

LEARNED COLOR CONSTANCY FROM LOCAL CORRESPONDENCES

Tijmen Moerland* and Frédéric Jurie

LEAR - INRIA/CNRS, 655 avenue de l'Europe, Montbonnot, 38334 Saint Ismier Cedex, France

ABSTRACT

The ability of humans for color constancy, i.e. the ability to correct for color deviation caused by a different illumination, is far beyond computer vision performances: nowadays, automatic color constancy is still a difficult problem. This article proposes a new step forward towards solving this color constancy problem. Basically, it consists in learning how illumination can affect some reference objects. During a learning stage, images are taken under various illuminations, allowing for automatic building of a model explaining color changes. The model can explain complex non-linear color transformations with only a few parameters. Therefore, the observation of color variations in a few reference regions (e.g. known object) is enough to estimate the global color changes.

1. INTRODUCTION

Color is not an intrinsic property of objects. Illumination change can severely modify colors acquired by a camera and can make the use of color information for subsequent tasks difficult. The active work done by our brain to assign color to an observe scene is difficult to reproduce, even with complex computer vision techniques.

The color constancy problem can be expressed as that of deriving an image of the scene as it would appear under a canonical illuminant, given the image of this scene under an unknown illuminant. The mapping has to account for the change in relative spectral power distribution between the unknown and canonical illuminant.

Usually, two different ways to investigate color constancy are distinguished:

- the first approach is based on psychophysical investigations [1]. The problem with this kind of approach is that conclusions are hard to draw due to the complex experimental environment necessary to examine color constancy.
- the alternative approach tries to understand the principles involved with the spectral image formation [2,3]. The knowledge of physical laws allows to build - for example - illumination independent color ratios

[4,5]. These color invariants unfortunately have several drawbacks, since these transformations can be singular for some sensor values or unstable at others [3].

A third - and new - way to address this problem is as a machine learning problem. By observing how colors are mapped between two images of the same scene under various "real world" lighting changes, it is possible to automatically build a model of possible color changes. First investigations in this area were proposed by Miller and Tieu [6]. They proposed the *color eigenflows model* that statistically learns joint color changes. However, this promising approach suffers from a major drawback: color constancy is possible only if the image of the scene is aligned to the same scene viewed under the canonical illumination.

In this paper we try to go beyond the limitations of the Miller and Tieu's approach. The key idea is to evaluate color changes within a few reference regions of the scene (e.g. known objects) to compute the global color constancy. The presence of known objects in the image is clearly a restrictive assumption, which is however acceptable in numerous applications, including object recognition. We show that under this assumption much better results can be obtained.

The use of known objects through a matching process is not new in the color constancy area. Obdrzalek *et al.* [7] propose a similar approach, however they use a simple affine model for color change. Our approach is more powerful as our model is able to deal with complex non-linear transformations.

As it is often the case with many color constancy algorithms, our approach will only estimate the chromaticity of the color and will ignore the intensity component. There are many ways of normalizing RGB colors to eliminate the effect of the intensity. For this purpose we will use the standard HSV representation for colors and ignore the V component.

This article gives a description of the proposed method and presents results obtained for the color constancy problem.

2. LEARNED COLOR FLOWS AND LOCAL CORRESPONDENCES

The proposed approach consists of two separate stages:

*Tijmen Moerland is also with LIACS, Leiden University, Netherlands.

- Learning stage. During this stage a model of possible color changes is learned. In order to learn this model, the same scene is viewed under different illuminations allowing to build a statistical model of joint color changes. This model, described in section 2.1, can explain complex non-linear color variation with only a few parameters (usually between 2 and 5). Some *reference patches*, which are reference regions (e.g. known objects) in canonically illuminated images, are stored and used in the second stage.
- Color constancy stage. During this stage, the objective is to change the color of a target image in order to make it as if the illuminant was the canonical illuminant. The constancy is obtained by matching local regions of the target image with reference regions previously stored. These few matches provide a small set of color correspondences allowing to estimate the parameters of the model. The model is then able to explain color changes for the remaining colors (those which are not occurring in the reference regions) allowing to build an image of the scene under the canonical illumination.

The remainder of this section details different parts of our algorithm; we first present the color model (section 2.1), we then explain how to find correspondences between the scene and reference regions (section 2.2) and we finally show how robust estimation techniques allow the method to handle misalignment and noise in the parameter estimation process.

2.1. Color flow model

The color flow model, originally proposed by Miller and Tieu [6], aims to describe a joint color change from a source to target image. Using two aligned images depicting the same scene, changes in color between corresponding pixels can be observed. The color changes can be imagined as vectors in the RGB cube, where the vectors start at the color of a pixel in the source image, and has destination the color of the corresponding pixel in the target image. Direction and length of the vector determine the color change. For every pixel in the image we obtain a vector, thus defining a $R^3 \rightarrow R^3$ vector field. Since on the one hand not every possible RGB color will be present in the source image, and on the other hand some colors will have multiple occurrences, a 3D Gaussian kernel is applied to obtain a continuous, smooth vector field, called *full color flow*. Every dimension in the RGB cube is divided into 16 bins and the values of the color changes of the centers of those bins are stacked in a single, high dimensional, column vector. As this vector represents the full color flow, it is also referred to by the same name.

The described process can be repeated for any number of image pairs. We take random pairs from a database of

training images and determine the full color flows. The vectors corresponding to the full color flows have a length of $3 \times 16^3 = 12288$ and can be seen as a point in the 12288-dimensional space. Principal component analysis gives us directions of highest variance and we find that the total variance in the training sample can be well explained by a small number of principal components. The principal components represent full color flows and are dubbed *eigenflows*.

We have thus determined a group of general color changes that are suitable to model color changes observed in images. Or, observed color changes can be expressed in terms of the eigenflows. Our research has shown that this method is sensitive to photometric effects such as shadows and highlights. Such occurrences have their largest effect in the brightness color component, so we decided to move to hue and saturation space. The third component, brightness, can be ignored as for our purpose of matching other suitable methods exist to deal with brightness.

An interesting property of this method is the adaptability of the model to types of illumination variation. If a general model including all of the possible kind of illuminant is desired, this model can be used. If a very specific illuminant is considered, then the same model can also be used. The only difference will be the images used to train the system and the number of parameters.

2.2. Interest points and correspondences

Expressing the observed color flow in terms of eigenflows is easy for aligned images, as we can simply take the difference between the images and estimate the parameters by means of a least squares estimation. The parameters are referred to as *flow parameters*.

Things are different when there is no such correspondence between pixels. When a scene or object is photographed from different positions, transformations in scale, translation and rotation will have occurred. Furthermore, objects might have moved and certain parts of the scene have now disappeared or new parts have appeared, possible occluding other regions. This is obviously how things are and should be considered. The only assumption is that some reference regions known under a canonical illumination can be seen in the image.

By detecting and matching interest points we can find corresponding regions between the test image and the reference regions. We will show how to estimate the flow parameters from these *invariant* regions. Before going into details of the parameter estimation process, we will detail the detection and matching of interest points.

Interest points are detected using the Difference-of-Gaussian (DoG) detector [8] of David Lowe. The detector provides us with location, scale and orientation of the detected point. Mahalanobis distances in SIFT descriptor space are used to find mutual matches. Experiments have



Fig. 1. Images and result from the Newspaper set.

shown that, under considerable illumination changes, the true positive rate is low, around 50% to 60%. We normalize regions around the interest points, according to their specific scale and orientation, into fixed size image patches. This allows us to compare the regions directly. We use an adapted version of color ratio histogram matching [4] to obtain an estimation of the match probability. A threshold can be set on this probability to reject false matches and increase true positive rate. In our experiments we could usually get this rate up to 90%, sufficient for our purposes.

2.3. Robust estimation

Now that we have pairs of corresponding regions, one coming from the test image, the other coming from a reference image with the reference regions under the canonical illumination, we can try to estimate flow parameters. Remember that we need to observe starting colors and their color changes from the image pair. Furthermore, we should be able to express these color changes in terms of the eigenflow model. This requires corresponding pixels to really correspond to the same physical point in the scene, so that color changes are only due to illumination changes. These pixels are called inliers, while pixels where color changes are not due to illumination changes are called outliers. Note that while we estimate from local regions, we still estimate a *global* color flow.

Unfortunately, there are small differences between position, scale and orientation of the matched interest points, so that their normalized regions will not be completely aligned. But even with small deviations, the percentage of outliers reaches very high levels, possibly 80% to 90%. One of the reasons we choose the DoG detector is that it favors blob-like regions over sharp edges. Thus, while pixels might not correspond to the same physical point, they might have the same or similar colors and contribute to a reliable estimation of the flow parameters.

Nonetheless, a simple least squares estimation is too much affected by outliers and a robust estimation method is

required. We have tested three techniques for robust estimation, two of which are popular within computer vision: M-estimation, RANSAC and MDPE [9]. Testing with synthetically generated reference images showed that M-estimator is remarkably robust to high fractions of outliers. It clearly outperformed the random sampling techniques, probably due to the large fraction of outliers so that too many draws are required to find a clean subset. Estimation of outlier variance is done by computing the Median Absolute Deviation during the first three iterations. Initialization of the M-estimator was done by a simple least squares estimation, outperforming and more stable than initialization by the random sampling methods.

The number of minimal correspondences required to estimate the parameters depend on several factors; the most important ones are the number of parameters of the model and the number and diversity of the colors included in the local regions.

3. RESULTS

The experimental results presented in this section substantiate two key points. First, the color flow method is superior to standard color change models. Secondly, estimating global color change from local regions as opposed to full aligned images does not decrease results. The local method is robust enough and does not invoke a penalty on the estimation results so that we obtain viewpoint invariance for free.

3.1. Training and test sets

For the experiments shown in this article, we have created two pairs of training and test images, one, called 'Tapis', based on a children's book containing many colors in complicated structures, another of newspapers with color images spread out in a random fashion on the floor, called 'Newspapers'. For each pair, two sets (of similar but different scenes) of 10 images were taken under different illu-




Fig. 2. Experimental results: comparing different methods applied on different datasets (see text for details).

mination conditions. Images were taken in a normal office, where the amount of daylight could be adjusted by controlling blinds and fluorescent lights could be turned on and off. Figure 1 shows examples of the Newspaper set.

3.2. Experiments

The color flow model was trained twice separately on 4 aligned image pairs taken from the first sequence of both sets. From this we obtain 4 eigenflows, which were all used for estimation and the model has 4 parameters. The estimation was done on six pairs from the second sequence. For comparison we ran three other color change models. The Gray World model matches the separate color band means from the first image to the second image and derives the three scaling factors from it (3 parameters). The Diagonal model estimates these scaling parameters directly (also 3 parameters), while the affine model additionally estimates offsets in each color band (6 parameters). The last two models can be estimated robustly in the same way as the color flow model. In these cases 30 reference regions are extracted from the training images.

As an error measurement we consider the error in each component of the full image and calculate the Root Mean Squared error in HS space. This measures the model’s performance on all relevant colors, as opposed to the method employed by [7] that only calculates the error in the chromatic components of a white patch. The first two histograms in figure 2 show mean RMS error and the standard deviation for each of the four methods. The cluster on the left is obtained by estimating directly from the full image, while the second cluster is the results of estimation from 30 local corresponding regions (the same robust estimation is used for each method). Comparing these two histograms we find that results remain the same for the color flow and affine model, but get somewhat worse for the gray world and diagonal model. Note that the gray world model can not be estimated robustly. The color flow model outperforms all others in both cases.

To show the invariance to changes in the scene composition, we added in another experiment a randomly colored block to the target image and estimated the parameters. The block had a size of 40% the image size in both dimensions, such that it covered a surface of only 16%. For each pair, 20 such synthetic target images were generated. Errors are calculated in respect to the original, block-less target image and results are displayed in the two histogram clusters on the right of figure 2. Even though the block is not too large, estimation on the global image is severely disturbed by the block’s presence. However results of the local method are unaffected.

Figure 1(c) visualizes the performance of the method when the scene changes. The scenes are different due to the reshuffling of newspapers, but even so our approach is able to catch the color change.

4. CONCLUSION

In this article we demonstrate how the assumption of having a few known reference elements in the scene allows to compensate easily and efficiently for illuminant changes. In our approach the effect of illumination changes on the reference objects is modeled in a training stage. Subsequently, a few correspondences between these reference objects and some canonical views of these objects allow to estimate accurately the dense set of color changes. Experimental results on real images show the efficiency of this method.

5. REFERENCES

- [1] E.A. Day, R.S. Berns, L.A. Taplin, and F.H. Imai, “A psychophysical experiment evaluating the color and spatial image quality of several multispectral image capture techniques,” *The Journal of Imaging Science and Technology*, vol. 48, no. 2, pp. 93–104, 2004.
- [2] J.M. Geusebroek, R. van den Boomgaard, A.W.M. Smeulders, and H. Geerts, “Color invariance,” *PAMI*, vol. 23, no. 12, pp. 1338–1350, December 2001.
- [3] T. Gevers and H. Stokman, “Robust histogram construction from color invariants for object recognition,” *IEEE Transactions on Pattern Analysis and Machine Intelligence*, vol. 26, no. 1, pp. 113–118, 2004.
- [4] B.V. Funt and G.D. Finlayson, “Color constant color indexing,” *PAMI*, vol. 17, no. 5, pp. 522–529, May 1995.
- [5] T. Gevers and A. W. M. Smeulders, “Color based object recognition,” *Pattern Recognition*, vol. 32, pp. 453–464, Mar. 1999.
- [6] E. Miller and K. Tieu, “Color eigenflows: Statistical modeling of joint color changes,” in *Proceedings of the 8th International Conference on Computer Vision, Vancouver, Canada, 2001*, vol. 1, pp. 607–614.
- [7] S. Obdrzalek, M. Matas, and O. Chum, “On the interaction between object recognition and colour constancy,” in *IEEE Workshop on Color and Photometric Methods in Computer Vision, 2003*, pp. xx–xx.
- [8] D.G. Lowe, “Distinctive image features from scale-invariant keypoints,” *IJCV*, vol. 60, no. 2, pp. 91–110, November 2004.
- [9] R.R. Wilcox, *Introduction to Robust Estimation and Hypothesis Testing*, Academic Press, 2004.