Color Constancy Using Faces

Jiun-Hung Chen, Hwann-Tzong Chen, *Member, IEEE*

Abstract

This paper presents an empirical study on the problem of solving color constancy using faces. We investigate the effectiveness of using the human face in a scene as the source of chromatic information to achieve color constancy. State-of-the-art algorithms are evaluated to see whether the chromatic information provided by faces may help to estimate the scene illuminant. We create a dataset containing 472 pairs of images of 472 people, and each pair includes one image with a neutral-color reference card and the other image without the card. The neutral-color reference card can be used to retrieve the 'ground-truth' illuminant of the scene. The dataset covers diverse scenes with various lighting conditions, and the face regions of the people in the images are manually labeled. Based on the dataset, we show that the color on the face highly correlates with the scene illuminant. The results of empirical evaluation on the dataset suggest that, when face information is available, better performance of color constancy can be achieved by inferring the scene illuminant from the face instead of the whole scene.

I. Introduction

To perceive the true reflectance properties of visible surfaces under different lighting conditions is an important issue in designing effective machine vision systems. However, estimating the illumination and reflectance from a single image is an under-determined inverse problem. Constraints or assumptions need to be made to allow feasible analyses. For example, the Retinex theory [18] aims to address the problem of separating the illumination from the reflectance within a scene in which the reflectance is piecewise constant. Multiple observations and the natural-image statistics can also impose constraints on the decomposition problem [26]. On the other hand, computational color constancy often assumes that the illumination within a scene is approximately uniform. The problem can thus be defined as estimating the *illuminant* in a scene, *e.g.*, [1], [3], [8], [9], [11], [13], [14], [15], [20], [22], [23].

Different assumptions on the scene reflectance have been proposed in the literature of computational color constancy. The widely used *gray-world* algorithm [4] assumes that the spatial average of reflectances in a scene is 'gray'. It follows that the illuminant can be estimated by averaging the observed values in

J.-H. Chen is with Amazon.com; H.-T. Chen is with the Department of Computer Science, National Tsing Hua University, Taiwan.

each color channel. As a limiting case of the gray-world algorithm, a similar assumption based on the Retinex theory [17], [18] yields the *scale-by-max* algorithm [1], which estimates the illuminant by the maximum response in each color channel. Finlayson and Trezzi [10] have shown that the scale-by-max algorithm is equivalent to applying the L_∞ Minkowski-norm to the observed image, and the gray-world algorithm is equivalent to applying the L_1 Minkowski-norm. As a more general version of the gray-world algorithm and based on the assumption that the average of 'reflectance differences' in a scene is achromatic, Van de Weijer and Gevers present the *gray-edge* algorithm [24], which applies L_p norm to the color derivatives instead of the pixel values.

Forsyth's *gamut-mapping* approach [11] and its variants [6], [9] estimate the scene illuminant by computing the mappings from 'colors observed under an unknown illuminant' to 'a gamut of reference colors taken under a known canonical illuminant'. Recently Gijsenij *et al*. [16] further extend the gamut mapping approach such that, in addition to pixel values, higher-order derivatives can also be incorporated into gamut mapping. Their algorithm is able to achieve state-of-the-art performance.

Besides gamut-mapping algorithms, several statistical approaches to color constancy have also shown promising results, *e.g*., [2], [3], [7], [12], [14], [22], [23]. The *color-by-correlation* approach presented by Finlayson *et al*. [7] aims to model the correlation between the illuminant with the image color. The color space is quantized into bins to build histograms, and totally 37 histograms are built for 37 commonly occurring illuminants. Each histogram is associated with a specific illuminant, and each color bin of the histogram records the relative frequency of observing the corresponding color under the associated illuminant. The unknown illuminant of a new scene is then estimated by finding the illuminant that correlates most strongly with the color histogram of the image. Rosenberg *et al*. [22] combine the idea of Bayesian color constancy [3], [12] with nonparametric models (similar to those used in the color-by-correlation approach but based on reflectance histograms rather than illuminant histograms), and obtain very good results of illuminant estimation. Gehler *et al*. [14] further apply the algorithm of [22] to a dataset of more diverse images with accurate illuminant labels. They show that, with such a good dataset for training and evaluation, the Bayesian approach can achieve significant improvement in estimating illuminants.

Although many color-constancy algorithms have been developed, the problem of estimating illuminant is still hard in practice due to the wide variety of illumination and reflectance in real scenes. Collecting image data that cover various observations of different reflectance surfaces under different illuminants might help to build a better color-constancy model. However, since the dataset should be sufficiently diverse for training, it would require great effort to obtain accurate estimates of illuminants as the 'ground truth'. A large dataset comprising about 11,000 images with measured illuminants is presented in [5], and has been used in several studies, *e.g*., [15], [25]. The images are extracted from video clips with a gray sphere appearing in the lower-right field of view for recovering the ambient illuminant. However, as pointed out by Gehler *et al*. [14], the dataset of [5] has limitations that the images are of low resolution. Moreover, nearby frames in the video are highly correlated, and thus only around 600 images extracted from the full set can be considered uncorrelated. Gehler *et al*. further build a dataset containing 568 high-quality images in RAW format. Each image includes a Macbeth color chart for retrieving the ground-truth illuminant. Gehler *et al*. show that color constancy can be effectively learned from such a dataset using a Bayesian approach.

A. Overview

Our aim is to investigate the effectiveness of using the human face in a scene as a source of information to achieve color constancy. We evaluate state-of-the-art algorithms to see whether the color information provided by faces may help to estimate the scene illuminant. The idea of extracting illuminant information from faces for auto white-balance has been mentioned in [19] and [21], but both only use a small number of simple data and lack detailed empirical analysis. In [19], faces are assumed to be taken under standard illuminations and only 10 training images are used. The evaluation of [21] is also limiting and the method requires face recognition to get per-person skin color models. In this paper, we present the first empirical study and performance evaluation of solving color constancy with faces.

We manage to carry out a detailed evaluation with a large dataset of faces in complex scenes, and to ensure that the color measured on a face is indeed a reliable source for estimating the scene illuminant. We collect a dataset containing labeled face images of 472 people. The images are taken under various lighting conditions to cover a wide range of spectra in indoor and outdoor scenes. We use a neutral-color reference card to obtain the 'ground-truth' illuminant of each scene. While modeling the surface reflectance of general objects in a natural scene would require a huge number of data, modeling a single object category (the human face) is much more feasible. To our best knowledge, our dataset is the first one that provides labeled faces with ground-truth illuminants in real scenes.

We use our dataset to evaluate five existing color-constancy algorithms, as well as five regression-based algorithms using different machine learning techniques. The main idea of the regression-based algorithms is to learn a model for skin color distributions under different lighting conditions. We test the performances of different color-constancy algorithms under the settings of including or excluding face information during illuminant estimation. The experimental results show that the chromatic information extracted from faces can improve the accuracy of illuminant estimation. Our dataset can also be used as a benchmark for further evaluation of new color-constancy algorithms and for analysis of skin-color statistics.

(a) (b) (c)

Fig. 1. (a) & (b) A pair of images in our dataset. For each human subject, we take two photos with and without a neutral-color reference card. A part of the reference card is cropped for estimating the ground-truth illuminant, and the region of the face is also manually labeled. (c) A collage of the gray reference cards extracted from our dataset. The blue patches in the last row are mere padding.

II. The Face Dataset

In this section, we describe in detail the face dataset that we build for the evaluation of color-constancy algorithms. The face dataset contains 472 pairs of photos of 472 people. For each human subject, we take two photos: In the first photo the subject is holding a neutral-color reference card near the face, and the second photo is taken without the reference card being shown. The photos with a reference card can be used to generate the ground truth, and the photos without a reference card will be used for testing. The human subjects are from different ethnic groups, and the photos are taken under various indoor/outdoor lighting conditions. We use the WhiBal gray cards as the reference cards, which are certified to have a uniform response distribution under visible spectrum. The photos are captured and stored in RAW format using Canon 1D and Canon 400D cameras. We then export the RAW data as sRGB images with the default white-balance setting using the *dcraw* software. Note that the experiments shown in this paper are done in linear RGB. We convert sRGB to linear RGB when we evaluate various color-constancy algorithms, and henceforth, if not particular specified, RGB refers to linear RGB.

The faces, eyes, and reference cards are manually marked (*e.g.*, Figs. 1a & 1b). To compute the ground-truth illuminant, we assume that only one dominant illuminant is in the scene and the two paired photos are taken under the same lighting condition. We also assume that the lighting on the reference card is roughly constant. The ground-truth illuminant is thus estimated by the mean RGB values $\boldsymbol{\ell} = [\ell_r \; \ell_g \; \ell_b]^\mathsf{T} = \sum_i [y_r(i) \;\; y_g(i) \;\; y_b(i)]^\mathsf{T}$ within the marked region of the reference card, and the rg-chromaticities $\mathbf{c} = [c_r \; c_g]^\mathsf{T} = [\ell_r/(\ell_r + \ell_g + \ell_b) \; \ell_g/(\ell_r + \ell_g + \ell_b)]^\mathsf{T}$ are also used for comparisons in the experiments. We show in Fig. 1c a collage of the reference cards extracted from the 472 images in the dataset. It illustrates the distribution of various scene illuminants. As can be seen, the dataset includes images taken under a wide variety of lighting conditions.

To verify whether the face contains useful information for color constancy, we show in Fig. 2 some

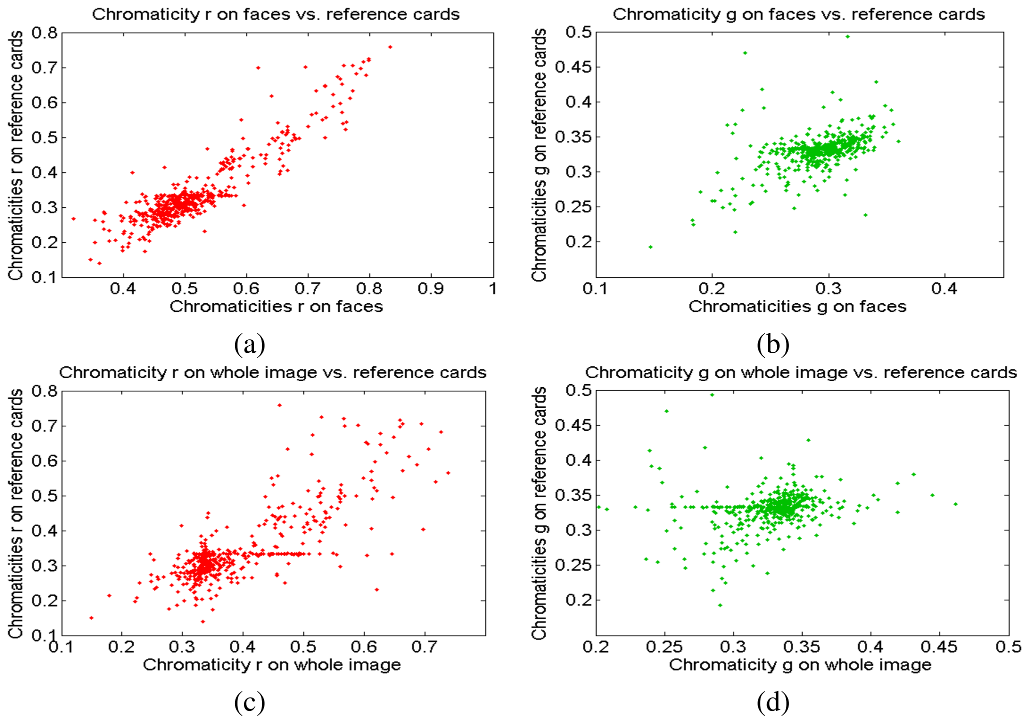

Fig. 2. Comparisons between 'chromaticities on face vs. reference card' (first row) and 'chromaticities measured from the whole image vs. reference card' (second row). (a) The correlation coefficient of 'chromaticity r on face vs. reference card' is 0.91 with p-value 0. (b) The correlation coefficient of 'chromaticity g on face vs. reference card' is 0.50 with p-value 0. (c) The correlation coefficient of 'chromaticity r measured on whole image vs. reference card' is 0.79 with p-value 0. (d) The correlation coefficient of 'chromaticity g measured on whole image vs. reference card' is 0.21 with p-value 0. In sum, the correlation coefficients of 'chromaticities on face vs. reference card' are higher than those of 'chromaticities measured from the whole image vs. reference card'.

comparisons between 'chromaticities on face vs. reference card' and 'chromaticities measured from the whole image vs. reference card'. For this analysis, we ignore the variation of the illumination on each face since we are more interested in the statistics of the entire dataset rather than the properties of individual faces. The results are listed as follows: The correlation coefficient of 'chromaticity r on face vs. reference card' is 0.91, the correlation coefficient of 'chromaticity g on face vs. reference card' is 0.50, the correlation coefficient of 'chromaticity r measured from the whole image vs. reference card' is 0.79, and the correlation coefficient of 'chromaticity g measured from the whole image vs. reference card' is 0.21, with all p-values being 0. Consequently, the correlation coefficients of 'chromaticities on face vs. reference card' are much higher than those of 'chromaticities measured from the whole image vs. reference card'. This result implies that faces could be a stable source of illuminant information for solving color constancy.

A by-product of this evaluation dataset is to provide rectified facial chromaticities for further skin color analysis and detection. We may remove the influence of illumination on facial chromaticities according

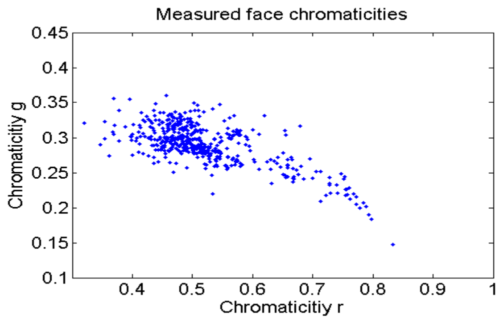

(a)

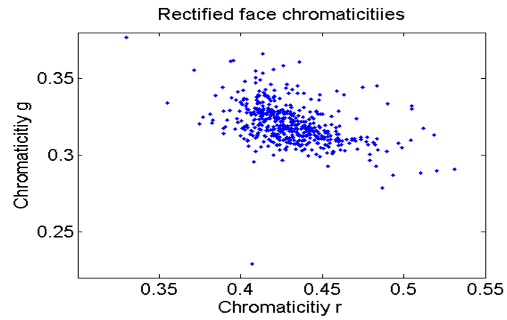

(b)

Fig. 3. (a) The distribution of measured facial chromaticities. (b) The distribution of rectified facial chromaticities.

to the reference card. Fig. 3 shows the comparisons between the distribution of measured facial rg-chromaticities and the distribution of rectified facial rg-chromaticities. The mean and the covariance matrix of the rectified facial rg-chromaticities are

$$\begin{bmatrix} 0.43 \\ 0.32 \end{bmatrix} \quad \text{and} \quad 10^{-3} \times \begin{bmatrix} 0.60 & -0.14 \\ -0.14 & 0.18 \end{bmatrix}. \tag{1}$$

The variance in r-chromaticity is larger than the variance in g-chromaticity, which conforms to our daily experience of human face color.

III. Evaluation

We use our dataset to evaluate five existing color-constancy algorithms and another five regression-based algorithms. The purpose of including regression-based algorithms for evaluation is to see if the mapping from skin color to scene illuminant can be effectively learned by applying standard machine learning techniques to our training data. Overall, we evaluate ten algorithms, which are summarized as follows.

a) Regression-based algorithms:: We may consider color constancy as the problem of estimating the scene illuminant from an input image. Instead of estimating a color illuminant $\hat{\boldsymbol{\ell}} = [\hat{\ell}_r \; \hat{\ell}_g \; \hat{\ell}_b]^\mathsf{T} \in \mathbb{R}^3$, typically the problem is simplified to estimate rg-chromaticities, $\hat{\mathbf{c}} = [\hat{c}_r \; \hat{c}_g]^\mathsf{T} = [\hat{\ell}_r/(\hat{\ell}_r+\hat{\ell}_g+\hat{\ell}_b) \; \hat{\ell}_g/(\hat{\ell}_r+\hat{\ell}_g+\hat{\ell}_b)]^\mathsf{T}$. Let $\mathcal{F} \subseteq \mathbb{R}^D$ be an image feature space and $\mathcal{C} \subseteq \mathbb{R}^2$ be the space of all rg-chromaticities. The color-constancy problem is formulated as learning a regression function $\phi : \mathcal{F} \to \mathcal{C}$ from a given set of N training examples $\mathcal{T} = \{(\mathbf{f}_n, \mathbf{c}_n) \in \mathcal{F} \times \mathcal{C} \,|\, n = 1, \ldots, N\}$. In our experiments, we use the mean RGB values within a facial area and their variances as the image features. Therefore the dimension of the feature space is 6 in our case ($\mathcal{F} \subseteq \mathbb{R}^6$). We choose five standard algorithms for experiments, including *linear regression*, *Gaussian process*, *support vector regression with polynomial kernel*, *support* *vector regression with Gaussian kernel*, and *multilayer perceptrons*.

b) Scale-by-max:: The scale-by-max algorithm (also referred to as white-patch algorithm) estimates the illuminant by the maximum response in each color channel, *i.e.*,

$$\hat{\ell}_r = \max_i y_r(i)\,,\ \hat{\ell}_g = \max_i y_g(i)\,,\ \hat{\ell}_b = \max_i y_b(i)\,, \tag{2}$$

where $y_r(i)$, $y_g(i)$, and $y_g(i)$ are the three color channels of pixel i.

c) Gray-world:: The gray-world algorithm estimates the illuminant by the average value in each color channel:

$$\hat{\ell}_r = \frac{1}{M}\sum_i y_r(i)\,,\ \hat{\ell}_g = \frac{1}{M}\sum_i y_g(i)\,,\ \hat{\ell}_b = \frac{1}{M}\sum_i y_b(i)\,, \tag{3}$$

where M is the number of pixels in an image.

d) Gray-edge:: The gray-edge algorithm is based on the assumption that the light source can be computed from the average color derivative in the image:

$$\hat{\boldsymbol{\ell}} \propto \left(\frac{\int |\mathbf{y_x}|^p d\mathbf{x}}{\int d\mathbf{x}}\right)^{\frac{1}{p}}\,, \tag{4}$$

where $|\mathbf{y_x}| = [|\nabla_\mathbf{x} y_r|, |\nabla_\mathbf{x} y_g|, |\nabla_\mathbf{x} y_b|]^\mathsf{T}$, and p is the order of the Minkowski-norm. In our experiments we choose $p = 1$ and $p = 6$, where $p = 6$ is the suggested value in [24].

e) Generalized gamut mapping and Bayesian color constancy.: We also evaluate two state-of-the-art algorithms: generalized gamut mapping [16] and Bayesian color constancy [14], [22]. Both algorithms can learn from the training data. For generalized gamut mapping, we have tested the zeroth-order and first-order settings. For Bayesian color constancy, we use the ground-truth illuminants to generate the prior, and compute the empirical reflectance distribution based on the assumption of reflectance exchangeability. The details of the two algorithms and their implementations are available in [16] and [14], [22].

A. Experimental Results

Two types of error measures are used for evaluation: *i*) the squared error ρ between the estimated rg-chromaticities and the ground truth, and *ii*) the angular error θ between the estimated illuminant vectors and the ground truth. The squared error is defined by

$$\rho(\mathbf{c}, \hat{\mathbf{c}}) = \|\mathbf{c} - \hat{\mathbf{c}}\|^2\,, \tag{5}$$

TABLE I

THE PERFORMANCES OF COLOR-CONSTANCY ALGORITHMS OVER FIVE RUNS OF THREE-FOLD CROSS VALIDATION. THE ABBREVIATIONS OF THE ALGORITHMS AND SETTINGS ARE SUMMARIZED AS FOLLOWS: GRAY-WORLD (GW), GRAY-EDGE WITH L_1 NORM (GE), GRAY-EDGE WITH L_6 NORM (GE6), WHITE-PATCH (WP), BAYESIAN COLOR CONSTANCY (BCC), ZEROTH-ORDER GAMUT MAPPING (GM0), FIRST-ORDER GAMUT MAPPING (GM1), LINEAR REGRESSION (LR), GAUSSIAN PROCESS (GP), SUPPORT VECTOR REGRESSION WITH POLYNOMIAL KERNEL (SVRP), SUPPORT VECTOR REGRESSION WITH GAUSSIAN KERNEL (SVRG), AND MULTILAYER PERCEPTRONS (MLP). THE AVERAGE SQUARED ERROR ρ IS IN AN ORDER OF 10^{-4} AND THE AVERAGE ANGULAR ERROR θ IS MEASURED IN DEGREES. THE BEST PERFORMANCE IS IN BOLDFACE AND THE SECOND BEST ONE IS IN ITALIC. THE GRAY-WORLD, GRAY-EDGE, AND WHITE-PATCH ALGORITHMS CAN ONLY USE THE WHOLE IMAGE AS THE INPUT DUE TO THEIR ASSUMPTIONS.

	GW	GE	GE6	WP	BCC	GM0	GM1	LR	GP	MLP	SVRP	SVRG
Whole image (ρ)	103	72	112	96	135	94	98	*67*	**62**	161	73	85
Face (ρ)	—	—	—	—	58	75	84	*49*	**34**	52	52	78
Whole image (θ)	9.6	8.1	9.4	8.5	7.2	10.9	8.6	9.0	*6.8*	12.3	**6.8**	7.3
Face (θ)	—	—	—	—	7.2	9.4	10.8	6.4	**5.1**	6.9	*6.0*	7.2

where $\mathbf{c}$ is the vector of the ground-truth rg-chromaticities and $\hat{\mathbf{c}}$ contains the estimated rg-chromaticities. Both $\mathbf{c}$ and $\hat{\mathbf{c}}$ are assumed to be normalized. The angular error is given by

$$\theta(\boldsymbol{\ell}, \hat{\boldsymbol{\ell}}) = \frac{180 * \arccos(\boldsymbol{\ell}^{\mathsf{T}} \hat{\boldsymbol{\ell}})}{\pi} , \tag{6}$$

where $\boldsymbol{\ell}$ is the ground-truth illuminant and $\hat{\boldsymbol{\ell}}$ is the estimated illuminant. Note that the illuminant is represented as an RGB vector, and can be easily derived from the rg-chromaticities using

$$\hat{\boldsymbol{\ell}} = [\hat{c}_r \;\; \hat{c}_g \;\; (1 - \hat{c}_r - \hat{c}_g)]^{\mathsf{T}} . \tag{7}$$

The five regression-based algorithms can learn the mapping of scene illuminant using either the face regions or the whole images as the training data. We thus perform two experiments for each regression-based algorithm, one with the faces and the other one with the whole images. The two state-of-the-art algorithms, Bayesian color constancy and generalized gamut mapping, can also infer the illuminant using face color as the input, and therefore we also evaluate these two algorithms using the face regions and the whole images. Furthermore, for generalized gamut mapping, we have tested the zeroth-order and first-order settings. The gray-world, gray-edge, and white-patch algorithms can only use the whole image as the input due to their assumptions. We have tested gray-edge with the settings of L_1 norm and L_6 norm. For each experiment on a given algorithm with a specified setting, we perform five runs of three-fold cross validation. Alternately, two thirds of the data with the ground-truth illuminants are used for training, and the remaining one third of the data are used for prediction. Each image will be used for prediction once during each run of three-fold cross validation. The average squared error and the average angular error over the five runs are are summarized in Table I. Note that, for gray-world, gray-edge, and white-patch, the results over the five runs should be the same since they do not require training, and therefore we only perform one run on the whole dataset for these algorithms.

From the results shown in Table I, we may observe that most of the color-constancy algorithms tested in this paper can achieve better performance by using the faces as the input than by using the whole images as the input. Among the results of all settings (with the whole images, with the faces, and using different parameters), the best results are produced by the regression-based algorithms, in particular Gaussian process, with the face information being used. The significant improvement in the performance of Bayesian color constancy with faces indicates that learning from a specific prior is more suitable for the Bayesian framework. To learn a more general reflectance distribution of the whole scene would need a lot more training data, which is not easy in real situation. Since the color on faces has a high correlation with the scene illuminant as shown previously in Section II, it is easier to model the reflectance distribution by the color histograms of the faces than those of the whole scenes. The inference with the learned face color prior can also be more accurately performed.

To further investigate the performances of the algorithms on individual inputs, we plot the error curves of the algorithms and compare the results. Some example results are shown in Figs. 4- 7. For easier comparison and better visualization, we sort the errors of different methods and plot them in groups. In each figure, the x-axis is the rank of an image and the y-axis is its squared error or angular error. Each curve shows the sorted errors of an algorithm in an increasing order. We use the result of Gaussian process with faces as the indicator and plot it in every graph in tandem with the results of other algorithms for comparison. As can be seen in these figures, about 10% of the images in the dataset yield very large error values for all algorithms. Their illuminants are hard to estimate due to the dim yellow light. Nevertheless, the overall performances still imply that color-constancy algorithms may benefit from the chromatic information of faces. The regression-based algorithms and Baysian color constancy can effectively learn and use the color model of faces to infer the scene illuminants. On the other hand, gamut mapping does not seem to gain much from the face information, especially for the eaiser cases (the first 90%, see Fig. 6). It might imply that the 'gamut of faces' cannot be stably modeled via convex optimization. We believe that the performance of gamut mapping can be improved if the method of modeling the gamut is specifically adjusted for faces.

Furthermore, to check whether the idea of solving color constancy with faces can be applied to real scenarios, we repeat the experiments of Gaussian process regression, but this time we perform the prediction phase of cross validation using the face regions detected by the OpenCV face detector instead of using the manually labeled bounding boxes. Among the 472 images of our data, the OpenCV face detector finds faces in 467 images (359 true positives + 108 false positives) and finds nothing in 5 images. We define a true positive as one that has a least 50% overlap with the manually labeled bounding

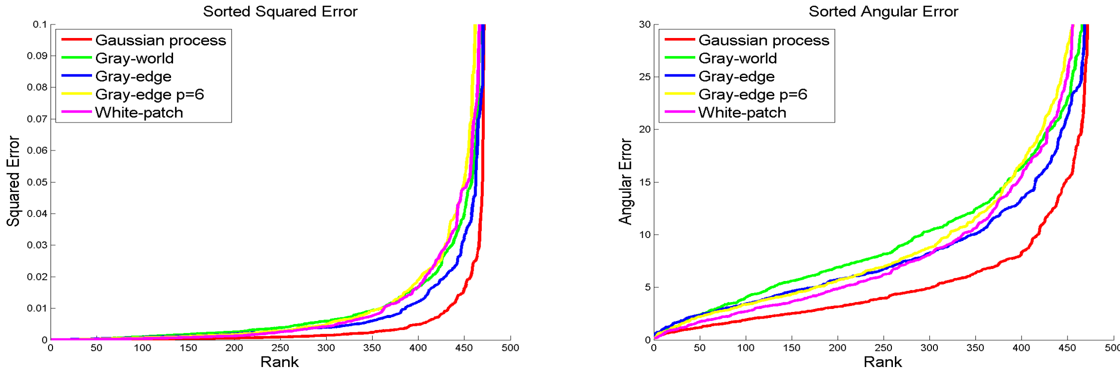

Fig. 4. The error comparisons among Gaussian process with faces, gray-world using the whole images, gray-edge using the whole images, gray-edge (L_6 norm) using the whole images, and white-patch using the whole images.

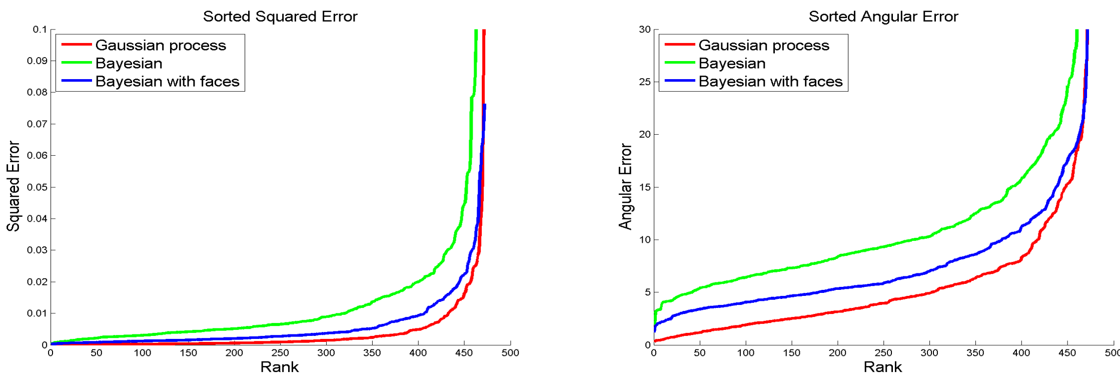

Fig. 5. The error comparisons among Gaussian process with faces, Bayesian color constancy using the whole images, and Bayesian color constancy with faces.

box. Note that if the OpenCV face detector finds nothing, we use the whole image as input. The squred error obtained by Gaussian process is 0.0042 and the angular error is 5.7. The results are a bit worse than those for the manually labeled data, but are still better than the results of estimating the illuminant using the whole image. This experiment suggests that estimating scene illuminants from faces is robust to the variations in the locations of faces and should be useful in real applications.

IV. Conclusions

This paper presents an evaluation of solving color constancy using faces. We test the idea of modeling the color on human faces for illuminant estimation, in contrast to modeling the reflectance of general objects in natural scenes. We have shown that the color on the face has a high correlation with the scene illuminant, and the chromatic information provided by the face can be used to improve the performances

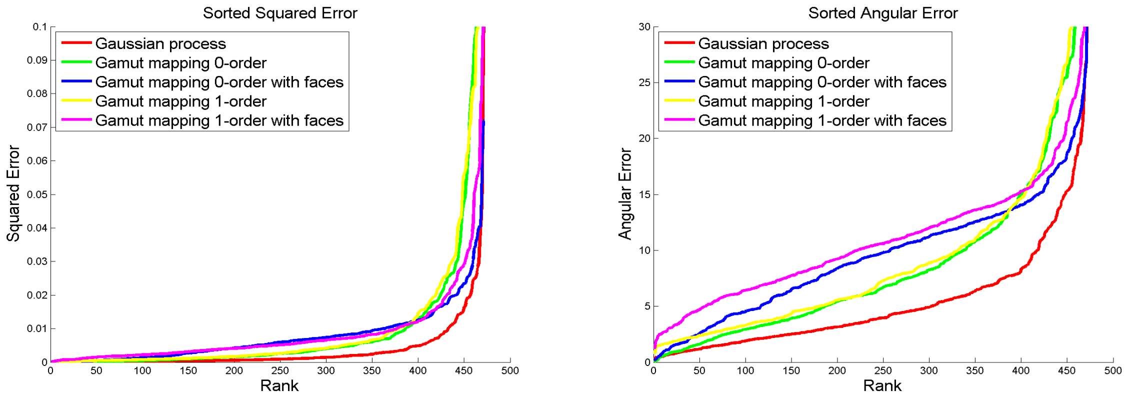

Fig. 6. The error comparisons among Gaussian process with faces, zeroth-order gamut mapping using the whole images, zeroth-order gamut mapping with faces, first-order gamut mapping using the whole images, and first-order gamut mapping with faces.

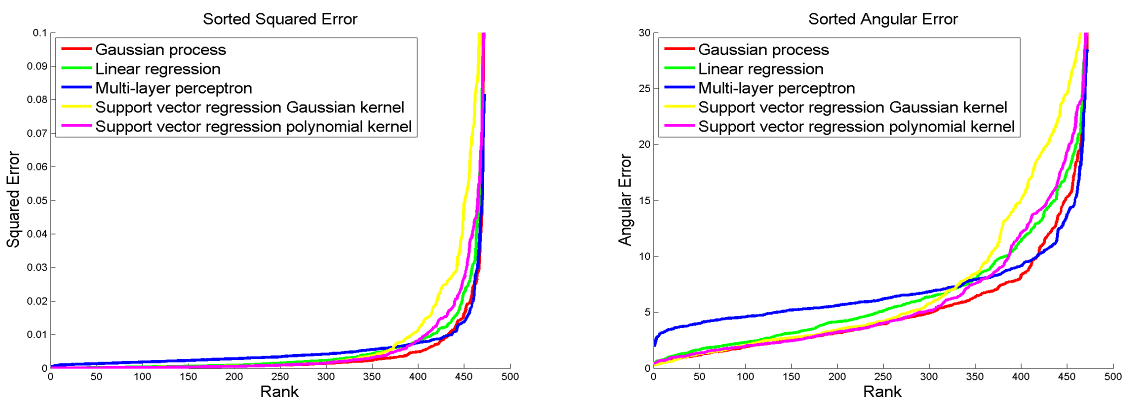

Fig. 7. The error comparisons among Gaussian process with faces, linear regression with faces, multilayer perceptrons with faces, support vector regression by Gaussian kernel with faces, and support vector regression by polynomial kernel with faces.

of color-constancy algorithms. The experimental results indicate that using faces as the input improves the performances of Bayesian color constancy and regression-based algorithms, in terms of the average squared error and the average angular error. Since the function of face detection is available on most digital cameras, smartphones, and webcam systems, it should be promising and suitable for practical use if the face information is employed to facilitate auto white-balance on those devices. In addition to the use as a benchmark for evaluating color-constancy algorithms, our dataset might also be helpful for skin-color analysis. Finally, in this work we assume that only one dominant illuminant is present in a photo. How to deal with multiple illuminants would be an interesting and challenging problem for further research.

References

[1] K. Barnard, L. Martin, A. Coath, and B. V. Funt. A comparison of computational color constancy algorithms. ii. experiments with image data. *IEEE Transactions on Image Processing*, 11(9):985–996, 2002.

[2] K. Barnard, L. Martin, and B. V. Funt. Colour by correlation in a three-dimensional colour space. In *ECCV (1)*, pages 375–389, 2000.

[3] D. Brainard and W. Freeman. Bayesian color constancy. *J. Opt. Soc. Am. A*, 14(7):1393–1411, 1997.

[4] G. Buchsbaum. A spatial processor model for object colour perception. *Journal of the Franklin Institute*, 310(1):1– 26, 1980.

[5] F. Ciurea and B. V. Funt. A large image database for color constancy research. In *Color Imaging Conference*, pages 160–164, 2003.

[6] G. D. Finlayson and S. D. Hordley. A theory of selection for gamut mapping color constancy. In *CVPR*, pages 60–65, 1998.

[7] G. D. Finlayson, S. D. Hordley, and P. M. Hubel. Colour by correlation: A simple, unifying approach to colour constancy. In *ICCV*, pages 835–842, 1999.

[8] G. D. Finlayson, S. D. Hordley, and P. M. Hubel. Color by correlation: A simple, unifying framework for color constancy. *IEEE Trans. Pattern Anal. Mach. Intell.*, 23(11):1209–1221, 2001.

[9] G. D. Finlayson, S. D. Hordley, and I. Tastl. Gamut constrained illuminant estimation. *International Journal of Computer Vision*, 67(1):93–109, 2006.

[10] G. D. Finlayson and E. Trezzi. Shades of gray and colour constancy. In *Color Imaging Conference*, pages 37–41, 2004.

[11] D. A. Forsyth. A novel algorithm for color constancy. *International Journal of Computer Vision*, 5(1):5–35, 1990.

[12] W. T. Freeman and D. H. Brainard. Bayesian decision theory, the maximum local mass estimate, and color constancy. In *ICCV*, pages 210–217, 1995.

[13] B. Funt, V. Cardei, and K. Barnard. Learning color constancy. In *Proc. IS&T/SID Fourth Color Imaging Conference: Color Science, Systems and Applications*, pages 58–60, 1996.

[14] P. V. Gehler, C. Rother, A. Blake, T. P. Minka, and T. Sharp. Bayesian color constancy revisited. In *CVPR*, 2008.

[15] A. Gijsenij and T. Gevers. Color constancy using natural image statistics. In *CVPR*, 2007.

[16] A. Gijsenij, T. Gevers, and J. van de Weijer. Generalized gamut mapping using image derivative structures for color constancy. *International Journal of Computer Vision*, 86(2-3):127–139, 2010.

[17] E. H. Land. The retinex theory of color vision. *Sci Am.*, 237(6):108–128, 1977.

[18] E. H. Land and J. J. McCann. Lightness and retinex theory. *Journal of the Optical Society of America*, 61(1):1–11, 1971.

[19] J.-W. Lee and K.-I. Sohng. Auto white balance using chromatic coordinates of detected human faces. In *The 23rd International Technical Conference on Circuits/Systems, Computers and Communications Technical Program*, 2008.

[20] L. Maloney and B. Wandell. Color constancy: A method for recovering surface spectral reflectance. *JOSA-A*, 3:29–33, 1986.

[21] J. Montojo. Face-based chromatic adaptation for tagged photo collections, 2009. http://www.cs.toronto.edu/ jmontojo/papers/chromatic09.pdf.

[22] C. R. Rosenberg, T. P. Minka, and A. Ladsariya. Bayesian color constancy with non-gaussian models. In *NIPS*, 2003.

[23] H. J. Trussell and M. J. Vrhel. Estimation of illumination for color correction. In *IEEE ICASSP*, volume 4, pages 2513–2516, 1991.

[24] J. van de Weijer and T. Gevers. Color constancy based on the grey-edge hypothesis. In *ICIP (2)*, pages 722–725, 2005.

[25] J. van de Weijer, C. Schmid, and J. J. Verbeek. Using high-level visual information for color constancy. In *ICCV*, pages 1–8, 2007.

[26] Y. Weiss. Deriving intrinsic images from image sequences. In *ICCV*, pages 68–75, 2001.