Effective Learning-Based Illuminant Estimation Using Simple Features

Dongliang Cheng[1], Brian Price[2], Scott Cohen[2], Michael S. Brown[1]
[1]School of Computing, National University of Singapore. [2]Adobe Research.

Illumination estimation is the process of determining the chromaticity of the illumination in an imaged scene in order to remove undesirable color casts through white-balancing. While computational color constancy is a well-studied topic in computer vision, it remains challenging due to the ill-posed nature of the problem. One class of techniques relies on low-level statistical information in the image color distribution and works under various assumptions (e.g. Grey-World, White-Patch, etc). These methods have an advantage that they are simple and fast, but often do not perform well. More recent state-of-the-art methods employ learning-based techniques that produce better results, but often rely on complex features and have long evaluation and training times. Figure 1 helps to illustrate this with a plot of various statistics-based and learning-based methods in terms of accuracy versus computation time.

In this paper, we present a learning-based method based on color features and show how to use this with an ensemble of regression trees to estimate the illumination. We demonstrate that our approach is not only faster than existing learning-based methods in terms of both evaluation and training time, but also gives the best results reported to date on modern color constancy data sets. Our work is inspired in part by the recent successful method [1] that showed that relatively simple features (color/edge moments) could be used to give good performance in a learning-based framework. In this paper, we simplify the learning-based procedure further to use only four simple features. A key technical contribution of our paper is a method for training an ensemble of decision trees on these features that can accurately predict the chromaticity of the illumination.

An overview of our method is shown in Figure 2. Given an image, four features are extracted: average color chromaticity, brightest color chromaticity, dominant color chromaticity, and chromaticity mode of the color palette. Each of these features are used in a bank of regression trees to generate many illuminant candidates. Results from the multiple regression trees that are in agreement are combined to estimate the illumination. Our learning-based method is based on variance reduction regression trees that have been shown to be a powerful nonlinear predictive model. In particular, our approach estimates K pairs of trees per feature to predict illuminant chromaticity (r, g). Each of these trees is computed from samples in the training data that are biased to a local region in chromaticity space of the ground truth illuminations. The K trees per feature form an ensemble of decision trees from which the final results need to be estimated. The power in the ensemble comes when the different features' trees estimations are in agreement. When there are cross-feature consensus estimations, we take all the output from those trees and add to the output sets. The final estimated illuminant chromaticity is taken as the median of the output set.

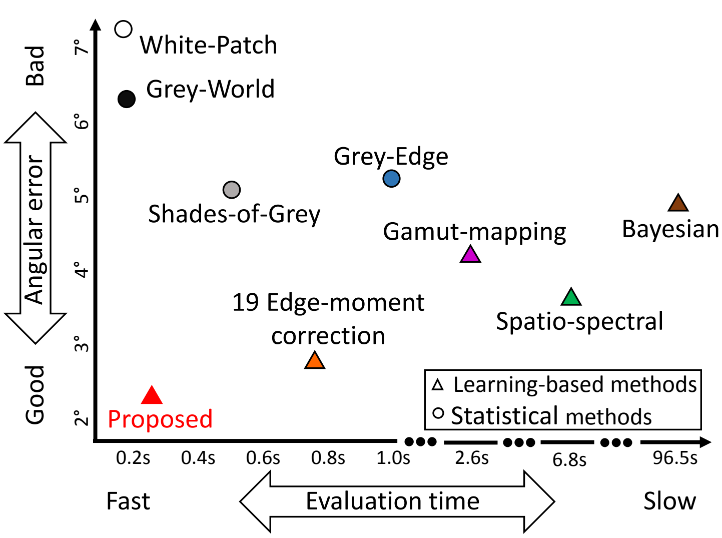

Figure 1: Computational time vs. performance of representative illuminant estimation methods. Statistics-based methods are fast but have lower accuracy than learning-based methods. Our proposed learning-based method achieves both high accuracy and fast computation.

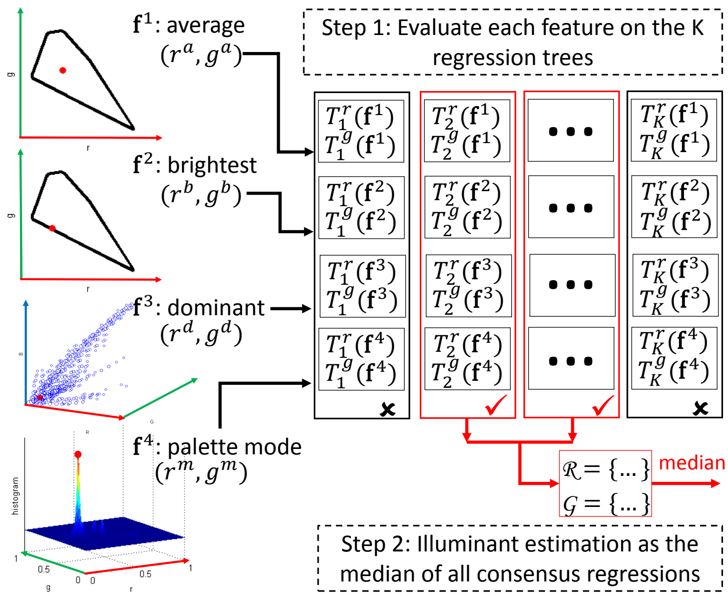

Figure 2: An overview of our proposed framework. Given an input image, four features are extracted and a bank of K regression trees is evaluated. Each regression tree outputs a prediction of the illumination. The final illumination is estimated by combining the results of regression trees that have cross-feature consensus.

Our approach is demonstrated on three standard data sets and produces excellent results with a running-time on par with statistical methods. Our fast running time is attributed to our features simple 2D descriptors computed on the input image's RGB color distribution. In addition, the K tree pairs can be evaluated very quickly given the binary tree structure. Moreover, the training of these trees is reasonably fast. To summarize, this paper has demonstrated a learning-based approach that gives excellent results with running time comparable with statistical methods. The larger implication of this work is that learning-based methods can be viable real-time options and suitable for onboard camera processing.

[1] Graham D Finlayson. Corrected-moment illuminant estimation. In *ICCV*, 2013.

This is an extended abstract. The full paper is available at the Computer Vision Foundation webpage.