

Color and Color Constancy in a Translation Model for Object Recognition

Kobus Barnard¹ and Prasad Gabbur²

¹Department of Computer Science, University of Arizona

Email: kobus@cs.arizona.edu

²Department of Electrical and Computer Engineering, University of Arizona

Abstract

Color is of interest to those working in computer vision largely because it is assumed to be helpful for recognition. This assumption has driven much work in color based image indexing, and computational color constancy. However, in many ways, indexing is a poor model for recognition. In this paper we use a recently developed statistical model of recognition which learns to link image region features with words, based on a large unstructured data set. The system is general in that it learns what is recognizable given the data. It also supports a principled testing paradigm which we exploit here to evaluate the use of color. In particular, we look at color space choice, degradation due to illumination change, and dealing with this degradation. We evaluate two general approaches to dealing with this color constancy problem. Specifically we address whether it is better to build color variation due to illumination into a recognition system, or, instead, apply color constancy preprocessing to images before they are processed by the recognition system.

Introduction

Color is of interest to those working in computer vision largely because it is assumed to be helpful for recognition. Color has been studied in detail for specific recognition tasks such as skin.¹⁻³ Color is also possibly the most useful of the features typically used in content based image retrieval (CBIR) systems.⁴⁻⁷ In these systems the usual task is to find images in a database which are similar in appearance to a query image. There has also been much written about the similar endeavor of image indexing as suggestive of object recognition.⁸⁻¹³ This body of work is similar to CBIR in that an exemplar image is used for querying, but here the query image generally is thought to come from a database of single objects, typically on a controlled background which is easily removed by pre-processing.

Conceptually, images which are examined for the presence of the object might contain it amid clutter. However, since the indexing paradigm makes queries based on the properties of an entire image, finding the object

among clutter requires searching over image windows. These considerations are usually deferred, and most work has been simply searching for “object images”. Thus in practice, indexing is only suggestive of object recognition.

Even if clutter is dealt with, there are still other key difficulties. First, while this form of recognition might be able to find a particular, known, multi-colored ball, it does nothing to help find a different multi-colored ball. More to the point, because there is no notion in the training set that these objects are of the same class, there is no way to learn the variance of the colors to be expected in the world of multicolored balls, nor to bootstrap the learning of shape which is essential for a general theory of recognition.

A second problem surfaces when one tries to deal with illumination change. When objects and scenes are imaged under different lights their colors can change dramatically. This presents problems for color based indexing and recognition systems in general, and is a large part of the motivation for the large research effort on computational color constancy. In the case of indexing, the strategy is usually some form of normalization,^{9,11,12} where the database of object images which is searched is normalized for illumination change in a pre-processing step. The same normalization is applied to the object image or window thereof to be matched. For example, with the gray-world method, we assume that the average of any scene is a specific color (“gray”) and deviations from that statistic are due to illumination effects. Assuming the diagonal model of illumination change,¹⁴⁻¹⁷ image colors are then scaled independently so that the overall image is the specific gray. Notice that even if the process does not determine the actual illumination change, it can still work for indexing. The key point is that images are mapped into a different space where they can be matched regardless of whether there is a difference in their imaging conditions.

The problem with normalization is that in order to deal with the illumination variation, useful information is discarded. For example, consider uni-colored objects. Grey world normalization maps them all into the same representations. Thus the system cannot distinguish between a blue ball and a white one. Nonetheless, for multi-colored objects the method can be effective because as the number

of colors increases, the space of possibilities increases rapidly because these methods typically take into account the relative amounts (histogram) of each color.

One alternative is to assume that the objects in the reference dataset are all imaged under the same (canonical) illumination. Color constancy processing is then applied to the image under consideration, but not to any sub-windows tested for the object. If the image is only of a single object (typical in research in this topic), then the situation is not much different than the previous case. However, if the image is a more complex scene (typical in many applications), then color constancy assumptions are likely to hold, and looking for the object in sub-windows might work. In this scenario, it is thus possible to distinguish a blue object from a white one, despite illumination change.

A second alternative¹⁸⁻²¹ is to represent the range of colors possible under expected illumination changes for the objects of interest in the recognition system. This suggests a question we address below. Specifically, is it better to build color variation due to illumination into a recognition system, or to apply color constancy preprocessing to images before they are considered by the recognition system.

In what follows we will review a recently developed model for recognition, and then use the performance of that model to quantify the effects of color spaces for recognition. Next we will consider how the performance of the system degrades with color variation. Then we will evaluate the two strategies for dealing with the degradation: simple color constancy pre-processing, and exposing the training of the model to illumination variation.

Object Recognition as Translation

We adopt a model of object recognition where words must be placed on image regions,²²⁻²⁴ illustrated in Figure 1. This is achieved in practice by exploiting large image data sets with associated text. Critically, we do not require that the text be associated with the image regions, as such data is rare. Considering processes which translate from images (visual representation) to words (semantics) gives a handle on a number of difficult computer vision problems. In part, this is because translation performance can be measured on a large scale, by comparing the proposed translation (predicted words) with the actual translation (associated text). In this work, we use word prediction performance to evaluate the efficacy of color spaces for recognition, as well as the two strategies discussed above for handling illumination variation.

A number of methods have recently been described for predicting words from segmented images.²²⁻²⁵ For the results reported in this paper we use a special case of one of these. Specifically, we model the joint probability of words and image regions as being generated by a collection of nodes, each of which has a probability distribution over both words and regions. The word probabilities are provided by simple frequency tables, and the region probability distributions are Gaussians over feature vectors. We restrict the Gaussians to have diagonal covariance.

Given an image region, its features imply a probability of being generated from each node. These probabilities are then used to weight the nodes for word emission. Thus words are emitted conditioned on image regions. In order to emit words for an entire image (auto-annotation), we simply sum the distributions for the N largest regions. Thus each region is given equal weight, and the image words are forced to be generated through region labeling.

To be consistent with the more general models referenced above, we index the nodes by “levels”, l . Given a region (“blob”), b , and a word w , we have

$$P(w|b) = \sum_l P(w|l)P(b|l)P(l)/P(b) \quad (1)$$

where $P(l)$ is the level prior, $P(w|l)$ is a frequency table, and $P(b|l)$ is a multivariate Gaussian over region features. To estimate the conditional density of words given blobs for the entire image these probabilities are summed over the N largest blobs. In the experiments reported in this paper, N was 8.

The parameters of the model are estimated from the word-blob co-occurrence data using Expectation Maximization.²⁶ In particular, we learn the level priors, $P(l)$, the frequency tables for each level, $P(w|l)$, and the means and the variances of the multivariate Gaussians for computing $P(b|l)$. For all experiments reported in this paper we used 500 nodes.

Experimental Protocol

We used images from 160 CD's from the Corel image data set. Each CD has 100 images on one relatively specific topic such as "aircraft". From the 160 CD's we drew samples of 80 CD's, and these sets were further divided up into training (75%) and test (25%) sets. The images from the remaining CD's formed a more difficult “novel” held out set. Predicting words for these images is difficult, as we can only reasonably expect success on quite generic regions such as “sky” and “water”—everything else is noise.




Figure 1. Illustration of labeling. Each region is labeled with the maximally probable word, but a probability distribution over all words is available for each region.

Each such sample was given to each process under consideration, and evaluated on the basis of at least 1000 images. The results of 10 such samples were further averaged. This controls for both the input data and EM initialization. Words occurring less than 20 times in the training set were excluded. The number of words in the vocabulary varied from 153 to 174 over the 10 runs.

Images were segmented using Normalized Cuts.²⁷ We used a modest selection of features for each region, including size, position, average region color, standard deviation of color over the region, average oriented energy (12 filters), average differential response of 2 different Gaussian filters, and a few simple shape features. The features were chosen to be consistent with recent work on linking words with images,²³⁻²⁵ where we purposely chose to represent color redundantly by (R,G,B), (r,g,S) defined by $S=R+G+B$, $r=R/S$, and $g=G/S$, and L^*a^*b . In the color space evaluation, we restrict the color features to be from only one of these choices. We normalize all features so that in the training data each has mean zero and variance one.

Varying the Illumination

Unfortunately, appropriate large scale data sets with controlled illumination variation are not available. As a compromise, we constructed a semi-synthetic data set as follows. We began with a comprehensive controlled illumination data set.^{28,29} This data set was constructed to be representative of the changes in illumination chromaticity generally encountered. We scaled each pixel in those images so that the overall brightness, $R+G+B$ was the same. We then computed the best, in the least squares sense, 3 by 3 matrices mapping the images under each of the 11 illuminants to one chosen canonical illuminant (Sylvania 50MR16Q). (One of these matrices is the identity). We then used these 11 matrices to simulate illumination changes in the Corel data set. For each image, we removed the gamma correction, scaled the overall brightness of each pixels so that all $R+G+B$ were the same, applied one of the 11 matrices to the (R,G,B), and then re-scaled the (R,G,B) so that it was set to the same value in the original image. We then computed images features for the images based on the new (R,G,B) values. This process produced some (R,G,B) values which were above the usual maximum value of 255. When color constancy processing was applied to such images values over 255 were made available to that process, all pixel values were truncated to 255 before being used for recognition experiments.

Our simulation of illumination change is only a gross approximation of what would occur if the illumination striking the scene underwent analogous changes. For example, the process makes no sense for sources, such as the sky. However, the procedure is more justified if we think of the “scenes” as being prints of the images, not the scenes themselves. The key point is that we want to capture the distribution of colors in a real data set which is also appropriate for large scale recognition experiments.

Performance Measures

Several ways to quantify word prediction performance have been proposed²⁴ Here we use the simplest measure. Specifically, we allow the model to predict M words, where M is the number of words available for the given test image. In our data M varies from 1 to 5. The number correct divided by M is the score.

We express word prediction relative to that for the empirical word distribution—i.e., the frequency table for the words in the training set. This reduces variance due to varied test sample difficulty. Exceeding the empirical density performance is required to demonstrate non-trivial learning. Doing substantially better than this on the Corel data is difficult. The annotators typically provide several common words (e.g. “sky”, “water”, “people”), and fewer less common words (e.g. “tiger”). This means that annotating all images with, say, “sky”, “water”, and “people” is quite a successful strategy. Performance using the empirical word frequency would be reduced if the empirical density was flatter. Thus for this data set, the increment of performance over the empirical density is a sensible indicator.

Color Space Evaluation

Color space choice is often difficult. Clearly, the choice should reflect the application. One issue is the degree to which the three values are correlated. For example, in natural images, R , G , and B , tend to be quite correlated because variation in illumination intensity and direction (shading) tend to effect the three channels similarly. (r,g,S), where $S=R+G+B$, $r=R/S$, and $g=G/S$ is less correlated. R , G , and B can be further decorrelated using PCA and ICA.

A second issue is the degree to which the color space aligns with human perception. In computer vision, L^*a^*b is often used where the connection to human vision is weak. However, one could make a generic argument that the human vision system has evolved to accomplish tasks like the ones we are interested in, and that emulating it where possible makes sense.

Since we have taken care to develop a comprehensive test strategy, we can evaluate which color space is best for our approach. Further, since our system focuses on the canonical computer vision task—linking image features with semantics—it is likely that our findings apply to other systems as well.

We consider adding color as encoded in three different ways—straight RGB, L^*a^*b , and chromaticity with brightness, specifically, $S=R+G+B$, $r=R/S$, and $g=G/S$, in addition to using them all as in the original work.²³⁻²⁵ In all cases we used both the average color and its variance over the region, and we kept the number of features the same by duplicating the chosen color features appropriately. Word prediction performance using each color space is reported in Table 1.

We found that using either L^*a^*b or (r,g,S) is substantially better than using straight RGB. The difference between using L^*a^*b and (r,g,S) was negligible. These results suggest that for our task, it is helpful to decorrelate

brightness and chromaticity, but beyond this step, color space may not be very important.

The Effect of Illumination Variation

To investigate the effects of illumination variation on our recognition system, we trained the models using images from the original Corel images, but tested on the images which had a simulated illumination change set by any of the 11 possible ones in roughly equal proportions.

The results (Table 2, row 2) show that for this application, the range of illumination expected in natural images causes substantial degradation in performance. This is expected as color is an important cue for our system.

Training with Illumination Variation

Studying incorporating illumination variation into the training of the recognition system leads to an important design choice. It could be argued that the training set should consist of every training image from the previous experiment, but under each of the 11 illuminants, making the training set 11 times large. However, we would likely now require a larger model. Thus to avoid this confound, and to match the processing costs and model size with the other experiments, we trained the models on exactly the same number of images as before, and each image was subjected to one of the 11 illumination changes. Each of the 11 illumination changes received roughly equal representation. This should not be an overly large burden, as the system learns from multiple examples—and now it simply sees more color variation in those examples. Recall from (1) that the variance of the feature is part of what is learnt by the system.

The results (Table 2, row 3) show that exposing the training process to the expected illumination variation is helpful, reducing the negative effect of varying illumination by about 40% in the case of first held out set, and 60% in the novel held out set.

Color Constancy Pre-Processing

As discussed in the introduction, the obvious, and often assumed solution to the illumination variation problem in object recognition is color constancy pre-processing. For this work we test this idea with the two simple color constancy methods: gray-world (GW) and scale-by-max (SBM). Many better methods exist (see, for example^{16,30-37}), but here we are more interested in first establishing whether color constancy processing helps at all. For the gray world method we computed the appropriate expected value of the average (R,G,B) over all 34,000 Corel images. (For this data set, gray is (52.9, 51.0, 43.0)). We then removed the color cast from the images by assuming that the average (R,G,B) for each image was gray, and that the diagonal model was valid. For the Corel images, the diagonal model is not likely to be particularly good, but for our simulation experiments where the modeled illumination change was set to emulate

the SFU data, the diagonal model is reasonably accurate. However, we purposely allowed for some variation from the diagonal model by using the 3 by 3 linear transformations on chromaticity only to create the data set.

Table 1. Word prediction performance for the most common color spaces in computer vision. The numbers are amount by which word prediction exceeds that of using the empirical distribution (bigger is better).

| Feature set | Word prediction performance on the various data sets (error is roughly 0.003) | | |
|---------------------|---|----------|-------|
| | Training | Held out | Novel |
| RGB, L*a*b, and rgS | 0.140 | 0.090 | 0.055 |
| RGB | 0.112 | 0.064 | 0.044 |
| L*a*b | 0.148 | 0.096 | 0.059 |
| rgS | 0.149 | 0.094 | 0.060 |

For the scale-by-max method, we simple scale each channel so that the maximum in the images is that observed in the entire dataset, which, for the Corel data set, is not surprisingly 255 for each channel.

The results (Table 2, rows 4 and 5) show that the scale-by-max normalization is very helpful, whereas the gray-world normalization is not. An examination of the images reveals that the color balance of many or most of them is consistent with the maximum in each channel being close to 255. There are obvious exceptions, such as the entire CD of sunsets, but each CD makes up less than 1% of our data. By contrast, the gray world assumption is not particularly valid for this data set, and attempting to deal with illumination change by exploiting it did not yield good results.

Color Normalization

In our final experiment, we applied the same normalization, either GW or SBM, to the training data as well as the test data. This scenario is thus similar to the indexing paradigm discussed in the introduction. There we argued that indexing may not make sense if the reference data set is simple objects, because normalization removes too many degrees of freedom.

Notice, however, that in our approach, we are neither using the training images as objects to be recognized, nor images to be found. Rather we are using them to learn about image regions from images which typically have a wide range of colors. Thus training in a normalized space might make sense if illumination variation is expected and this is what is suggested by the results. Using this strategy improves upon that possible using color constancy processing for the test images only. In the case of SBM, the absolute improvement is small because the result obtained without normalizing the training set was already good. In the case of grey world normalization, the improvement was substantial. This makes sense because it in effect alters the data so that the gray world assumption is more valid. However, the performance is still below that of using scale-by-max both with and without extending the normalization to the test data.

Table 2. The effect of illumination change and subsequent processing to deal with it on word prediction performance. The numbers are amount by which word prediction exceeds that of using the empirical distribution (bigger is better). The held out test set was composed of images hidden from the training process but from the same Corel CD's as the training data. The novel test set was composed of images from CD's different from those used in training. Errors were estimated based on the variance of the 10 samples taken. The results confirm our expectation that the range of color variation from typical illumination variation significantly degrade our recognition system where color is an important cue, and that the right color constancy processing can help. In this data set, the conditions for scale-by-max are good, and it is clearly better than the gray world method. Further, if it makes sense for the application, applying color constancy to the training data (bottom two rows) can improve performance even further. This is the "normalization" strategy discussed in the introduction. As explained further in the text, this might not make sense when the training images are meant to be objects, but in our application, where object labels are learnt from images, the strategy appears helpful.

| Experiment | Word prediction performance on the various data set (error estimates are shown in parentheses) | | |
|--|---|---------------|---------------|
| | Training | Held out | Novel |
| No illumination variation | 0.140 (0.003) | 0.090 (0.002) | 0.055 (0.005) |
| Train with no illumination variation and test with illumination variation | 0.092 (0.0025) | 0.060 (0.002) | 0.030 (0.004) |
| Train and test with illumination variation | 0.121 (0.003) | 0.072 (0.002) | 0.045 (0.005) |
| Train with no illumination variation and test with illumination variation and GW color constancy pre-processing | 0.062 (0.003) | 0.038 (0.003) | 0.039 (0.003) |
| Train with no illumination variation and test with illumination variation and SBM color constancy pre-processing | 0.122 (0.003) | 0.082 (0.003) | 0.053 (0.004) |
| Train with no illumination variation and GW normalization and test with illumination variation and GW color constancy pre-processing | 0.121 (0.003) | 0.073 (0.002) | 0.053 (0.004) |
| Train with no illumination variation and SBM normalization and test with illumination variation and SBM color constancy pre-processing | 0.135 (0.002) | 0.086 (0.002) | 0.059 (0.003) |

Conclusion

We posit that our system for translating image regions to words is more representative of the general object recognition problem than the often used indexing task, and thus it is a good platform to study the use of color for object recognition. Using this platform we have confirmed the notion that illumination variation can pose problems for object recognition systems, and have looked at several classes of approaches for dealing with it.

In our somewhat artificial dataset, conditions were good for the scale by max algorithm, and using it gave results approaching that where there was no color constancy problem. This is encouraging, because, as argued in the introduction, color constancy is required when objects to be identified have limited color ranges. In this case, simple normalization does not work. When normalization is appropriate, as was the case in our artificial test setting, our results gave a small improvement in conjunction with scale by max, and a large improvement in conjunction with the gray world method. However, the performance with the grey world method was still less than the scale by max method, without normalization.

References

1. N. Tsumura, H. Haneishi, and Y. Miyake, Independent component analysis of skin color image, IS&T/SID Sixth Color Imaging Conference: Color Science, Systems and Applications, The Society for Imaging Science and Technology, Springfield, Va., 177-184 (1998).
2. J. Yang, W. Lu, and A. Waibel, Skin-color modeling and adaptation, Third Asian Conference on Computer Vision, II:687-694 (1998).
3. M. J. Jones and J. M. Rehg, Statistical color models with application to skin detection, Int. J. of Comp. Vis., 46, 81-96 (2002).
4. 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, 28, 22-32 (1995).
5. W. Y. Ma and B. S. Manjunath, NeTra: A toolbox for navigating large image databases., Multimedia Systems, 7, 84-198 (1999).
6. J. Z. Wang, J. Li, and G. Wiederhold, SIMPLIcity: Semantics-Sensitive Integrated Matching for Picture Libraries, IEEE Trans. Patt. Anal. Mach. Intell., 23, 947-963 (2001).
7. C. Carson, S. Belongie, H. Greenspan, and J. Malik, Blobworld: Color and Texture-Based Image Segmentation Using EM and Its Application to Image Querying and

- Classification, *IEEE Trans. Patt. Anal. Mach. Intell.*, 24, 1026-1038 (2002).
8. M. J. Swain and D. H. Ballard, Color Indexing, *Int. J. of Comp. Vis.*, 7, 11-32 (1991).
 9. B. V. Funt and G. D. Finlayson, Color Constant Color Indexing, *IEEE Trans. Patt. Anal. Mach. Intell.*, 17, 522-529 (1995).
 10. M. S. Drew, J. Wei, and Z.-N. Li, Illumination-Invariant Color Object recognition via Compressed Chromaticity Histograms of Normalized Images, Simon Fraser University School of Computing Science Technical Report, CMPT-TR-97-09 (1997).
 11. B. Funt, K. Barnard, and L. Martin, Is Colour Constancy Good Enough?, 5th European Conference on Computer Vision, Springer, I:445-459 (1998).
 12. G. D. Finlayson, B. Schiele, and J. L. Crowley, Comprehensive colour image normalization, 5th European Conference on Computer Vision, Springer, I:475-490 (1998).
 13. G. D. Finlayson, S. S. Chatterjee, and B. V. Funt, Color Angular Indexing, 4th European Conference on Computer Vision, Springer, II:16-27 (1996).
 14. J. A. Worthey and M. H. Brill, Heuristic analysis of von Kries color constancy, *Journal of the Optical Society of America A*, 3, 1708-1712 (1986).
 15. G. West and M. H. Brill, Necessary and sufficient conditions for von Kries chromatic adaptation to give colour constancy, *J. Math. Biol.*, 15, 249-258 (1982).
 16. D. Forsyth, A novel algorithm for color constancy, *Int. J. of Comp. Vis.*, 5, 5-36 (1990).
 17. G. D. Finlayson, Coefficient Color Constancy, Simon Fraser University, School of Computing, Ph.D. thesis, (1995).
 18. J. Matas, R. Marik, and J. Kittler, On representation and matching of multi-coloured objects, Fifth International Conference on Computer Vision, 726-732 (1995).
 19. J. Matas, R. Marik, and J. Kittler, Illumination Invariant Colour Recognition, 5th British Machine Vision Conference (1994).
 20. J. Matas, Colour-based Object Recognition, University of Surrey, Ph.D. thesis, (1996).
 21. J. G. Matas, D. Koubaroulis, and J. Kittler, Colour Image Retrieval and Object Recognition Using the Multimodal Neighbourhood Signature, 6th European Conference on Computer Vision, Springer, 48-64 (2000).
 22. K. Barnard and D. Forsyth, Learning the Semantics of Words and Pictures, International Conference on Computer Vision, II:408-415 (2001).
 23. P. Duygulu, K. Barnard, J. F. G. D. Freitas, and D. A. Forsyth, Object recognition as machine translation: Learning a lexicon for a fixed image vocabulary, The Seventh European Conference on Computer Vision, IV:97-112 (2002).
 24. K. Barnard, P. Duygulu, N. D. Freitas, D. Forsyth, D. Blei, and M. I. Jordan, Matching Words and Pictures, *Journal of Machine Learning Research*, 3, 1107-1135 (2003).
 25. K. Barnard, P. Duygulu, and D. Forsyth, Clustering Art, *IEEE Conference on Computer Vision and Pattern Recognition*, II:434-441 (2001).
 26. A. P. Dempster, N. M. Laird, and D. B. Rubin, Maximum likelihood from incomplete data via the EM algorithm, *Journal of the Royal Statistical Society. Series B (Methodological)*, 39, 1-38 (1977).
 27. J. Shi and J. Malik, Normalized Cuts and Image Segmentation, *IEEE Trans. Patt. Anal. Mach. Intell.*, 22, 888-905 (2000).
 28. K. Barnard, L. Martin, B. Funt, and A. Coath, A Data Set for Color Research, *Color Research and Application*, 27, 147-151 (2002).
 29. K. Barnard, Lindsay Martin, B. Funt, and A. Coath, Data for computer vision and colour research, available from <http://kobus.ca/research/data/index.html>.
 30. B. Funt, V. Cardei, and K. Barnard, Learning Color Constancy, *IS&T/SID Fourth Color Imaging Conference: Color Science, Systems and Applications*, The Society for Imaging Science and Technology, Springfield, Va., 58-60 (1996).
 31. G. Finlayson, S. Hordley, and P. Hubel, Color by Correlation: A Simple, Unifying Framework for Colour Constancy, *IEEE Trans. Patt. Anal. Mach. Intell.*, 23, 1209-1221 (2001).
 32. K. Barnard, V. Cardei, and B. Funt, A Comparison of Computational Colour Constancy Algorithms. Part One. Methodology and Experiments with Synthesized Data, *IEEE Transactions on Image Processing*, 11, 972-984 (2002).
 33. K. Barnard, L. Martin, A. Coath, and B. Funt, A Comparison of Computational Colour Constancy Algorithms. Part Two. Experiments on Image Data, *IEEE Transactions on Image Processing*, 11, 985-996 (2002).
 34. K. Barnard, Improvements to Gamut Mapping Colour Constancy Algorithms, 6th European Conference on Computer Vision, Springer, 390-402 (2000).
 35. K. Barnard, L. Martin, and B. Funt, Colour by correlation in a three dimensional colour space, 6th European Conference on Computer Vision, Springer, 375-389 (2000).
 36. S. Tominaga and B. A. Wandell, Natural scene-illuminant estimation using the sensor correlation, *Proceedings of the IEEE.*, 90, 42-56 (2002).
 37. S. Tominaga, S. Ebisui, and B. A. Wandell, Scene illuminant classification: brighter is better, *J. Opt. Soc. Am. A*, 18, 55-64 (2001).

Biographies

Kobus Barnard is an assistant professor at the University of Arizona, Tucson. Before arriving at U of A, he was a post doctoral fellow in computer vision at the University of California at Berkeley. He received his Ph.D. in computer science from Simon Fraser University, where he specialized in colour constancy. His current research interests include object recognition, image databases, and colour issues in vision and image reproduction.

Prasad Gabbur is a graduate student pursuing Master of Science in electrical and computer engineering at the University of Arizona, Tucson. Earlier, he received his Bachelor of Engineering degree from the Electronics and Communications department of the National Institute of Technology, Karnataka (NITK), India. He was a Young Engineering research fellow at the Indian Institute of Science in the field of biometrics. He has also worked as a software developer for Eyematic Interfaces Inc., Los Angeles, California. His current research interests are in the fields of computer vision/pattern recognition, color constancy for object recognition, perceptual grouping and biometrics. Email: pgsangam@ece.arizona.edu