

Color Constancy with Spatio-Spectral Statistics

Ayan Chakrabarti, Keigo Hirakawa, and Todd Zickler

Abstract—We introduce an efficient maximum likelihood approach for one part of the color constancy problem: removing from an image the color cast caused by the spectral distribution of the dominating scene illuminant. We do this by developing a statistical model for the spatial distribution of colors in white balanced images (*i.e.* those that have no color cast), and then using this model to infer illumination parameters as those being most likely under our model. The key observation is that by applying spatial band-pass filters to color images one unveils color distributions that are unimodal, symmetric, and well-represented by a simple parametric form. Once these distributions are fit to training data, they enable efficient maximum likelihood estimation of the dominant illuminant in a new image, and they can be combined with statistical prior information about the illuminant in a very natural manner. Experimental evaluation on standard datasets suggests that the approach performs well.

Index Terms—Color Constancy, Statistical Modeling, Spatial Correlations, Maximum Likelihood, Illumination Statistics



1 INTRODUCTION

In addition to its intrinsic reflectance properties, the observed color of a material depends on the spectral and spatial distributions of its surrounding illumination. Thus, in order to use color as a reliable cue for recognition, we must somehow compensate for these extrinsic factors and infer a color descriptor that is stable despite changes in lighting. The ability to make this inference—termed *color constancy*—is exhibited by the human visual system to a certain degree, and there are clear benefits to building it into machines.

One important part of computational color constancy, and the part we consider in this paper, is compensating for the *color cast* that affects a scene as a whole. For this, one ignores spatial variations in lighting spectra and makes the assumption that the spectrum of the illuminant is approximately uniform throughout the scene. The problem is then one of inferring the map $M: \mathbb{R}^3 \rightarrow \mathbb{R}^3$ from the space of observed tristimulus vectors (RGB, human cone, etc.) to the space of *canonical colors*—those that would have been obtained for the same scene by a standard observer under a standard illuminant spectral power distribution.

Like most vision tasks, this version of the color constancy problem is ill-posed, so solving it requires a prior model for the set of canonical colors that exist in a scene. Once equipped with such a model, and optionally a prior model for illuminants as well, we can solve for the color cast (the map M) that best explains the image under this model.

Traditional approaches to this problem use models based on statistics of per-pixel colors, and the principal challenge they face is how to define useful statistics in a

tractable manner. The set of per-pixel colors varies significantly from scene to scene, so methods based on simple first-order statistics like the mean color (*grey world* [1]) or the color of the brightest pixel (*white patch* [2]) can fail quite dramatically. Improving reliability requires more sophisticated models for the set of per-pixel colors in a scene, such as the shape of their convex hull in color space [3] or a non-parametric representation of the empirical probability distribution [4], [5], [6]. These methods improve performance, but at a significant computational cost.

This paper introduces a model that goes beyond statistics of per-pixel colors and leverages joint spatial-color structure. Motivated by the success of filter-based methods [7], [8], we begin by decomposing an input color image into distinct spatial sub-bands, and then we model the color statistics separately in each sub-band. We show that unlike per-pixel colors, the empirical probability distributions of the colors in each sub-band can be well represented using simple parametric (radial exponential) forms, and that these allow inferring the color cast of an image quite effectively. The resulting method outperforms existing approaches on standard databases, operates very efficiently, and allows incorporating a prior model for illuminants in a statistical framework.

1.1 Problem Formulation

Assuming a Lambertian model, we denote the effective spectral reflectance of a surface patch observed at pixel $\mathbf{n} \in \mathbb{Z}^2$ by $\kappa(\lambda, \mathbf{n})$, where $\lambda \in \mathbb{R}$ denotes wavelength. Here, κ accounts for both the material reflectance and surface orientation with respect to the illuminant. We assume a single dominant scene illuminant and represent its spectral distribution by $\ell(\lambda)$. As depicted in Fig. 1, the spectral distribution of the radiance that is emitted toward the observer is then the product of ℓ and κ .

• AC and TZ are with the Harvard School of Engineering and Applied Sciences, Cambridge, MA 02138. KH is with the University of Dayton, Dayton, OH 45469. E-mails: {ayan@eecs.harvard.edu, k.hirakawa@notes.udayton.edu, zickler@eecs.harvard.edu}.

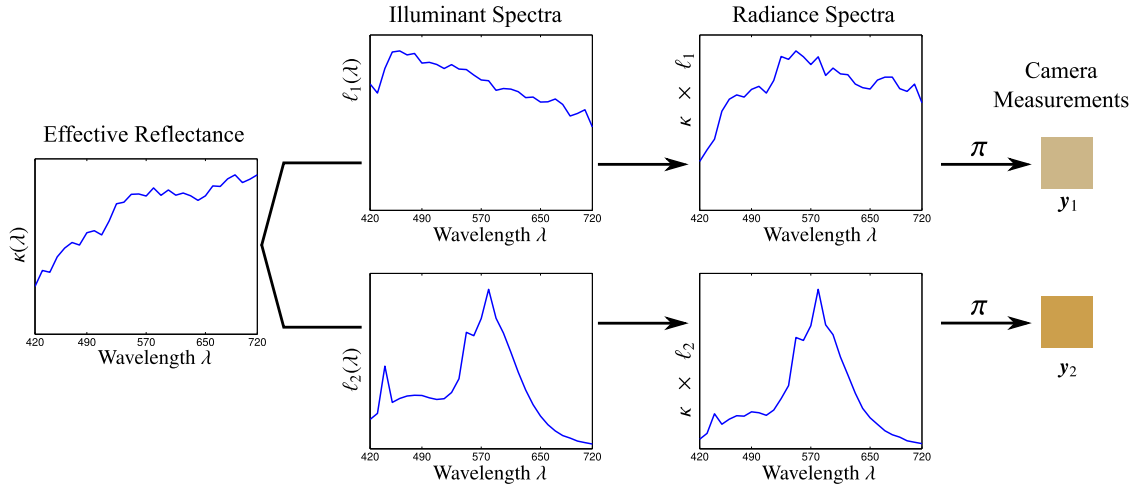


Fig. 1. Camera measurements under different illuminants. The effective surface reflectance (κ) gets shaped by the spectrum (ℓ) of the scene illuminant. Different illuminants can lead to significantly different radiance spectra ($\kappa \times \ell$) being incident on the camera, in turn leading to different trichromatic measurements by the camera’s sensors (π).

We assume that we are given three spectral measurements at each pixel n :¹

$$\mathbf{y}(n) = \begin{bmatrix} y^{\{1\}}(n) \\ y^{\{2\}}(n) \\ y^{\{3\}}(n) \end{bmatrix} := \int \pi(\lambda) \kappa(\lambda, n) \ell(\lambda) d\lambda, \quad (1)$$

where $\pi(\lambda) = [\pi^{\{1\}}(\lambda), \pi^{\{2\}}(\lambda), \pi^{\{3\}}(\lambda)]^T$ are the spectral transmittance distributions of the sensor’s color filters. From image $\mathbf{y}(n)$ our task is to estimate an illuminant-invariant representation, or “canonical color” image, $\mathbf{x}(n)$ given by

$$\mathbf{x}(n) := \int \pi(\lambda) \ell_0(\lambda) \kappa(\lambda, n) d\lambda. \quad (2)$$

We refer to ℓ_0 as our “canonical illuminant”, and for the specific choice $\ell_0(\lambda) = 1$, each canonical color $\mathbf{x}(n)$ can be interpreted as the trichromatic projection (via π) of the effective reflectance $\kappa(\lambda, n)$ at the corresponding surface point.

The colors \mathbf{y} and \mathbf{x} in (1) and (2) are trichromatic reductions of the full spectral distribution that cannot be exactly reversed, even when illuminants ℓ and ℓ_0 are known. However, prior work [9], [10], [11] suggests that a linear mapping of the form

$$\mathbf{x}(n) = \mathbf{M}^{-1} \mathbf{y}(n), \quad \mathbf{M} \in \mathbb{R}^{3 \times 3}, \quad (3)$$

can achieve accurate chromatic adaptation in many cases, and that it can often be approximated further as a diagonal transformation (i.e., $M_{ij} = 0$, if $i \neq j$). In this paper, we apply such diagonal mappings directly in the color space defined by π , even though applying them in an optimized (often called “sharpened”) color space can improve performance [9], [10], [11]. For this reason, our results may provide a conservative estimate of what can be achieved by our algorithm.

With some abuse of notation, we let \mathbf{M} represent a diagonal matrix for the rest of this paper; and with some abuse of terminology, we refer to its recovery as *estimating the illuminant*.

1.2 Related Work

Having settled on a linear diagonal form for the mapping from input color \mathbf{y} to canonical color \mathbf{x} , the color constancy problem reduces to estimating the three diagonal entries of \mathbf{M} . Since both the illuminant and scene reflectances are unknown, this is an under-determined problem, and prior work in this area typically address this by introducing a model or statistical prior for surface reflectance κ or canonical color \mathbf{x} . *White Patch* [2] and *Grey World* [1] are two well-known methods with intuitive interpretations. Based on the observation that color-neutral surfaces (i.e. with constant spectral reflectance) are the most efficient reflectors, the White Patch algorithm posits that pixels observed to have the greatest intensity correspond to a color-neutral surface patch. Similarly, based on the model that the mean surface reflectance in a scene is color-neutral, the Grey World method assumes that the sample mean of $\{\mathbf{x}(n)\}_n$ is achromatic or *grey*². The diagonal elements of \mathbf{M} are then estimated by mapping $\max_n(\mathbf{y}(n))$ or $\text{avg}_n(\mathbf{y}(n))$ to white.

An alternative to these approaches is to represent the set of canonical colors expected in a natural scene by their convex hull, or *gamut* [3]. In a new input image $\mathbf{y}(n)$, the corresponding corrected colors $\{\mathbf{M}^{-1} \mathbf{y}(n)\}_n$ are expected to lie within this gamut, and this induces a set of constraints on the unknown illuminant \mathbf{M} . In many cases, multiple choices of \mathbf{M} may satisfy these constraints for a given image, and different ap-

1. For typical digital cameras, this assumes idealized demosaicking.

2. In fact, Buchsbaum [1] did not require the average color to be grey—only that it be some known system parameter.

proaches [3], [12] have been proposed to arrive at a unique solution from this feasible set.

Another alternative is to define a probability distribution over the space of observed or canonical colors in a scene. The *Color by Correlation* [13] algorithm begins by discretizing the chromaticity space into a number of bins, and then assigns a probabilities to each bin being observed under every candidate illuminant. The scene illuminant M is then chosen to maximize the likelihood of all observed colors in the input image y . Statistical approaches also permit the inclusion of a prior distribution over illuminant spectra. Statistical models for illuminants are combined with those for canonical pixel colors, based on Gaussian [4] or, more effectively, non-parametric [5], [6] distributions, to yield Bayesian estimates of the illuminant M .

Related to these are a number of learning-based approaches that effectively use a training database, of color images with ground truth illuminant data, to learn a function f that maps an observed image y to an estimate of the corresponding illuminant M . For example, Shi *et al.* [14] learn this function f using thin-plate spline interpolation from colors in a training database. Other approaches include learning f in terms of linear filters [15], neural networks [16] and using support-vector regression [17].

All of the above methods can be considered pixel-based, because they model the set of individual pixel colors without considering each pixel’s spatial context. One can arbitrarily re-order the pixels of any input image, for example, without affecting the resulting estimate of illuminant M . Because they ignore spatial structure, these methods can be negatively influenced by scene content that skews the color histogram, such as the presence of large colorful objects.

A variety of work suggests that improvements might be gained by employing spatial image features, such as segmentations or linear filter responses, that incorporate spatial information in a tractable manner. For example, Gershon *et al.* [18] assume that the average of mean colors of segmented regions of an image, rather than of individual pixels, is color neutral. Alternatively, van de Weijer *et al.* [7] and Gijssen *et al.* [8] respectively apply the Grey World and gamut constraint procedures described above to the outputs of linear filters instead of the individual pixels they were originally designed for. (The *Grey Edge* method [7] posits that image gradients are on average color neutral, while the generalized gamut mapping algorithm [8] proposes strategies for combining cues from the expected gamuts of various linear filter coefficients.) Finally, Singh *et al.* [19] reason about the illuminant using a linear spatio-spectral basis for small spatial patches of color images.

The goal of this paper is to leverage joint spatial and spectral structure in a more efficient and effective manner. We achieve this by representing the prior probability distribution over canonical color images in terms of the color coefficients in distinct spatial sub-bands. We show

empirically that, unlike per-pixel colors, the distributions of these sub-band coefficients can be well-represented using convenient parametric forms, and that these parametric forms can be used in a maximum-likelihood framework to estimate the illuminant M quite efficiently and accurately. The paper builds on an earlier version of this work [20], which used Gaussian distributions and tested them on non-linearly transformed (*i.e.* gamma-corrected) color images. Here, we work directly on linear data as defined in (1) and (2) and show that distributions with heavier tails are more appropriate.

2 SPATIO-SPECTRAL MODELING

As stated above, we move beyond statistics of individual pixels colors, drawing inspiration from previous studies [7], [8], [19], and seek to exploit information about a pixel’s spatial context. We take an approach similar to the grey-edge [7] and generalized-gamut [8] methods, by looking at the properties of filter coefficients instead of individual pixels. These methods represent important first steps in looking beyond individual pixels and demonstrate that filter coefficients provide more robust cues for color constancy. However, they are based on the direct application of pixel-based color constancy techniques, grey world [1] and gamut mapping [3] respectively, to filter coefficients, so there is reason to believe that a new estimation procedure, that is specifically tailored to the statistical behavior of these coefficients, can improve performance.

We begin with the observation that statistics of filter coefficients show far more structure than those of individual pixels. Figure 2 compares the empirical histograms of red values of individual pixels to those of coefficients for a particular band-pass filter, in images under a canonical illuminant. We note that the histogram for the filter coefficients is uni-modal and symmetric, while individual pixels show less discernible statistical structure. It is because of this lack of structure that pixel-based color constancy methods have been constrained to either use properties of simple ensemble statistics like the mean (*eg.* grey world [1]), employ expensive non-parametric distributions (*eg.* Rosenberg [5], [6]) or model only the support, or convex hull, of pixel colors (*eg.* gamut mapping [3]).

In contrast, Fig. 2 suggests that filter coefficients have distributions that can be described accurately by parametric models. In this section, we propose such a model and describe a method to learn its parameters from training data. We then present an algorithm that uses this model to estimate the unknown illuminant from an observed image.

2.1 Image Model

We begin by defining a statistical model for an image $x(n)$ observed under canonical illumination. Since pixels in a spatial neighborhood are strongly correlated, we first apply a spatially-decorrelating transform by




Fig. 2. Distributions of individual pixels vs. filter coefficients. Shown are empirical log-histograms (over a standard database [6] of 12-bit color images) of the red values of individual pixels (Left) and coefficients (Right) of a sub-band filter (Inset), for images under canonical illumination. The coefficient histogram shows more structure and can be modeled by parametric distributions.




Fig. 3. Sub-band Decomposition. We use horizontal and vertical second-derivative Gaussian filters at multiple scales to spatially decorrelate the image, and model corresponding coefficients independently.

filtering the image with a series of spatial sub-band filters $\{f_k(\mathbf{n})\}_k$. While these filters can be chosen to correspond to any standard image decompositions, such as wavelet transforms, steerable pyramids, *etc.*, Gaussian derivative filters have proved successful in edge-based color constancy methods [7], [8]. Therefore, we use a decomposition based on horizontal and vertical second-derivative Gaussian filters at multiple scales as illustrated in Fig. 3. Each filter is applied separately to all channels of the image and we define the trichromatic color coefficient vector $\mathbf{x}_k(\mathbf{n})$ as

$$\mathbf{x}_k(\mathbf{n}) = \left[(x^{\{1\}} * f_k)(\mathbf{n}), (x^{\{2\}} * f_k)(\mathbf{n}), (x^{\{3\}} * f_k)(\mathbf{n}) \right]^T, \quad (4)$$

where $*$ denotes convolution. We assume that following the decomposition by the filter bank $\{f_k\}$, each sub-band canonical color image $\mathbf{x}_k(\mathbf{n})$ can be modeled as being independent from the rest.

We represent the prior probability of each sub-band canonical color image as the product of independent, identical distributions per trichromatic coefficient vector, $\prod_{\mathbf{n}} p(\mathbf{x}_k(\mathbf{n}))$, where the distribution $p(\mathbf{x}_k)$ encodes the joint statistics of the three per-channel coefficients in \mathbf{x}_k .




Fig. 4. Equiprobable contours of the distribution of filter coefficient vectors $\mathbf{x}_k(\mathbf{n})$, corresponding to the filter shown in Fig. 2. Contours are shown along the two major eigen-vectors of the covariance matrix Σ_k , for a multi-variate Laplace distribution (Left), a radial exponential distribution (Center), and the empirical distribution computed using kernel density estimation (Right). The empirical distribution has elliptical contours, making the radial exponential distribution a more appropriate choice for modeling these coefficients.

Specifically, we use the *radial exponential* distribution [21],

$$p(\mathbf{x}_k) = \frac{1}{\pi \sqrt{\det(\Sigma_k)}} \exp\left(-2\sqrt{\mathbf{x}_k^T \Sigma_k^{-1} \mathbf{x}_k}\right), \quad (5)$$

where Σ_k is a 3×3 positive-definite matrix corresponding to the covariance of \mathbf{x}_k .

Note that the radial exponential distribution used in (5) is more kurtotic, or heavy-tailed, than a multi-variate Gaussian, and can be thought of as a generalization of the Laplace distribution to the multi-variate case. It differs, however, from the standard multi-variate Laplace distribution [22] in that the multi-variate Laplace has equiprobable contours that are L_1 -spherical or *diamond* shaped (see Fig. 4 (left)) while those of the radial exponential distribution are ellipsoidal (see Fig. 4 (center)). This implies that the components of \mathbf{x}_k along the eigen-vectors of Σ_k are un-correlated but *not* independent under (5). Figure 4 (right) shows the equiprobable contours of the actual empirical distribution of \mathbf{x}_k computed using kernel-density estimation and we see that these contours are indeed ellipsoidal.

2.2 Learning Model Parameters

Having defined a parametric statistical model for canonical color images, $p(\mathbf{x}) = \prod_k \prod_{\mathbf{n}} p(\mathbf{x}_k(\mathbf{n}))$, we are in need of two algorithms. First, we require the ability to fit the model parameters Σ_k given a training set of T canonical color images $\{\mathbf{x}_t\}_{t=1}^T$ (indexed by t), and second, we require the ability to use this learned prior model to reason about the illuminant \mathbf{M} for a given input color image $\mathbf{y}(\mathbf{n})$. We deal with training in this section, and in the next we address illuminant estimation.

We treat each sub-band separately during training, so the problem amounts to independently inferring each Σ_k given the corresponding sub-band of the training set, $\{\mathbf{x}_{k,t}\}_{t=1}^T$. Although, Σ_k corresponds to the covariance matrix of \mathbf{x}_k , the maximum likelihood (ML) estimate of

Σ_k is *not* the empirical covariance of $\{\mathbf{x}_{k,t}\}$. The log-likelihood of the training set as a function of Σ_k is

$$\begin{aligned} J_k(\Sigma) &= \sum_t \log p(\mathbf{x}_{k,t}|\Sigma) \\ &= -\frac{T}{2} \log \det(\Sigma) - \sum_t 2\sqrt{\mathbf{x}_{k,t}^T \Sigma^{-1} \mathbf{x}_{k,t}}. \end{aligned} \quad (6)$$

Unfortunately, the ML estimate of Σ_k , which is given by $\Sigma_k = \arg \max J_k(\Sigma)$, does not have a closed-form solution. Therefore, we propose an iterative estimation procedure.

We initialize Σ_k as the identity matrix. Based on the current estimate Σ_k^* of Σ_k at any iteration, we define an approximation to $J_k(\cdot)$ as

$$J_k^*(\Sigma|\Sigma_k^*) = -\frac{T}{2} \log \det(\Sigma) - \sum_t 2 \frac{\mathbf{x}_{k,t}^T \Sigma^{-1} \mathbf{x}_{k,t}}{\sqrt{\mathbf{x}_{k,t}^T \Sigma_k^{*-1} \mathbf{x}_{k,t}}}. \quad (7)$$

We then update Σ_k in the following iteration as

$$\Sigma_k = \arg \max_{\Sigma} J_k^*(\Sigma|\Sigma_k^*) = \frac{4}{T} \sum_t \frac{\mathbf{x}_{k,t} \mathbf{x}_{k,t}^T}{\sqrt{\mathbf{x}_{k,t}^T \Sigma_k^{*-1} \mathbf{x}_{k,t}}}. \quad (8)$$

The true ML estimate is a fixed point of these iterations, and we find that the procedure converges quickly in practice (usually in less than five iterations).

2.3 Illuminant Estimation

Once we have learned the parameters of the prior model, we are ready to infer the illuminant M for an input color image $\mathbf{y}(\mathbf{n})$. As per our model in (3), we assume that $\mathbf{y}(\mathbf{n}) = M\mathbf{x}(\mathbf{n})$, where M is a diagonal 3×3 matrix with positive entries. Since convolution is a linear operation, it follows that $\mathbf{y}_k(\mathbf{n}) = M\mathbf{x}_k(\mathbf{n})$, where $\mathbf{y}_k(\mathbf{n})$ are the color sub-band coefficients of $\mathbf{y}(\mathbf{n})$ for filter f_k as in (4).

We begin by noting that under our prior model for \mathbf{x}_k , the likelihood of \mathbf{y}_k given M is given by

$$\begin{aligned} p(\mathbf{y}_k|M) &= \frac{\exp\left(-2\sqrt{\mathbf{y}_k^T (M\Sigma_k M)^{-1} \mathbf{y}_k}\right)}{\pi \sqrt{\det(M\Sigma_k M)}} \\ &\propto \frac{1}{\det(M)} \exp\left(-2\sqrt{\mathbf{y}_k^T (M\Sigma_k M)^{-1} \mathbf{y}_k}\right). \end{aligned} \quad (9)$$

Therefore, conditioned on M the coefficients $\mathbf{y}_k(\mathbf{n})$ also have a radial exponential distribution, with the covariance matrix $M\Sigma_k M$. Figure 5 compares the covariance matrices Σ_k and $M\Sigma_k M$ for a typical illuminant M . This difference between the k th sub-band coefficient distributions for the input image \mathbf{y}_k and the canonical color image \mathbf{x}_k , as embodied by these covariance matrices, is the fundamental cue that we will exploit for illuminant estimation.

For notational convenience, we define $\mathbf{m} = [m_1, m_2, m_3]$ and $\mathbf{w} = [w_1, w_2, w_3]$ to be the diagonal elements of M and M^{-1} respectively, where

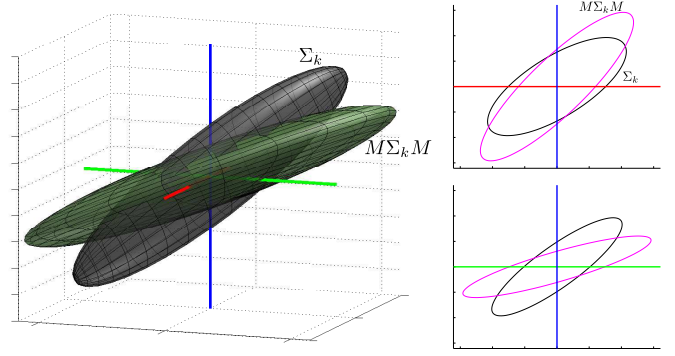


Fig. 5. Covariance matrices Σ_k and $M\Sigma_k M$, for coefficients observed under a canonical illumination and a typical illuminant M . (Left) Ellipsoids corresponding to the covariance matrices, and (Right) their projections onto the R-B and G-B planes. The illuminant causes a skew in the shape of the covariance matrix, which serves as a cue for our estimation method.

$m_i = w_i^{-1}$. Then, the ML estimate of M given all observed coefficients $\{\mathbf{y}_k(\mathbf{n})\}_{k,n}$ is given by

$$\begin{aligned} \hat{M}_{ML} &= \arg \max_{M \in \text{diag}(\mathbb{R}^3)} \sum_{k,n} \log p(\mathbf{y}_k(\mathbf{n})|M) \\ &= \arg \min \sum_{k,n} \log \det(M) + 2\sqrt{\mathbf{y}_k^T(\mathbf{n})(M\Sigma_k M)^{-1} \mathbf{y}_k(\mathbf{n})} \\ &= \arg \min N \log m_1 m_2 m_3 \\ &\quad + \sum_{k,n} 2\sqrt{\mathbf{w}^T [(\mathbf{y}_k(\mathbf{n})\mathbf{y}_k^T(\mathbf{n})) \circ \Sigma_k^{-1}] \mathbf{w}}, \end{aligned} \quad (10)$$

where $N = \sum_{k,n} 1$ is the total number of coefficients, and $A \circ B$ denotes the Hadamard or entry-wise product of matrices A and B .

We propose an iterative algorithm to solve (10) using a similar approach as the training method. Based on the estimate M^* of the illuminant at every iteration, we approximate the cost function in (10) as

$$\begin{aligned} L(M|M^*) &= N \log m_1 m_2 m_3 \\ &\quad + \sum_{k,n} 2 \frac{\mathbf{w}^T [(\mathbf{y}_k(\mathbf{n})\mathbf{y}_k^T(\mathbf{n})) \circ \Sigma_k^{-1}] \mathbf{w}}{\sqrt{\mathbf{y}_k^T(\mathbf{n})(M^* \Sigma_k M^*)^{-1} \mathbf{y}_k^T(\mathbf{n})}} \\ &= N \left[\log m_1 m_2 m_3 + \frac{1}{2} \mathbf{w}^T \mathbf{A}^* \mathbf{w} \right], \end{aligned} \quad (11)$$

where \mathbf{A}^* is a 3×3 symmetric matrix given by

$$\mathbf{A}^* = \frac{4}{N} \sum_k \left(\sum_n \frac{\mathbf{y}_k(\mathbf{n})\mathbf{y}_k^T(\mathbf{n})}{\sqrt{\mathbf{y}_k^T(\mathbf{n})(M^* \Sigma_k M^*)^{-1} \mathbf{y}_k^T(\mathbf{n})}} \right) \circ \Sigma_k^{-1}. \quad (12)$$

In the next iteration, the illuminant estimate M is set to $\arg \min L(M|M^*)$. There is a closed-form solution to this expression for each element of m_i of \mathbf{m} when the

others elements $m_j, j \neq i$ are fixed. This is given by

$$\begin{aligned} \frac{\partial}{\partial m_i} L(\mathbf{M}|\mathbf{M}^*) &= 0 \Rightarrow m_i^2 - \left(\sum_{j \neq i} \frac{A_{ji}^*}{m_j} \right) m_i - A_{ii}^* = 0 \\ \Rightarrow m_i &= \frac{1}{2} \left(\sum_{j \neq i} \frac{A_{ji}^*}{m_j} + \sqrt{\left(\sum_{j \neq i} \frac{A_{ji}^*}{m_j} \right)^2 + 4A_{ii}^*} \right). \end{aligned} \quad (13)$$

We therefore apply (13) on each channel i keeping the others fixed, and iterate until convergence.

Once we have estimated $\hat{\mathbf{M}}_{ML}$ through this procedure, the canonical image $\mathbf{x}(\mathbf{n})$ can be computed simply as $\mathbf{x}(\mathbf{n}) = \hat{\mathbf{M}}_{ML}^{-1} \mathbf{y}(\mathbf{n})$.

3 ILLUMINANT PRIOR

Implicit in the maximal likelihood estimation of (10) is the notion that all illuminants are equally likely. Drawing from the success of existing color constancy algorithms [4], [6] that leverage the *a-priori* knowledge of the illuminant distribution, however, we develop a strategy to incorporate an illuminant prior. The experimental validation in Section 4 will show that such a *regularized* estimate can improve robustness.

Consider an illuminant prior distribution $p(\mathbf{M})$ of the form

$$p(\mathbf{M}) = \left(\frac{2}{\pi} \right)^{3/2} \frac{\det(\mathbf{Q})^{-1/2}}{m_1^2 m_2^2 m_3^2} \exp\left(-\frac{1}{2} \mathbf{w}^T \mathbf{Q}^{-1} \mathbf{w}\right), \quad (14)$$

where \mathbf{Q} is a 3×3 positive-definite matrix, and $\mathbf{m}, \mathbf{w} > 0$ are as defined before. It is useful to note that this multivariate Gaussian distribution on \mathbf{w} (restricted to positive values) is a conjugate prior to the likelihood expression in (10). We note that (14) may be used to describe the statistics of scene illuminants in general, or those restricted to a category of illuminants (such as indoor v.s. outdoor).

Ideally, the covariance parameter \mathbf{Q} would be computed as a sample covariance of a training set $\{\mathbf{w}_t\}$ of illuminant coefficients. However, only the relative spectral radiance of illuminants is usually available, meaning the absolute luminance of the illuminants in the training set is unknown. Therefore, we compute the sample covariance \mathbf{Q}' over the set of illuminant vectors normalized such that $\|\mathbf{m}_t\| = 1$ during training. During estimation, we first compute the ML estimate $\hat{\mathbf{M}}_{ML}$ of (10) (that is, without the illuminant prior), and then reweight \mathbf{Q}' covariance parameter as

$$\mathbf{Q} = \|\mathbf{m}_{ML}\|^2 \mathbf{Q}'. \quad (15)$$

This reweighting has the effect of normalizing the training set $\|\mathbf{m}_t\|$ to the initial estimate $\|\mathbf{m}_{ML}\|$.

Let \mathbf{y}_k be the k th sub-band coefficients of input image as before. The final illuminant estimate incorporating $p(\mathbf{M})$ is obtained by solving

$$\hat{\mathbf{M}} = \arg \max \left[\sum_{k, \mathbf{n}} \log p(\mathbf{y}_k(\mathbf{n})|\mathbf{M}) \right] + \alpha \log p(\mathbf{M}). \quad (16)$$

Here, α is a scalar parameter that weights the relative contribution of the prior to the image evidence, and is learned using cross-validation to minimize estimation error on a training set. While for the special case of $\alpha = 1$, (16) corresponds to maximum a-posteriori (MAP) estimation, we find that a larger weight on the prior yields better estimates in practice.

Solving (16) in a manner similar to (13), we define an approximated cost as

$$L_p(\mathbf{M}|\mathbf{M}^*) = (N + 2\alpha) \left[\log m_1 m_2 m_3 + \frac{1}{2} \mathbf{w}^T \mathbf{A}_p^* \mathbf{w} \right], \quad (17)$$

where \mathbf{A}_p^* is a 3×3 matrix given by

$$\mathbf{A}_p^* = \frac{N \mathbf{A}^* + \alpha \mathbf{Q}}{N + 2\alpha}. \quad (18)$$

Applying (13) iteratively to \mathbf{A}_p^* (instead of \mathbf{A}^*) solves $\mathbf{M} = \arg \min L_p(\mathbf{M}|\mathbf{M}^*)$. If \mathbf{M}^* is set to $\hat{\mathbf{M}}_{ML}$, we find that further iterations to update L_p or \mathbf{Q} are unnecessary.

4 EXPERIMENTAL EVALUATION

We primarily evaluate the proposed method on the ‘‘Color Checker’’ database collected by the authors of [6], which includes 568 color images of which 246 were labeled as captured indoors, and 322 as captured outdoors. Each image has a color checker chart at manually marked co-ordinates which serves as ground truth, and is masked out during evaluation. We use the version of the database made available by Shi and Funt [23]. It consists of 12-bit linear images in the sensor color space generated directly from the RAW captured data—without using the camera’s auto-white balance, gamma-correction or any demosaicking (the sensor data is sub-sampled to provide a trichromatic vector at each pixel). While we focus most of our discussion on the performance of various algorithms on this database, we also include results for two other common databases in Sec. 4.3.

A recent survey paper [24] provides a comprehensive evaluation of a large number of color constancy methods on this database (*i.e.* the re-processed version from [23]), allowing us to compare the proposed algorithm to these methods. Performance is measured in terms of the angular deviation between the estimated and true illuminant, where the later is computed from the brightest grey patch in the color chart. We use the same error metric and ground truth to evaluate our method and compare performance to the state of the art: grey world [1], grey edge [7], gamut mapping [3], generalized gamut mapping [8] and Rosenberg [5], [6].

4.1 Implementation Details

We implement the proposed method with second-derivative Gaussian filters at three different scales ($\sigma = 1, 2$ and 4). We report performance both with and without the illumination prior. For the case with the prior,

TABLE 1
Angular error quantiles for various methods on “Color Checker” database [6]

	All Images (568)			Indoor (246)			Outdoor (322)		
	Mean	Median	Worst-25%	Mean	Median	Worst-25%	Mean	Median	Worst-25%
Grey World [1]	4.6°	4.1°	8.7°	5.8°	5.5°	9.9°	3.7°	3.1°	7.1°
Grey Edge (1 st Order) [7]	4.1°	3.5°	8.0°	4.8°	4.1°	8.8°	3.5°	2.9°	7.1°
Grey Edge (2 nd Order) [7]	4.0°	3.4°	7.8°	4.7°	4.0°	8.5°	3.5°	2.8°	7.1°
Gamut Mapping ($\sigma = 5$) [3]	4.1°	2.5°	10.3°	5.5°	4.3°	12.3°	3.1°	1.8°	7.9°
Generalized Gamut ($\sigma = 5, 1-jet$) [8]	4.1°	2.5°	10.3°	5.5°	4.4°	12.4°	3.1°	1.8°	7.9°
Rosenberg [6]	4.8°	3.5°	10.5°	6.5°	5.9°	12.0°	3.5°	2.4°	7.9°
Proposed (ML Estimate)	3.7°	3.0°	7.6°	4.2°	3.6°	8.1°	3.3°	2.5°	7.1°
Proposed (with general Prior)	3.6°	3.0°	7.4°	4.2°	3.6°	7.9°	3.2°	2.4°	6.9°
Proposed (with category-wise Prior)	3.1°	2.3°	6.5°	4.1°	3.7°	7.8°	2.3°	1.9°	4.6°

we evaluate a general prior over all illuminants, as well as separate priors for the indoor and outdoor images. To train the parameters of our model $\{\Sigma_k\}$, Q' and α , we use three-fold cross validation where the database is split into three equal folds, and when testing on each fold, we use the remaining images for training. MATLAB source code for this implementation, and the estimated illuminants for all images in this database are available for download at <http://vision.seas.harvard.edu/colorconstancy/>.

Since the authors of [24] have made the estimated illuminants and errors for all evaluated methods available, we report performance for most other methods directly from their data. However, we incorporate an important enhancement to the evaluation of grey world and grey edge that improves their performance. While grey world is often interpreted as assuming the mean color of an image to be along $[1, 1, 1]$ (as is the case in [24]), we posit that the mean is along a 3×1 unit vector \hat{g} . The illuminant estimate is then given by:

$$\hat{m} = [\text{diag}(\hat{g})]^{-1} \text{avg}(\mathbf{y}(n)), \quad (19)$$

where $\text{avg}(\mathbf{y}(n))$ is the trichromatic average of the observed image. This approach is sometimes referred to as *database grey world* [25], and we apply the same interpretation to grey edge, and learn the vector \hat{g} for each algorithm from a training set of images as the mean of the unit-vectors corresponding to the pixel or edge averages of each image. We use three-fold cross validation for this training, and to pick the optimal smoothing and norm parameters for grey edge as described in [24].

4.2 Results

Table 1 reports the mean and median errors for all algorithms, as well as the *Worst-25%* error which is a measure of robustness and refers to the mean of the 25% highest error values (note that these may correspond to different images for different methods). These quantiles are reported for the entire database, as well as separately for the indoor and outdoor sets. Figures 6 and 7 show

TABLE 2
Sign test results on the “Color checker” database [6]

	Grey World 1	Grey Edge (1 st Order) 2	Grey Edge (2 nd Order) 3	Gamut Map. ($\sigma = 5$) 4	Gen. Gamut ($\sigma = 5, 1-jet$) 5	Rosenberg 6	Proposed (ML) 7	Proposed (General Prior) 8	Proposed (categ. Prior) 9
1	0	-1	-1	-1	-1	0	-1	-1	-1
2	1	0	-1	-1	-1	0	-1	-1	-1
3	1	1	0	-1	-1	1	-1	-1	-1
4	1	1	1	0	0	1	0	-1	-1
5	1	1	1	0	0	1	-1	-1	-1
6	0	0	-1	-1	-1	0	-1	-1	-1
7	1	1	1	0	1	1	0	-1	-1
8	1	1	1	1	1	1	1	0	-1
9	1	1	1	1	1	1	1	1	0

Value at (row i , column j)	Interpretation at 95% confidence level
1	Algorithm i tends to have lower errors
-1	Algorithm j tends to have lower errors
0	No significant difference

examples, of indoor and outdoor images respectively, corrected according to the illuminant estimates of various methods. Table 2 shows results of a *sign test* [26], [27] comparing the performance of every pair of algorithms over the entire database. This test determines whether one algorithm tends to have lower errors compared to another, by using significance testing to reject the null hypothesis that $P(\text{err}_i > \text{err}_j) = P(\text{err}_j > \text{err}_i)$, where err_i and err_j denote angular errors from two algorithms i and j on a random image. Note that the sign test is based on only the sign of the difference in errors for each image, not the magnitude.

We first look at the relative performance of the ML estimator, without any prior information. The sign test results indicate that it typically has lower errors in comparison to most of the state of the art methods over the entire database. The only exception is traditional gamut-mapping which has higher errors than the proposed




Fig. 6. Indoor images from the color checker database, corrected using different algorithms. Angular errors for estimated illuminant are indicated below each image.

method on a majority of the images, but by a margin too small to be statistically significant. We also find that the ML estimates have lower mean and worst-25% error values amongst all state of the art methods. Both gamut-based methods have lower median errors, largely due to their superior performance on the outdoor images that were captured in daylight. They have the lowest

mean and median errors for those images, although the proposed ML estimator appears to be more robust with a lower worst-25% error. However, the gamut-based methods perform worse on indoor images and have the highest error quantiles on that set after Rosenberg. In general, all algorithms show poorer performance on the indoor set than on outdoor images, indicating that

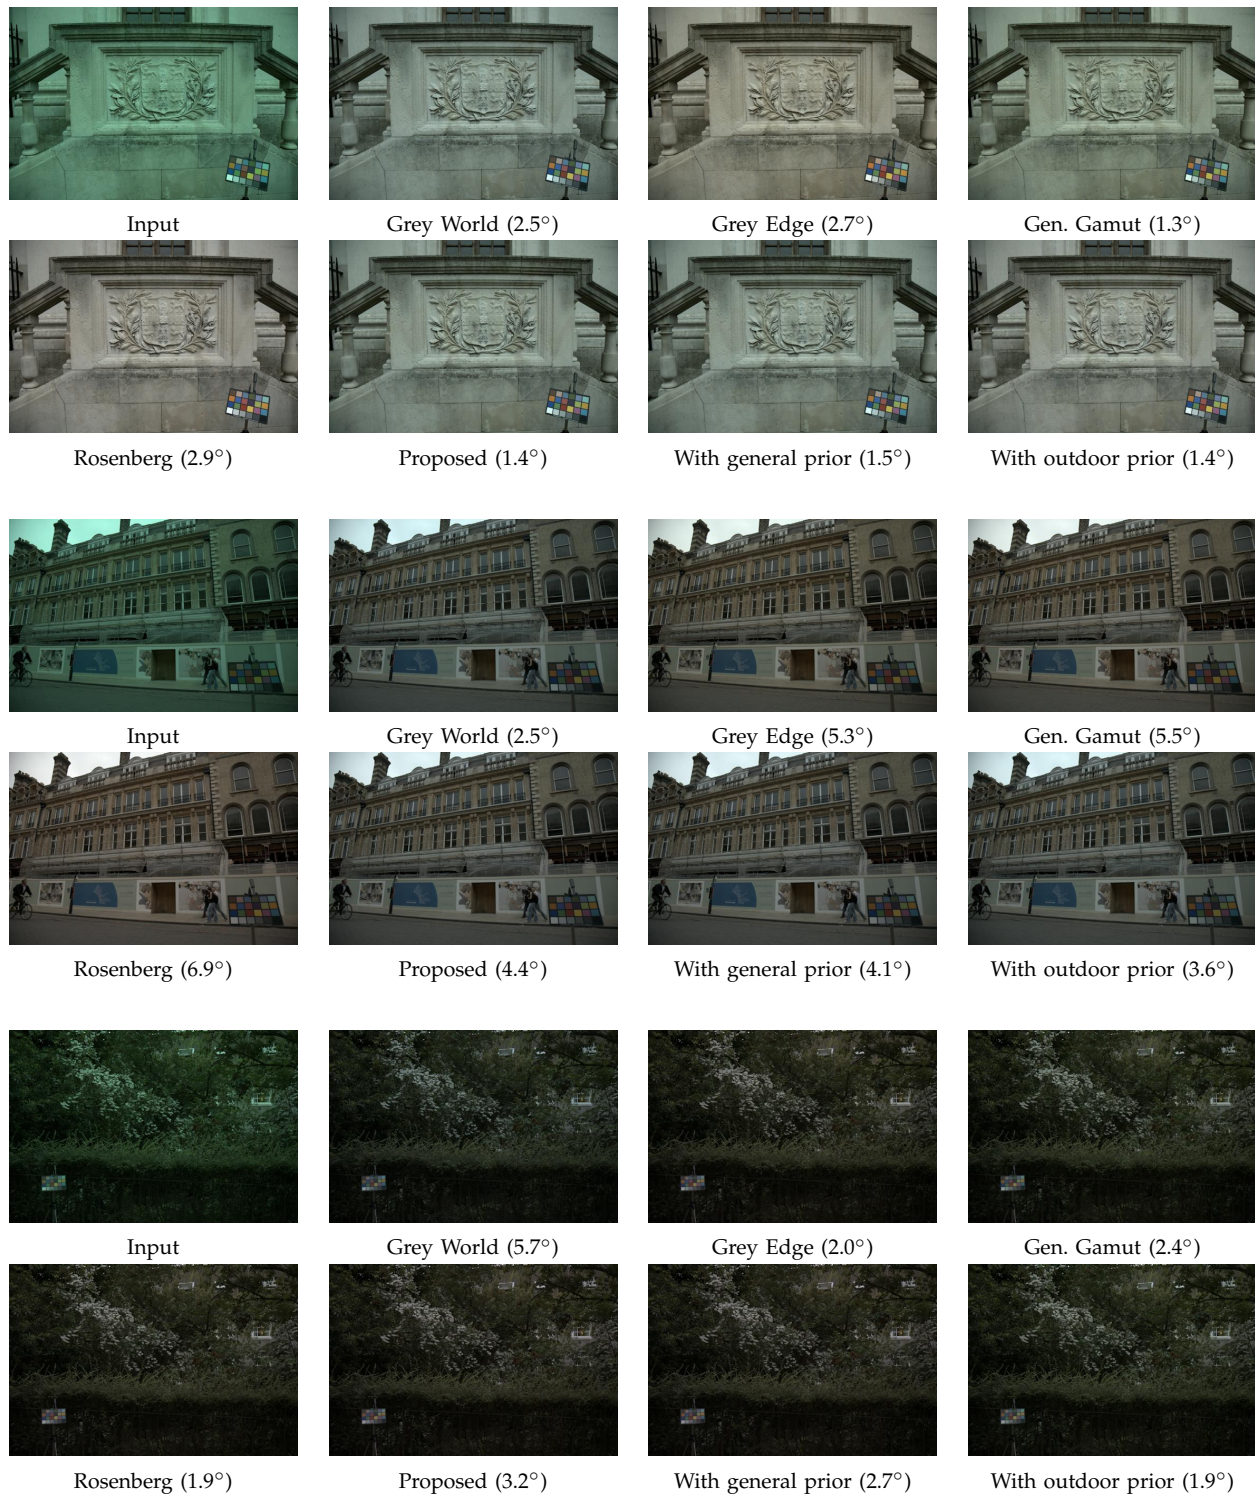


Fig. 7. Outdoor images from the color checker database, corrected using different algorithms. Angular errors for estimated illuminant are indicated below each image.

indoor scenes are more likely to contain regions that act as outliers. However, the ML estimator performs noticeably better than the state of the art, with the lowest values for all error quantiles on indoor images.

Next, we note that an illuminant prior, even a general one over all images, results in better performance. The sign test shows that the prior leads to a statistically

significant improvement in comparison to the ML estimator, and errors are now likely to be lower than for all other methods, including gamut mapping. However, note that while there is an overall improvement in accuracy, estimation errors on individual images may increase as is the case for some examples in Figs. 6-7. In Table 1, we see that the most significant improvement is

TABLE 3

Performance of grey edge and proposed method on the “Color checker” database [6], with both using the same set of second-derivative Gaussian filters with $\sigma = 1$

	Mean	Median	Worst-25%
Grey Edge	4.0°	3.4°	7.8°
Proposed (ML Estimate)	3.9°	3.2°	7.7°
Proposed (with gen. prior)	3.7°	3.0°	7.5°

to the worst-25% quantile for all sets, indicating that the prior’s major contribution is in making the estimation process more robust.

When the illuminant prior is defined separately over sets of indoor and outdoor images (and it is possible that these labels may be obtained during estimation from a light-meter on the camera or a scene-classification algorithm), the improvement in overall performance is dramatic. We note that most of this improvement is on the outdoor images, which is likely because of relatively less variability in the set of daylight illuminants.

While the spatio-spectral model allows a natural way to combine cues across sub-bands at multiple scales, we seek to gauge how well the proposed estimator succeeds at exploiting the information in each sub-band. Therefore, we perform a more direct comparison of our method to the grey edge algorithm, by restricting the former to work on the same set of coefficients as the latter. Table 3 shows the error quantiles of the two methods (over all images) when both use second-derivative Gaussian filters at scale $\sigma = 1$, which corresponds to the best choice for grey edge. We find that the proposed method yields better estimates even when not using an illuminant prior, and appears to benefit from accurately modeling the statistical structure of sub-band coefficients.

4.3 Evaluation with Other Databases

Finally, we report performance on two older datasets considered in [24]. The first is a large set of real-world scenes captured by a consumer video camera [28], and the second is a relatively small set of objects captured under various illuminants in laboratory conditions [29]. We use the evaluation results from [24] as an indicator of the performance of state-of-the-art methods, but again, we differ in applying the learning step to grey world and grey edge, as described in Sec. 4.1.

The “Grey Ball” database [28] is a collection of 11346 images captured using a video camera. A diffuse grey sphere is attached to the camera and is present at the same location in all images, providing a ground truth estimate of the illuminant. The original database includes gamma-corrected images, and we follow the procedure in [24] by measuring performance after applying inverse gamma-correction to the images and recomputing the ground truth accordingly. We also adopt the approach suggested in [24] for cross-validation: since the database

TABLE 4

Performance of various methods on “Grey Ball” set [28]

	Mean	Median	Worst 25%
Grey World [1]	12.2°	9.7°	25.4°
Grey Edge (1 st Order) [7]	10.8°	9.2°	20.9°
Grey Edge (2 nd Order) [7]	11.4°	9.9°	21.5°
Gamut Mapping [3]	11.8°	8.9°	24.9°
Generalized Gamut Mapping [8]	11.8°	8.9°	24.9°
Proposed (ML Estimate)	10.3°	8.9°	20.3°

is divided into 15 videos, we carry out estimation on the frames of each video using the remaining videos as training. Table 4 shows error quantiles of the proposed method and compares them to those of other algorithms. Note that we only report errors for the ML estimator, since the prior-based estimates have identical error quantiles.

Like with the color checker database [6], we find that the proposed method yields more accurate estimates than other algorithms. However, results on this database should be interpreted with caution as the images are lower in quality, with possible artifacts from compression and the approximate inverse-gamma map. Further, the ground truth information may be unreliable in some cases because the grey ball is close to the camera and can be lit by an illuminant that is different from the dominant illuminant of the scene.

We also report results on the “SFU laboratory” database [29], which consists of 321 images of 31 different objects captured under different illuminants in a laboratory. For training, we take the approach suggested in [24] and fit model parameters to the images of all 31 objects captured under illuminant *syl-50MR16Q* before using these parameters to measure estimation accuracy on the entire database. Note that this means the training and testing images overlap. Further, since the training set contains all objects that appear in the database, albeit under different illuminants, this training scheme may be prone to overfitting. Finally, since the training data contains only one unique illuminant, we can only evaluate the maximum likelihood version of our method, which does not make use of a prior over illuminants.

We report error quantiles for various methods on this database in Table 5. We find that the proposed method performs better than grey world and first-order grey edge. It has a higher mean error and median error than second-order grey edge, but it appears to be more robust with a lower worst-25% error.

Interestingly, the gamut mapping methods perform the best with significantly lower error quantiles. One possible explanation is that object-scale images tend to have a smaller number of independent samples of pixels/coefficients than scene-scale images, so that modeling the support (or convex hull) of the observed set is more robust than attempting to describe and exploit its distribution, as done by grey world, grey edge and

TABLE 5
Performance of various methods on “SFU Lab” set [29]

	Mean	Median	Worst 25%
Grey World [1]	8.8°	5.6°	20.7°
Grey Edge (1 st Order) [7]	6.6°	4.3°	15.6°
Grey Edge (2 nd Order) [7]	5.1°	2.5°	13.6°
Gamut Mapping ($\sigma = 4$) [3]	3.7°	2.3°	9.3°
Gen. Gamut ($\sigma = 4,2$ -jet) [8]	3.6°	2.1°	9.4°
Proposed (ML Estimate)	5.6°	3.5°	12.9°

the proposed algorithm. However, drawing strong conclusions from images of only thirty-one distinct objects would be premature. Since they learn a larger number of parameters and are trained on the entire set of reflectances that occur in the database, it is also possible that the gamut mapping methods are simply over-fitting to the convex hull of this small object set.

5 CONCLUSION

In this work, we describe an effective way to exploit spatial dependencies among pixels in a color image for the color constancy problem. We model these dependencies by first decomposing the image using a set of spatially-decorrelating filters, and then analyzing the statistics of their coefficients. These color coefficients are *better-behaved* than colors of individual pixels, and they enable color constancy techniques that are quite accurate and efficient. When no prior information about the illuminant is available, it can be estimated in a maximum likelihood framework; and for cases when one has prior illuminant information, we introduce a prior model that can be incorporated during estimation.

Interesting directions for future work include relaxing the two main assumptions made in this work: spatially-constant illuminant spectrum and diagonal transforms for chromatic adaptation. Even when there is a single illuminant in the scene, mutual illumination causes the *effective* illuminant spectrum to vary from point to point. A first step that could lead to improvement would be to remove regions that appear to be outliers to the general statistics of the image, in the hope that these are likely to correspond to areas with significant inter-reflections. Alternatively, one could explore variants of the proposed method that assume the scene to be lit by a mixture of n illuminants. The problem would then be to estimate which parts of the image are lit by which illuminant, as well the color of each of these illuminants.

It would be interesting to explore avenues to use the spatio-spectral model introduced here to go beyond diagonal transforms. One path is to estimate general linear transforms under appropriate constraints. Another is to consider a set of registered training images taken under different illuminants, and then learn the functional form of the map between corresponding pixels. The parameters of this map can then be estimated from a test image to do color constancy.

Another challenge is to adapt this method for applications that need color constancy to work on images captured using consumer digital cameras that do not provide linear (RAW) image data. In this work, we have assumed that training and testing are performed with images taken from the same or similar cameras. This is important, because as shown in [30], the color spaces and non-linear processing done in cameras can vary significantly, affecting both the sub-band statistics and the illuminant statistics in Q . Saenko *et al.* [31] look at the problem of adapting visual category models learned from one camera to another. It would be interesting to investigate whether the same can be done for color constancy in a way that would allow training data from one camera or linear images to be adapted to a new domain, with minimal additional information.

Finally, we note that computational color constancy algorithms are often motivated by observations from the human visual system [2], [32]. The relative success of this method therefore naturally raises questions about whether the human visual system employs related processing for adaptation and color constancy. Indeed, psychophysical experiments have shown a strong interaction between the spatial orientation and frequency of a stimulus and the chromatic adaptation it induces [33], [34]. The mechanisms that govern these interactions are poorly understood [33], and the machinery developed in this paper might be useful in creating experiments to analyze them further. Separate experiments have linked the textures of familiar objects to human color perception [35], [36], and while texture is currently thought to affect color perception through object memory, it is worth exploring the contribution of spatial correlations to this effect.

ACKNOWLEDGMENTS

The authors would like to thank the associate editor and reviewers for their thoughtful comments. A.C. and T.Z. were supported by ONR award N000140911022 and the US Army Research Laboratory and the US Army Research Office under contract/grant number 54262-CI.

REFERENCES

- [1] G. Buchsbaum, “A spatial processor model for object colour perception,” *Journal of the Franklin Institute*, vol. 310, no. 1, pp. 1–26, 1980.
- [2] E. Land, “The retinex theory of color vision,” *Scientific American*, vol. 237, no. 6, pp. 108–128, 1977.
- [3] D. Forsyth, “A novel algorithm for color constancy,” *International Journal of Computer Vision*, vol. 5, no. 1, pp. 5–36, 1990.
- [4] D. Brainard and W. Freeman, “Bayesian color constancy,” *Journal of the Optical Society of America A (JOSA A)*, vol. 14, no. 7, pp. 1393–1411, 1993.
- [5] C. Rosenberg, T. Minka, and A. Ladsariya, “Bayesian color constancy with non-gaussian models,” in *Advances in Neural Information Processing Systems (NIPS)*, 2003.
- [6] P. Gehler, C. Rother, A. Blake, T. Minka, and T. Sharp, “Bayesian color constancy revisited,” in *Proc. IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, 2008.
- [7] J. van de Weijer, T. Gevers, and A. Gijzenij, “Edge-based color constancy,” *IEEE Transactions on Image Processing*, vol. 16, no. 9, pp. 2207–2214, 2007.

- [8] A. Gijzenij, T. Gevers, and J. van de Weijer, "Generalized gamut mapping using image derivative structures for color constancy," *International Journal of Computer Vision*, vol. 86, no. 2, pp. 127–139, 2010.
- [9] G. West and M. H. Brill, "Necessary and sufficient conditions for von Kries chromatic adaptation to give color constancy," *Journal of Mathematical Biology*, vol. 15, no. 2, pp. 249–258, 1982.
- [10] G. Finlayson, M. Drew, and B. Funt, "Diagonal transforms suffice for color constancy," in *Proc. International Conference on Computer Vision (ICCV)*, 1993.
- [11] H. Chong, S. Gortler, and T. Zickler, "The von Kries hypothesis and a basis for color constancy," in *Proc. International Conference on Computer Vision (ICCV)*, 2007.
- [12] K. Barnard, "Improvements to gamut mapping colour constancy algorithms," in *Proc. European Conference on Computer Vision (ECCV)*, 2000.
- [13] G. Finlayson, S. Hordley, and P. Hubel, "Color by correlation: A simple, unifying framework for color constancy," *IEEE Transactions on Pattern Analysis and Machine Intelligence*, vol. 23, no. 11, pp. 1209–1221, 2001.
- [14] L. Shi, W. Xiong, and B. Funt, "Illumination estimation via thin-plate spline interpolation," *Journal of the Optical Society of America A (JOSA A)*, vol. 28, no. 5, pp. 940–948, 2011.
- [15] A. Hurlbert and T. Poggio, "Synthesizing a color algorithm from examples," *Science*, vol. 239, no. 4839, pp. 482–485, 1988.
- [16] V. Cardei, B. Funt, and K. Barnard, "Estimating the scene illumination chromaticity using a neural network," *Journal of the Optical Society of America A (JOSA A)*, vol. 19, no. 12, pp. 2374–2386, 2002.
- [17] W. Xiong and B. Funt, "Estimating illumination chromaticity via support vector regression," *Journal of Imaging Science and Technology*, vol. 50, p. 341, 2006.
- [18] R. Gershon, A. Jepson, and J. Tsotsos, "From $[r, g, b]$ to surface reflectance: computing color constant descriptors in images," *Perception*, vol. 17, pp. 755–758, 1988.
- [19] B. Singh, W. Freeman, and D. Brainard, "Exploiting spatial and spectral image regularities for color constancy," in *Proc. Workshop on Statistical and Computational Theories of Vision*, 2003.
- [20] A. Chakrabarti, K. Hirakawa, and T. Zickler, "Color constancy beyond bags of pixels," in *Proc. IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, 2008.
- [21] I. Selesnick, "The estimation of laplace random vectors in additive white gaussian noise," *IEEE Transactions on Signal Processing*, vol. 56, no. 8, pp. 3482–3496, 2008.
- [22] F. Sinz, S. Gerwinn, and M. Bethge, "Characterization of the p-generalized normal distribution," *Journal of Multivariate Analysis*, vol. 100, no. 5, pp. 817–820, 2009.
- [23] L. Shi and B. Funt, "Re-processed version of the gehler color constancy dataset of 568 images," accessed from <http://www.cs.sfu.ca/~colour/data/>.
- [24] A. Gijzenij, T. Gevers, and J. van de Weijer, "Computational color constancy: Survey and experiments," *IEEE Transactions on Image Processing*, vol. 33, no. 9, pp. 2475–2489, 2011.
- [25] K. Barnard, "A comparison of computational color constancy algorithm-part 1: Theory and experiments with synthetic data," *IEEE Transactions on Image Processing*, vol. 11, no. 9, pp. 972–984, 2002.
- [26] R. Hogg and E. Tanis, *Probability and statistical inference*. Macmillan New York, 1977.
- [27] S. Hordley and G. Finlayson, "Reevaluation of color constancy algorithm performance," *Journal of the Optical Society of America A (JOSA A)*, vol. 23, no. 5, pp. 1008–1020, 2006.
- [28] F. Ciurea and B. Funt, "A large image database for color constancy research," in *Proc. IS&T/SID Color Imaging Conference*, 2003.
- [29] K. Barnard, L. Martin, A. Coath, and B. Funt, "A comparison of computational color constancy algorithms. ii. experiments with image data," *IEEE Transactions on Image Processing*, vol. 11, no. 9, pp. 985–996, 2002.
- [30] A. Chakrabarti, D. Scharstein, and T. Zickler, "An empirical camera model for internet color vision," in *Proc. British Machine Vision Conference (BMVC)*, 2009.
- [31] K. Saenko, B. Kulis, M. Fritz, and T. Darrell, "Adapting visual category models to new domains," in *Proc. European Conference on Computer Vision (ECCV)*, 2010.
- [32] K. Hirakawa and T. Parks, "Chromatic adaptation and white-balance problem," in *Proc. IEEE Conference on Image Processing (ICIP)*, 2005.
- [33] A. Werner, "The spatial tuning of chromatic adaptation," *Vision Research*, vol. 43, no. 15, pp. 1611–1623, 2003.
- [34] M. Webster, "Human colour perception and its adaptation," *Network: Computation in Neural Systems*, vol. 7, no. 4, pp. 587–634, 1996.
- [35] T. Hansen, M. Olkkonen, S. Walter, and K. Gegenfurtner, "Memory modulates color appearance," *Nature Neuroscience*, vol. 9, no. 11, pp. 1367–1368, 2006.
- [36] M. Olkkonen, T. Hansen, and K. Gegenfurtner, "Color appearance of familiar objects: Effects of object shape, texture, and illumination changes," *Journal of Vision*, vol. 8, no. 5, p. 13, 2008.



Ayan Chakrabarti received the B. Tech and M. Tech degrees in electrical engineering from the Indian Institute of Technology Madras, Chennai, India, in 2006, and the S.M. and Ph.D. degrees in engineering sciences from Harvard University, Cambridge, MA, in 2008 and 2011, respectively.

He is currently a post-doctoral fellow at Harvard University where he works with Todd Zickler as a member of the Harvard Computer Vision Laboratory, and the Graphics Vision and Interaction group. His research focuses on the development of signal processing and statistical methods for applications in computer vision and computational photography.



Keigo Hirakawa (S'00–M'05) received the B.S. degree in electrical engineering from Princeton University, Princeton, NJ, in 2000 with high honors, the M.S. and Ph.D. degrees in electrical and computer engineering from Cornell University, Ithaca, NY, in 2003 and 2005, respectively, and the M.M. degree in jazz performance studies from the New England Conservatory of Music, Boston, MA, in 2006 with high honors.

He was a research associate at Harvard University, Cambridge, MA, from 2006 until 2009. He is currently an assistant professor at the University of Dayton, Dayton, OH. He heads Intelligent Signal Systems Laboratory at the University of Dayton, where the group focuses on statistical signal processing, color image processing, and computer vision. He has received honors from the IEEE, the IS&T, and the IAPR.



Todd Zickler received the B.Eng. degree in honors electrical engineering from McGill University, Montreal, QC, Canada, in 1996 and the Ph.D. degree in electrical engineering from Yale University, New Haven, CT, in 2004, under the direction of P. Belhumeur.

He joined the School of Engineering and Applied Sciences, Harvard University, Cambridge, MA, as an Assistant Professor in 2004 and was appointed John L. Loeb Associate Professor of the Natural Sciences in 2008. He is the Director of the Harvard Computer Vision Laboratory, and his research is focused on modeling the interaction between light and materials and developing algorithms to extract scene information from visual data. His work is motivated by applications in face, object, and scene recognition; image-based rendering; content-based image retrieval; image and video compression; robotics; and human-computer interfaces.

Dr. Zickler is a recipient of the National Science Foundation Career Award and a Research Fellowship from the Alfred P. Sloan Foundation. His research is funded by the National Science Foundation, the Army Research Office, the Office of Naval Research, and the Sloan Foundation. He is a member of the IEEE.