distance sets for shape filters and shape recognition. *IEEE Trans on Image Processing*, 12(10) (2003) 1274-1286.
25. Osada, R., Funkhouser, T., Chazelle, B., and Dobkin, D. Matching 3D models with shape distributions. *Proceedings of the Shape Modelling International Conference*, Genova, Italy, 2001, pp. 154-166.
26. Viola, P., and Jones, M. Rapid object detection using a boosted cascade of simple features. *Proceedings of Computer Vision and Pattern Recognition*, Vol. 1, 2001, pp 511-518.
27. Mojsilovic, A., Hu, J., and Soljanin, E. Extraction of perceptually important colors and similarity measurement for image matching, retrieval, and analysis. *IEEE Trans. on Image Processing*, 11(11) (2002) 1238-1248.
28. Law, M.H.C., Figueiredo, M.A.T., and Jain, A.K. Simultaneous feature selection and clustering using mixture models. *IEEE Trans. on Pattern Anal. Mach. Intell.* 26(9) (2004) 1154-1166.
29. Itti, L. Automatic foveation for video compression using a neurobiological model of visual attention. *IEEE Trans. on Image Processing*, 13(10) (2004) 1304-1318.
30. Frintrop, S., Nuchter, A., and Surmann, H. Visual attention for object recognition in spatial 3D data. *Proceedings of IEEE International Conference on Intelligent Robots and Systems*, Sendai, Japan, 2004.
31. Grigorescu, C., Petkov, N., and Westenberg, M.A. Contour detection based on nonclassical receptive field inhibition. *IEEE Trans. on Image Processing*, 12(7), (2003) 729-739.
32. Desimone, R. Visual attention mediated by biased competition in extrastriate visual cortex. *Phil. Trans. R. Soc. Lond. B*, 353, (1998) 1245 – 1255.
33. Stentiford, F.W.M. An attention based similarity measure with application to content based information retrieval. *Proceedings of SPIE Storage and Retrieval for Media Databases*, Vol 5021, Santa Clara, CA, USA, 2003.
34. Finlayson, G., Hordley, S., Schaefer, G., and Tian, G.Y. Illuminant and device invariant colour using histogram equalization. *Pattern Recognition*, vol 38, no 2, (Feb 2005), 179-190.
35. <http://www.staff.ncl.ac.uk/anya.hurlbert/>

36. [http://www.adobe.com/education/digkids/tips/pse/color\\_correction.html](http://www.adobe.com/education/digkids/tips/pse/color_correction.html)
37. <http://cilab.knu.ac.kr/research/Color/constancy.htm>
38. [http://www.panoguide.com/howto/panoramas/colour\\_correction.jsp](http://www.panoguide.com/howto/panoramas/colour_correction.jsp)
39. Z. Rahman, D. J. Jobson, and G. A. Woodell, "Retinex processing for Automatic Image Enhancement," Human Vision and Electronic Imaging VII, SPIE vol 4662, 2002.
40. F W M Stentiford, "Attention based facial symmetry detection," ICAPR, Bath, August 22-25, 2005.
41. F W M Stentiford, "A visual attention estimator applied to image subject enhancement and colour and grey level compression," ICPR 2004, Cambridge, 23<sup>rd</sup> – 26<sup>th</sup> August, 2004.
42. K Barnard, V Cardei, and B Funt, "A comparison of computational color constancy algorithms – Part 1: methodology and experiments with synthesised data", IEEE Trans on Image Processing, Vol. 11, no 9, pp 972-983, 2002
43. M Jackowski, A Goshtasby, S Bines, D Roseman and C Yu, "Correcting the geometry and color of digital images," IEEE Tans PAMI, Vol 19, no 10, pp1152- 1158, 1997.
44. E Reinhard, M Ashikhmin, B Gooch, and P Shirley, "Color transfer between images," IEEE Graphics and Applications, Vol 21, No. 5, pp 34-41, 2001.
45. G D Finlayson, S D Hordley, and P M Hubel, "Color by correlation: a simple, unifying framework for color constancy," IEEE Trans on PAMI, Vol 23, No 11, pp1209-1221, 2001.
46. Multimedia Understanding through Semantics, Computation and Learning, Network of Excellence. EC 6<sup>th</sup> Framework Programme. FP6-507752. <http://www.muscle-noe.org/>
47. <http://www.cs.sfu.ca/~colour/research/colour-constancy.html>