

CS 497 Marius and Dhruv's summary of "A combined physical and statistical approach to colour constancy"

Paper: Schaefer, Hordley, and Finlayson, "A combined physical and statistical approach to colour constancy", *CVPR*, 2005.

Summary by Marius Orehovschi and Dhruv Joshi

In this paper, the authors describe approaches for computational color constancy, the problem of estimating the scene illuminant given one image. They describe two main categories of approaches to this problem: statistics-based methods that use the insight that colors observed in the image constrain the set of possible illuminants, and physics-based methods that use the dichromatic reflectance model. Both of these categories of methods yield not just a single most likely illuminant, but a set of possible illuminants along with their respective probabilities. The main finding of the paper is that combining these approaches results in better performance than either one of the two categories of methods, with a reduction in error of at least 20%.

The statistical method:

The color of each pixel gives a better understanding to the algorithm on what color illuminant is for the scene in the image. A statistical method uses color by correlation which uses the color of the pixels to estimate the probability of each of N illuminants. Thus, using a Bayesian probability method, the model builds a vector that suggests the probabilities for each of the N illuminants.

This is done in 5 main steps: choosing the color space, characterizing most realistic illuminants that could reasonably be encountered in the image (both of which only need to be done once for a device), characterizing the input image, correlating the image information with the illuminant characteristics to predict a log probability for each of the N illuminants, and then finally selecting the best illuminant based on these probabilities.

The physics-based method:

The physics based method for determining the illuminant is based on the insight that the color signals of a dichromatic object fall on a two-dimensional plane in the RGB space which contains the illumination vector. Two such objects would yield two planes and their intersection would represent the illumination vector. When estimating the illuminant based on more than two surfaces, the illuminant is considered the "best" intersection of the planes, where "best" is the least-squares solution to an eigenvector problem.

But this intuitive model has been shown to rarely translate well into practice – noise and insufficient segmentation make it so that the model only works well under strict laboratory conditions, with highly saturated surfaces and controlled lighting. The authors propose a more robust method that involves comparing the intersections of pairs of planes. This allows for the elimination of unstable or unlikely results. For example, the intersection of two planes with similar orientation has been proven to be unstable; thus, an intersection of two planes that are separated by an angle smaller than a certain preselected value are ruled out. On the other hand, an intersection that is too far away from the convex hull of likely illuminants is unlikely to describe a plausible illuminant; thus, intersections that form an angle above a certain preselected value with any vector in the convex hull of likely illuminants are also ruled out.

The likely illuminants are then selected from the remaining intersections by a method similar to finding the nearest neighbors to a set of preselected reference lights. But because two intersections can be equally far away from a reference light, neither one would get the vote in this scenario. Thus, the authors propose an alternative method – increment the likelihood of each illuminant by the inverse of the distance between the intersection and the reference light.

Similarly to the correlation method, the physical approach to color constancy provides more than just the most likely illuminant – it provides a set of likely illuminants and their respective likelihoods.

Combining the two:

The authors use a straight-forward approach to combining the results of the two methods: the final result contains the vectors produced by the two methods along with their respective probabilities, where the probability of an illumination vector is simply a weighted sum of the probabilities

yielded by the two methods for the particular vector.

Strengths of the paper and the proposed techniques:

- The authors outline a combination of two strong algorithms that individually performed well and are fast, understandable, and have good results. Linearly combining these two algorithms, the results of the authors are still fast, understandable, and get better results than the individual algorithms.
- The explanation of the methods is accessible and well-organized: the authors first present a simpler, more intuitive model; then they acknowledge its limitations and present a more robust and sophisticated model that the reader can digest more easily after having understood the simpler model first.
- The paper also addresses the fact that even though the improvements in the table look small, they are equivalent to a near 20% improvement to the predecessors.

Weaknesses and limitations:

- The authors could have provided more data tables and more results to illustrate the improvements in performance of the new approach; they could have also organized table 1 by dataset, so that it is easier to make apples-to-apples comparisons.
- Although the paper uses the two algorithms to get better results, the paper does not introduce any significant novel concepts other than the combination of the two previously existing approaches.
- Although the authors show an improvement in performance over previously developed methods on two very particular datasets, they do not address how this improvement translates to real-world situations.

Overall, the paper proposes an important finding and is well-written and accessible. The paper also clearly shows connections to previous work – its new approach is based on combining previously developed methods. The description of the methods is thorough and makes it possible for a qualified reader to reproduce the results in practice. The authors address the fact that the improvements might look small at first sight and show why they are actually significant, with the reduction in error being over 20%. Although the paper's main finding is a simple one, the significant improvement in results deems its simplicity more a strength than a weakness.