Spotting Colours

Georgina Kwei, Kobus Barnard, and Brian Funt

Department of Computing Science,
Simon Fraser University,
8888 University Drive, Vancouver, BC, Canada, V5A 1S6
email: gkwei@ea.com, {kobus, funt}@cs.sfu.ca
phone: 604-291-3126; fax: 604-291-4424

Abstract

Spotting an object of a chosen colour, skin for example, in an image taken under uncontrolled lighting conditions requires some form of preprocessing to convert the image colours to an illumination-independent specification. Computational colour constancy algorithms are designed specifically to accomplish this task. We process images with different colour constancy algorithms and evaluate the results in terms of the resulting accuracy of object colour spotting. Since spotting a colour depends on the threshold used to define the colour, ROC analysis is used in evaluating the results. The experiments help to evaluate the performance of the colour constancy algorithms and to predict what can be expected in terms of being able to find particular colours in a cluttered environment.

Key Words. colour, object recognition, computational colour constancy

1. Introduction

This paper describes tests to evaluate the accuracy with which objects of a specific single colour can be identified in a complex image, and the effectiveness of the addition of colour constancy preprocessing. Although it is already known how well colour constancy algorithms perform in terms of RMS error[4], and there are numerous papers on colour-based object recognition [2][3][4], those works use and test objects that have more than one colour. In Yuille, Snow and Nitzberg [3], they perform the task of locating street signs using the two colours on the signs. Algorithms by Funt and Finlayson [4] and Ennesser and Medioni [2] use colour-edge histograms to do multi-coloured object recognition. This paper describes the results of spotting one distinct, and more or less uniform, colour. For example, a lemon is an object that has the distinct uniform colour of yellow. Locating human skin in an image by looking for flesh tones is another example.

Spotting of objects is the representative task used to measure the accuracy of the different colour constancy algorithms. The colour constancy algorithms used are 2D neural networks, greyworld, database greyworld and white-patch retinex. The control

Reprinted from: *Proceedings Colour Image Science 2000*, Univ. of Derby, Derby England, April 2000

image used is the image taken under the known illuminant with no colour constancy pre-processing.

2. The Test Images

The test images are taken from the object recognition image database used in Funt, Barnard and Martin [5] and have been adjusted against the camera calibration model described in that paper. In particular, this paper uses the images:

- Ball 1, which is a fuzzy, partially fluorescent, ball
- Ball 2, which is a shiny multicoloured ball.
- Tide, which is a detergent box.

Ball1 is fluorescent in the sense that it glows brightly under ultraviolet light.

These particular objects were chosen because they have fairly large areas of uniform and distinct colour that can be used for the spotting tests. From each image, one colour was segmented from the object and represented the **single-colour object** to be spotted in the later tests.

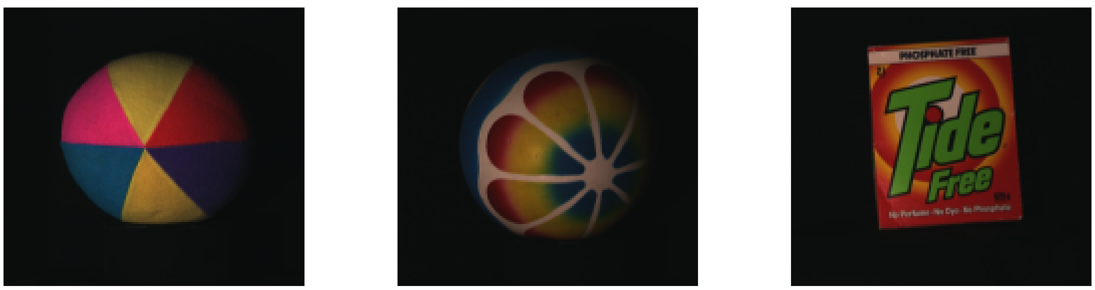

Ball1 Ball2 Tide

Figure 1. Test objects

In the Ball1 image, the blue and pink areas were each used separately as the single-colour object. In other words the goal is not to identify Ball1, but rather, to identify the pink region, for example, in various images of Ball1 taken under different illuminants. From Ball2, the white area, and from the tide image, the green lettering, were chosen as the single-colour object. This resulted in 4 "objects" for which to spot.

Each object in the database was placed in a Macbeth Judge II light booth, and an image was captured of the object under five different illuminants. The camera aperture was adjusted to prevent image clipping under the brightest lighting conditions for a given object, and not changed between illuminants. The five illuminants were:

Illuminant Name	Description
mb-5000	Macbeth 5000 K fluorescent tubes
mb-5000+3202	Macbeth 5000 tubes with Roscolux #3202 Full Blue filter
syl-cwf	Sylvania Cool White Fluorescent tube
Ph-ulm	Phillips Ultralume fluorescent tube
Halogen	Sylvania 75W halogen bulb

Images were taken of each object under each of the five illuminants. A copy of each image was left unprocessed as a control image and then each image was processed using each of the four colour constancy algorithms to generate 25 images. A total of 100 images were used.

3. Method

3.1 Colour Constancy Processing

The images stored in the calibrated database had been averaged over numerous frames to reduce noise and increase dynamic range. As well, they had been linearized and adjusted for camera black offset. The image data is stored in floating point to preserve the accuracy of the frame averaging.

The greyworld algorithm estimates the illumination chromaticity by comparing the average of the image chromaticities to an ideal grey defined as R=G=B. The database greyworld compares the image average to the average of the chromaticities taken over all images in the database. The “white-patch retinex” algorithm estimates the illumination chromaticity as the chromaticity obtained from the maximum values in each of the R, G, and B channels. The neural network method is based on a 2-layer feed-forward network that takes a binarized and coarsely sampled chromaticity histogram as input and outputs a chromaticity pair.

3.2 Spotting

The spotting algorithm requires as input a target spotting colour (the colour being searched for in an image) and a test image in which to look for the colour. Intensity is not considered relevant, so the spotting algorithm compares colours in standard chromaticity space:

$$r = R/(R+G+B) \text{ and } g = G/(R+G+B) \tag{1}$$

There are other ways of normalizing the RGB to eliminate the effects of intensity, such as $r = R/B$ and $r = R/\mathrm{sqrt}(R^2+G^2+B^2)$; however, from our tests formula (1) yields the most favourable results. That is, there was a higher degree of accuracy when this normalizing method for RGB was used instead of the others.

Before converting to chromaticity space, the test image was averaged using a 5x5 block of pixels. By smoothing the image, the individual pixels were more representative of the colour of the object and were less susceptible to noise.

The actual spotting task required determining whether the target colour existed in the test image. This was done by computing the Euclidean distance between each pixel of the target image and the target spotting colour. If the distance for the pixel and the target spotting colour is below some tolerance, then that pixel is considered a hit. The tolerance is defined as a percentage of the maximum value of a pixel in the standard chromaticity space.

Of course, the number of hits depends on the tolerance, and as the tolerance is increased, there will be more and more pixels incorrectly identified as being of the target colour. Therefore, the spotting performance is analyzed using receiver operating characteristic (ROC) curves. The true-positive (TP) proportion is plotted against the false-positive (FP) proportion for various possible thresholds. The TP proportion is defined as:

$$\frac{\text{Pixels identified as hits}}{\text{Pixels in ideal region}}$$

The FP proportion is defined as:

$$\frac{\text{Pixels identified as hits outside the ideal region}}{\text{Pixels outside the ideal region}}$$

Each of the nodes in the ROC curve represents a particular tolerance value. The tolerances used were 1%, 5%, 10%, 15%…50%.

To determine the target spotting colour, the average colour value of a patch of the single-colour object was calculated. In the case of Ball1, this was a 25-by-25 pixel patch.

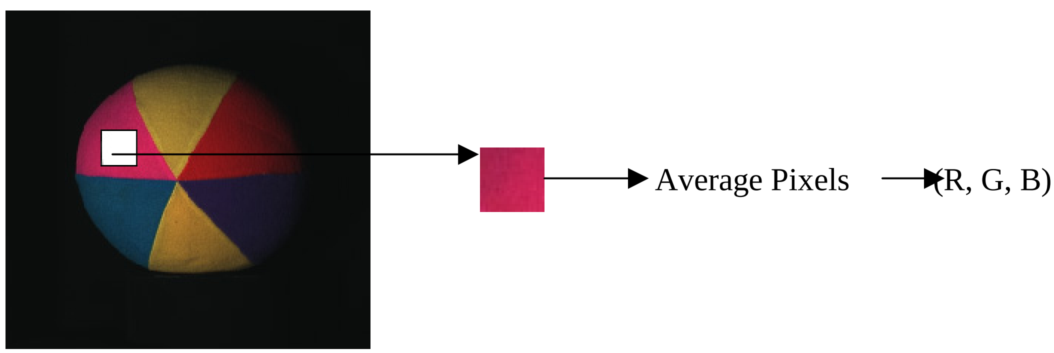

Figure 2. Determining the target spotting colour

Although, the pixels of the single-colour object vary slightly from one another, the averaging of a patch on the object helps to select a representative target spotting colour. Still, one concern was that due to shading variations across the image, the patch chosen would be a poor representation of the overall colour of the single-colour object. This could potentially have an adverse effect on the performance of the spotting algorithm. Thus, for each of the single-colour objects, we took five sample patches at different locations within the single-colour object to analyze the variation between the chromaticity values.

Table 1 illustrates the chromaticity values of each of the samples and their standard deviations.

		Sample 1	Sample 2	Sample 3	Sample 4	Sample 5		
Ball1 – Pink area	**r**	0.614492	0.605147	0.620408	0.609581	0.60995	**Stn dev:**	**0.005785**
	g	0.11729	0.120839	0.115259	0.119396	0.1194	**Stn dev:**	**0.002181**
Ball1 –Blue area	**r**	0.121209	0.120963	0.122714	0.122982	0.121192	**Stn dev:**	**0.000955**
	g	0.372324	0.371319	0.371341	0.372665	0.37118	**Stn dev:**	**0.000679**
Ball2 – White area	**r**	0.412187	0.413473	0.403474	0.411394	0.413821	**Stn dev:**	**0.004248**
	g	0.319539	0.315366	0.326269	0.317707	0.315363	**Stn dev:**	**0.004503**
Tide – Green area	**r**	0.315499	0.314403	0.323186	0.314637	0.318649	**Stn dev:**	**0.003714**
	g	0.542787	0.53366	0.525901	0.532917	0.535618	**Stn dev:**	**0.006053**

Table 1. Standard deviation of target spotting colours

As shown in the table, the standard deviation of chromaticity components of the sample patches is between 0.0006 and 0.006. The variation is small enough that a random patch of the single-colour object will yield similar results to any other randomly selected patch.

After determining the target spotting colour for each of the single-colour objects, we defined the “ideal” answer region within the image. In other words, this is the region of the image that we considered to be the correct answer in that all of its pixels appeared to be the target colour. For example, Figure 3 (right) shows the region considered to be the correct answer when spotting for the pink colour. Pixels identified within this region are considered correct hits; those outside as incorrect hits.

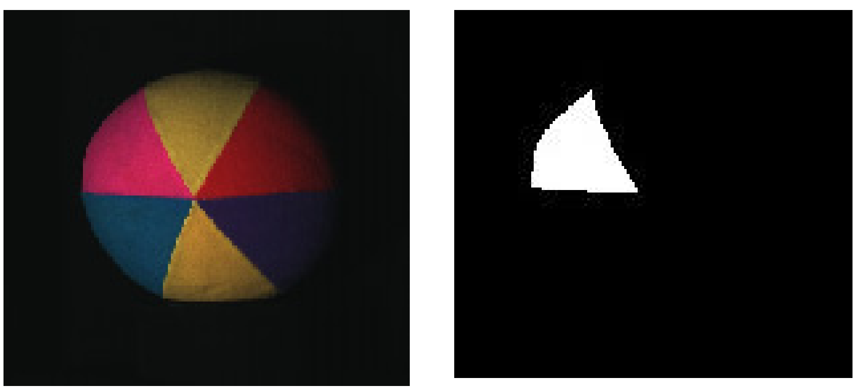

Figure 3. Ideal region of pink single-colour object in Ball 1 image

4. Results

An analysis of the performance of each of the different colour constancy algorithms under different illumination conditions was done. We plotted, but do not show them here, the ROC curves for each illuminant-object pair. For all the objects, except the blue in the ball1 image, the Macbeth 5000 with a blue filter appeared to cause the most difficulty for the colour constancy algorithms.

We then combined the ROC curves for each algorithm to form the charts shown in Figures 4, 5, 6 and 7. Each ROC curve in the chart represents the combined results of a colour constancy algorithm under different illuminants. Each plotted point on a curve represents an increase in the tolerance value over the preceding point on the curve.

For all the objects, colour constancy pre-processing appeared to give better results than no pre-processing at all. In fact, for the blue and pink objects in the ball1 image and the white object in the ball2 image, colour constancy pre-processing helped the spotting algorithm achieve a TP proportion greater than 90% at a 50% tolerance. While with no pre-processing, the spotting algorithm achieved a TP proportion less than 80% for the blue object and less than 60% for the pink object at a 50% tolerance.

Two of the four single-colour objects have a significant fluorescent component. These objects are the pink in the ball1 image and the green in the Tide image. Fluorescence of this sort is very common, especially in clothing laundered with detergents with added ‘whiteners.’ It is known that fluorescent surfaces can adversely affect the performance of colour constancy algorithms [6]. The effect of fluorescence on the colour constancy algorithms may explain why the 2D neural net and retinex algorithms mainly perform worse than no pre-processing at all on the green object in the Tide image. On the pink object in the ball1 image, the retinex, greyworld, and database greyworld algorithms achieve a superior number of TP but also have a larger number of FP than if no pre-processing were used.

The maximum tolerance of 50% was sufficient for all the colour constancy algorithms to achieve a TP proportion approaching 1.00 (100%). This appears to be a sufficient maximum value and if we require a tolerance greater than 50% when spotting for a colour, this may mean that the colour of the object we are spotting for is not sufficiently differentiated from the colours of its environment.

Figure 4. ROC curves for Ball1 - blue

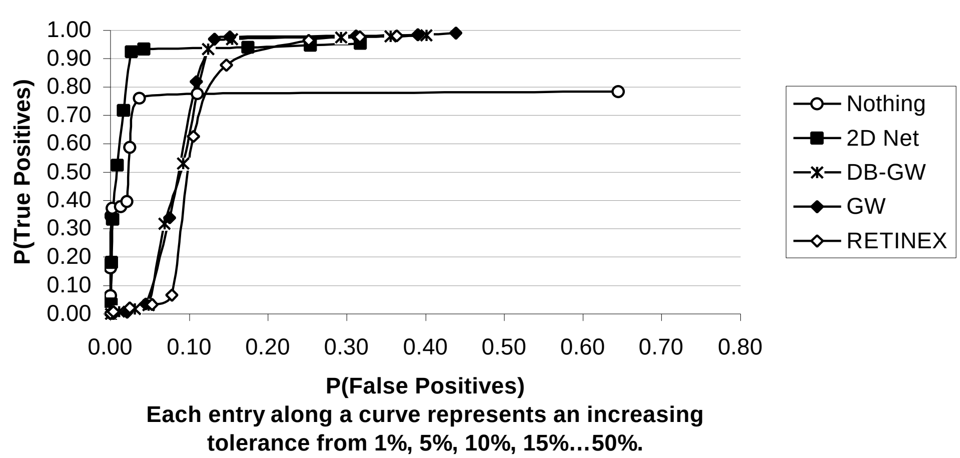

The range of FP proportions of the ROC curves is most likely a good indicator of the number of distracters in the scene. In our test cases, the range of FP values varied widely between the different scenes. At a tolerance of 50%, the blue object in the Ball1 image (Figure 4) and the white object in the Ball2 image (Figure 6) had a FP proportion of approximately 70% and 90%, respectively. In contrast, the pink object

Figure 5. ROC curves for Ball1 - pink

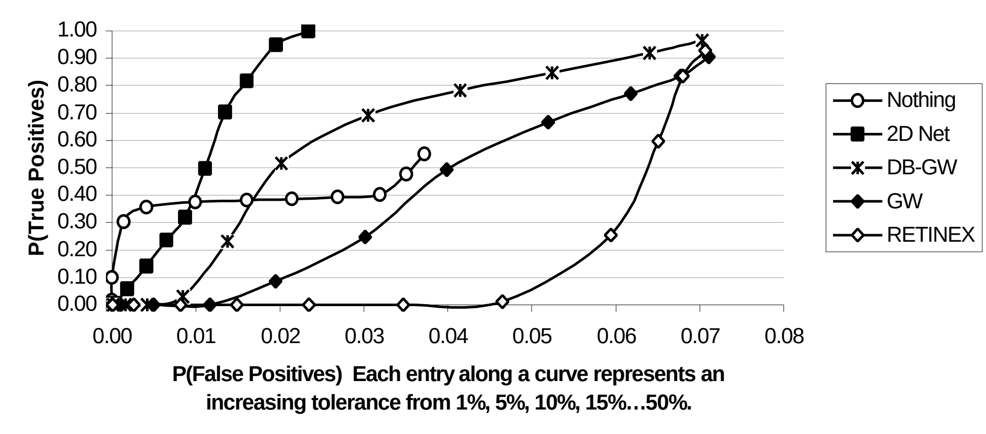

in the ball1 image (Figure 5) had a FP proportion around 7.5% and the green in the Tide image had a FP proportion around 3.5% at a tolerance of 50%.

Spotting was also performed on images taken under the canonical illuminant but which were pre-processed using a colour constancy algorithm. Figure 8 shows the average ROC analysis results from each of the four images. As expected, the algorithms adversely affected the spotting performance since the canonical contains

Figure 6. ROC curves for Ball2

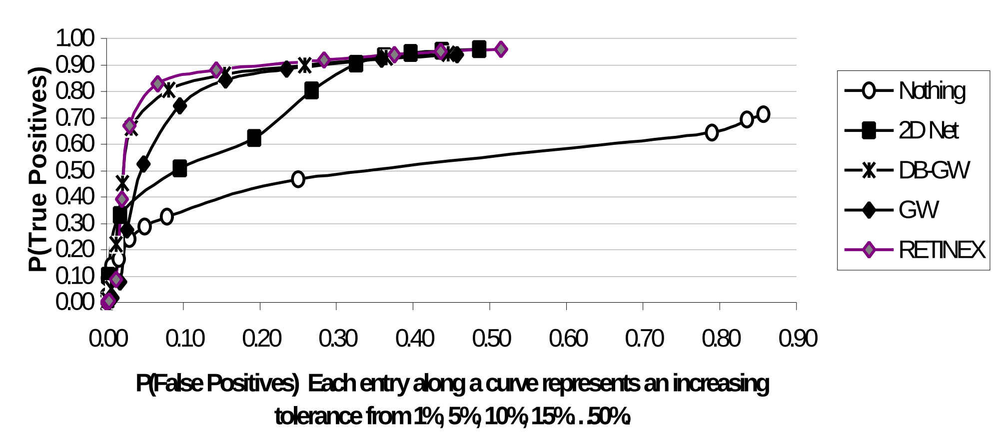

the correct colours to begin with.

Figure 7. ROC curves for Tide

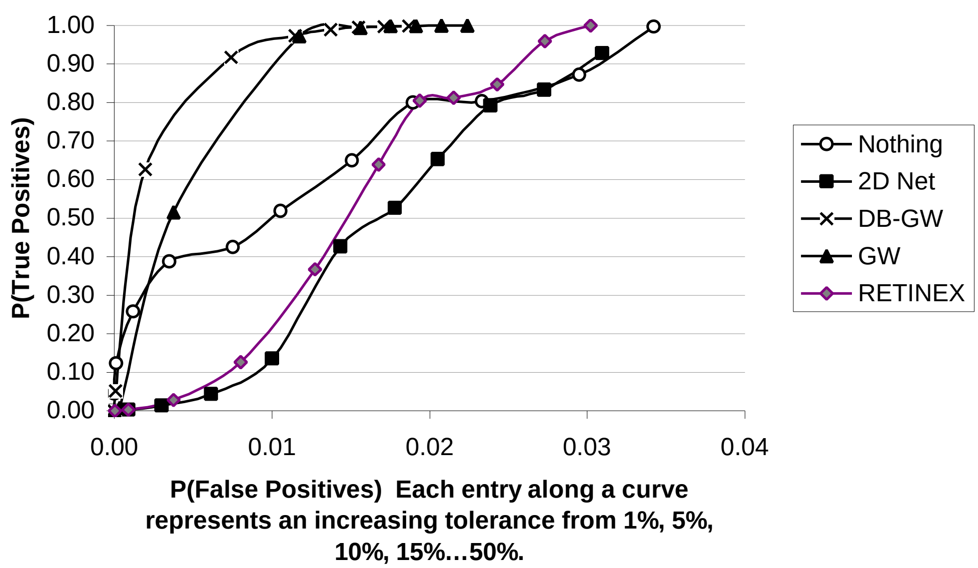

Figure 8. Average ROC curves

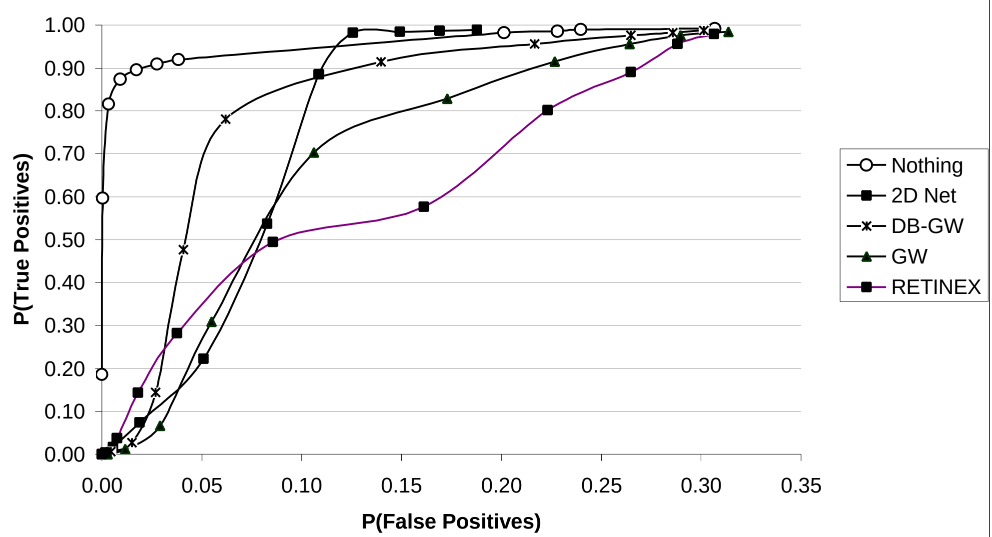

Figure 8 shows the receiver operator characteristic curve for the case of spotting for a colour in the original image with colour constancy processing. All algorithms degrade the performance somewhat. The 'nothing' case in which the image is left untouched still does not yield perfect spotting results because of the presence of similarly coloured pixels in other parts of the image. Also because the spotting algorithm performs some image smoothing it is comparing 2 very similar, but not identical, images.

5. Conclusion

Machine colour constancy algorithms were tested to find out how well they perform when applied as a pre-processing step for the task of spotting colours. In particular, the task focused on the location of an object with one distinct and uniform colour. As one would hope, the colour constancy algorithms generally improved the spotting performance and usually resulted in a higher number of true positives (hits) and a reasonable number of false positives. In addition, the presence of fluorescent surfaces also appeared to have an effect on the success of the colour constancy pre-processing. The results show that colour constancy pre-processing appears to be beneficial for the task of spotting colours and can greatly improve the results compared to when no pre-processing is done. However, the algorithms used for pre-processing are not guaranteed to improve the performance in all cases. A colour constancy algorithm may perform really well when the spotting task is for a particular colour, but the same algorithm may perform poorly when spotting for another colour.

6. References

[1] B. Funt, V. Cardei, and K. Barnard. "Learning Color Constancy", Proc. IS&T/SID Fourth Color Imaging Conf., pp. 58-60, Scottsdale, Nov. 19-22, 1996

[2] F. Ennesser and G. Medioni, "Finding Waldo, or Focus of Attention Using Local Color Information", IEEE Trans. On Pattern Anal. And Machine Int., pp. 805-809, August 1995

[3] A.L. Yuille, D. Snow, and M. Nitzberg, "Signfinder: Using Color to Detect, Localize and Identify Informational Signs," ICCV 98, pp. 629-633

[4] B. Funt and G. Finlayson, "Color Constant Color Indexing", IEEE, 1995 pp.522-529

[5] B. Funt, K. Barnard and L. Martin, "Is Machine Colour Constancy Good Enough?"
5th European Conference on Computer Vision, 1998, pages 445-459.

- Image database: http://www.cs.sfu.ca/~colour/image_db/

[6] K. Barnard, "Colour Constancy with Fluorescent Surfaces," Proc. IS&T/SID Seventh Color Imaging Conference, Scottsdale, Nov. 1999.

[7] M. Swain and D. Ballard, "Color Indexing," Int. J. Comp Vision, 7:11-32, 1991