Illuminant Estimation using Ensembles of Multivariate Regression Trees

Peter van Beek
Cheriton School of Computer Science
University of Waterloo
vanbeek@cs.uwaterloo.ca

R. Wayne Oldford
Department of Statistics & Actuarial Science
University of Waterloo
rwoldford@uwaterloo.ca

Abstract

White balancing is a fundamental step in the image processing pipeline. The process involves estimating the chromaticity of the illuminant source and using the estimate to correct the image to remove any color cast. Given the importance of the problem, there has been much previous work on illuminant estimation. Previous work is either more accurate but slow and complex, or fast and simple but less accurate. In this paper, we propose a method for illuminant estimation that uses (i) fast features known to be predictive in illuminant estimation and (ii) single feature decision boundaries in ensembles of multivariate regression trees, (iii) each of which has been constructed to minimize a multivariate distance measure appropriate for illuminant estimation. The result is an illuminant estimation method that is simultaneously fast, simpler, and more accurate.

1 Introduction

White balancing is either performed onboard the camera—if the image delivered by the camera is in JPEG format, for example—or in a post-processing phase—if the image delivered by the camera is in the camera's native RAW format. When done onboard the camera, real-time and space considerations place additional restrictions on white balancing algorithms. Most commercial cameras use simple algorithms based on the gray-world assumption [13], such as using the mean color of the image as an estimate of the illumination, as these algorithms have the advantage of being fast and simple, albeit at the expense of accuracy [11].

Given the importance of the problem, there has been much previous work on illuminant estimation (e.g., [3, 8, 15, 16, 18, 19, 21, 33, 34]; see [22] for a recent survey). Presently, convolutional neural networks (CNN) can provide state-of-the-art accuracy [4, 32]. However, in terms of simplicity and speed, the CNN approach is at a definite disadvantage. To get credible speed performance requires a GPU to process each image. Even then, a CNN can take approximately three seconds on a GPU to process a single image with present state-of-the-art methods [32]. While it is unclear whether this waiting time is tolerable for post-processing, it is certainly the case that the slow speed and the need to add specialized hardware is unacceptable for onboard a camera.

The most successful methods to date in illuminant estimation are based on regression and minimize the squared error loss when *constructing* their predictive models [11, 32]. However, when *evaluating* the performance of the predictive models, the squared error loss is not used and distance measures more appropriate for illuminant estimation, such as angular error, are used instead.

In this paper, we propose a method for illuminant estimation that uses (i) fast features known to be predictive in illuminant estimation and (ii) single feature decision boundaries in ensembles of multivariate regression trees, (iii) each of which has been *constructed* to minimize a multivariate distance measure appropriate for illuminant estimation. The result is an illuminant estimation method that is fast, simple, and accurate. Figure 1 summarizes the improvement in accuracy over Cheng et al. [11].

The rest of the paper proceeds as follows. In Section 2, we review distance or performance measures that have been proposed for evaluating the effectiveness of white balancing algorithms. In Section 3, we present our proposed method of constructing predictive regression models based on minimizing an appropriate distance measure. In Section 4,

Figure 1: Percentage improvement over Cheng et al. [11] for the (a) mean, (b) median, (c) mean of the best 25%, and (c) mean of the worst 25% of the distances measures for an image set, where each image set is captured by a single model of a camera, and there are 11 image sets and 5 distance measures for each image set.

we experimentally evaluate our proposed method, comparing against the most successful methods to date [11, 32]. Section 5 concludes the paper.

2 Distance Measures

Finlayson and Zakizadeh [17] have shown that the ranking of white balancing algorithms changes depending on the chosen distance or performance measure. The distance measures assume *normalized* RGB; i.e., $r = R/(R+G+B)$, $g = G/(R+G+B)$, and $b = B/(R+G+B)$, where R, G, and B are the red, green, and blue channel measurements and $r + g + b = 1$. In what follows, let all vectors be row vectors, $\hat{\mathbf{e}} = (\hat{r}, \hat{g}, \hat{b})$ be the estimated illuminant, and $\mathbf{e} = (r, g, b)$ be the ground truth illuminant.

The most widely used distance measure [20, 25], *recovery angular error*, measures the angular distance between the estimated illuminant and the ground truth illuminant,

$$dist_\theta(\hat{\mathbf{e}}, \mathbf{e}) = \cos^{-1}\left(\frac{\hat{\mathbf{e}} \cdot \mathbf{e}^T}{\|\hat{\mathbf{e}}\|\|\mathbf{e}\|}\right),$$

where $\hat{\mathbf{e}} \cdot \mathbf{e}^T$ is the dot product of the vectors, $\|\cdot\|$ is the Euclidean norm, and we assume the angular distance is measured in degrees. Gijsenij et al. [20] note that human observers sometimes judge the recovery angular error underestimates the perceived differences between two images. One reason is that the recovery angular error ignores the direction of the deviation from the ground truth, which can be important from a perceptual point of view.

Finlayson and Zakizadeh [17] propose *reproduction angular error*,

$$dist_{rep}(\hat{\mathbf{e}}, \mathbf{e}) = \cos^{-1}\left(\frac{r/\hat{r} + g/\hat{g} + b/\hat{b}}{\sqrt{(r/\hat{r})^2 + (g/\hat{g})^2 + (b/\hat{b})^2}\ \sqrt{3}}\right),$$

as a distance measure. In words, once the image has been corrected with the estimated illuminant, the angular error between the white balanced ground truth and the target of uniform gray $\mathbf{1} = (1, 1, 1)$ is determined.

Gijsenij et al. [20] discuss the use of the Minkowski distance for measuring the distance between $\hat{\mathbf{e}} = (\hat{r}, \hat{g}, \hat{b})$ and $\mathbf{e} = (r, g, b)$. Two special cases are the *Taxicab distance*,

$$dist_1(\hat{\mathbf{e}}, \mathbf{e}) = |\hat{r} - r| + |\hat{g} - g| + |\hat{b} - b|,$$

and the *Euclidean distance*,

$$dist_2(\hat{\mathbf{e}}, \mathbf{e}) = \sqrt{(\hat{r} - r)^2 + (\hat{g} - g)^2 + (\hat{b} - b)^2}.$$

Gijsenij et al. [20] note that the Euclidean distance treats each of the RGB channels uniformly whereas it is known that the sensitivity of the human eye to perceived differences varies across color channels. To this end, they define the *perceptual Euclidean error*,

$$dist_{ped}(\hat{\mathbf{e}}, \mathbf{e}) = \sqrt{w_r(\hat{r} - r)^2 + w_g(\hat{g} - g)^2 + w_b(\hat{b} - b)^2},$$

where weights w_r, w_g, and w_b capture this sensitivity and $w_r + w_g + w_b = 1$. In experiments, Gijsenij et al. [20] found that the weight vector $(w_r, w_g, w_b) = (0.21, 0.71, 0.08)$ gave a higher correlation with the judgment of human observers than that of $dist_\theta$, $dist_1$, and $dist_2$.

3 Our Proposal

In this section, we present our approach for white balancing based on ensembles of multivariate regression trees. We begin with a brief review of univariate (ordinary) regression trees and how they were applied by Cheng et al. [11] for white balancing.

3.1 Univariate regression trees

Univariate regression trees are constructed in a greedy, top-down manner from a set of labeled training examples $\{(\mathbf{x}_1, y_1), \ldots, (\mathbf{x}_n, y_n)\}$, where $\mathbf{x}_i = (x_{i1}, \ldots, x_{im})$ is a vector of feature values and y_i is a scalar response variable (see, e.g., [7, 29, 23]). As it is sufficient for our purposes, we assume that the feature values and the response variable are real-valued. The root node of the tree is associated with all the training examples. At each step in the construction of the tree, the training examples at a node are partitioned by choosing the feature $j \in \{1, \ldots, m\}$ and partition p that minimizes the total of the squared-error loss functions,

$$\operatorname*{argmin}_{c_1} \sum_{i \in L(j,p)} (y_i - c_1)^2 \; + \; \operatorname*{argmin}_{c_2} \sum_{i \in R(j,p)} (y_i - c_2)^2, \tag{1}$$

where the partition is into two subsets, $L(j, p)$ and $R(j, p)$, formed by branching on feature j according to $x_{ij} \leq p$ and $x_{ij} > p$, respectively. For any choice of feature j and partition p the minimization is solved by taking the scalar c_1 to be the mean of the y_i values in the left branch and the scalar c_2 to be the mean of the y_i values in the right branch. Once the best pair (j, p) is found, a left child node and right child node are added to the tree and are associated with the subsets $L(j, p)$ and $R(j, p)$. The partitioning continues until some stopping criterion is met, in which case the node is a leaf and is labeled with the scalar c associated with the subset of examples at the node.

To estimate an illuminant $\mathbf{e} = (r, g, b)$, Cheng et al. [11] predict the r chromaticity and the g chromaticity independently using separate univariate trees fit with the squared-error loss function and from these two values the estimate

3

of the b chromaticity can be determined using $b = 1 - r - g$. Cheng et al. [11] use four pairs of simple features in their univariate trees: (f_r^1, f_g^1), the mean color chromaticity as provided by the gray-world algorithm; (f_r^2, f_g^2), the brightest color chromaticity, an adaptation of the white-patch algorithm; (f_r^3, f_g^3), the bin average of the mode of the RGB histogram, and (f_r^4, f_g^4), the mode of the kernel density estimate from the normalized chromaticity plane.

Two important points are that Cheng et al.'s [11] method (i) predicts the r and g chromaticities *independently*, and (ii) minimizes the distance measure of interest—in their work, the recovery angular error—*indirectly* by minimizing the squared-error loss function. Our starting point is (i) the observation that our response variable is not a scalar but a vector $\mathbf{e} = (r, g, b)$ of chromaticities, and (ii) the related observation that independently fitting using the squared-error loss function is not necessarily a good surrogate for minimizing distance measures used in white balancing; Example 1 illustrates the point.

Example 1. *Let* $\mathbf{e} = (r, g, b)$ *be the ground truth illuminant. Consider the two estimates of the illuminant* $\hat{\mathbf{e}}_1 = (\hat{r}_1, \hat{g}_1, \hat{b}_1)$ *and* $\hat{\mathbf{e}}_2 = (\hat{r}_2, \hat{g}_2, \hat{b}_2)$*, where*

$$
\begin{aligned}
\hat{r}_1 &= r + \alpha, & \hat{r}_2 &= r + \alpha, \\
\hat{g}_1 &= g + \alpha, & \hat{g}_2 &= g - \alpha, \\
\hat{b}_1 &= 1 - \hat{r}_1 - \hat{g}_1 = b - 2\alpha, & \hat{b}_2 &= 1 - \hat{r}_2 - \hat{g}_2 = b,
\end{aligned}
$$

and $\alpha > 0$ *is some residual error. Here, the pair of estimates* $\hat{r}_1$ *and* $\hat{r}_2$ *and the pair of estimates* $\hat{g}_1$ *and* $\hat{g}_2$ *both have equal squared error* (α^2)*. Thus, from the point of view of independently fitting using the squared-error loss function, the two estimates of the illuminant have equal error. However, for distance measures that have been proposed for white balancing, the error of the estimate* $\hat{\mathbf{e}}_1$ *is larger (and can be much larger) than that of* $\hat{\mathbf{e}}_2$*. For example, let* $\mathbf{e} = (\frac{1}{3}, \frac{1}{3}, \frac{1}{3})$ *and* $\alpha = 0.05$*. For the recovery angular error,* $dist_\theta(\hat{\mathbf{e}}_1, \mathbf{e}) = 11.977$ *and* $dist_\theta(\hat{\mathbf{e}}_2, \mathbf{e}) = 6.983$*. Thus, minimizing the squared error is not necessarily a good surrogate for minimizing a distance measure of interest in white balancing.*

3.2 Multivariate regression trees

We show how multivariate regression trees [30, 12, 27], where each tree predicts multiple responses, can be used to effectively estimate an illuminant. In the case of multiple responses, multivariate trees are more compact than univariate trees and can be more accurate when the response variables are correlated [28]. Multivariate regression trees are constructed in a greedy, top-down manner from a set of labeled training examples $\{(\mathbf{x}_1, \mathbf{y}_1), \ldots, (\mathbf{x}_n, \mathbf{y}_n)\}$, where as before $\mathbf{x}_i$ is a vector of feature values but now $\mathbf{y}_i = (y_{i1}, \ldots, y_{ik})$ is a *vector* of response variables [30, 12, 27].

Our proposed method for illuminant estimation is for a single tree to simultaneously predict all three chromaticity components of an illuminant $\mathbf{e} = (r, g, b)$, rather than multiple trees predicting them independently, thus taking into account that the chromaticity components of the illuminant are correlated and constrained (see Figure 2). We fit a multivariate regression tree by directly minimizing a distance measure (loss function) of interest. Previous work on multivariate regression trees minimize only loss functions that are variants of the squared-error loss function [12, 30]. In our proposed method, at each step in the construction of the multivariate regression tree, the training examples at a node are partitioned by choosing a feature $j \in \{1, \ldots, m\}$ and partition p that minimizes the total of the loss functions,

$$
\operatorname*{argmin}_{\hat{\mathbf{e}}_1} \sum_{i \in L(j,p)} dist(\hat{\mathbf{e}}_1, \mathbf{e}_i) \; + \; \operatorname*{argmin}_{\hat{\mathbf{e}}_2} \sum_{i \in R(j,p)} dist(\hat{\mathbf{e}}_2, \mathbf{e}_i), \tag{2}
$$

where the partition is into two subsets, $L(j, p)$ and $R(j, p)$, formed by branching on feature j according to $x_{ij} \leq p$ and $x_{ij} > p$, respectively; $dist(\cdot, \cdot)$ is any distance measure that has been proposed for white balancing (see Section 2); and $\hat{\mathbf{e}}_1, \hat{\mathbf{e}}_2$, and $\mathbf{e}_i$ are normalized RGB vectors in $\mathbb{R}^3$, i.e., the chromaticities sum to one. The partitioning continues until some stopping criteria is met, in which case the node is a leaf and labeled with the estimated illuminant $\hat{\mathbf{e}}$ associated with the subset of examples at the node. When such a tree is used to estimate the illuminant of an image that has not been seen before, one starts at the root and repeatedly tests the feature at a node and follows the appropriate branch until a leaf is reached. The label of the leaf is the estimated illuminant of the image.
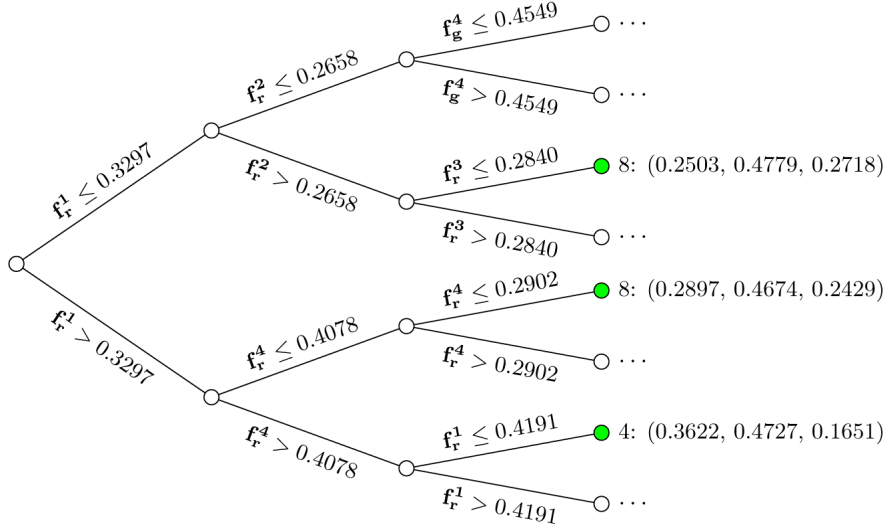
4

Figure 2: An example multivariate regression tree when the distance measure is recovery angular error, showing only the first four layers (out of 18 layers). The three (green) leaf nodes are labeled with the number of training examples at that node and the illuminant estimation if that node is reached. Also shown are the four images associated with the green bottom-most leaf node; the images have been gamma corrected for presentation.

Ensembles of multivariate regression trees. To improve predictive accuracy, an ensemble of trees is learned where each tree is a multivariate regression tree. Various methods have been proposed for constructing ensembles including manipulating the training examples, manipulating the input features, injecting randomness into the learning algorithm, and combinations of these methods (see [14]). We adopt the method for constructing ensembles that injects randomness into the construction of an individual tree. When constructing the tree, each feature j partition p pair considered in Equation 2 is accumulated along with the total of the loss functions associated with the pair. Rather than taking the (j, p) pair that minimizes the total of the loss functions, a (j, p) pair is chosen at random from all pairs within a given percentage of the minimum. This ensures that diverse yet accurate trees are constructed. We also experimented with bagging [5] and Cheng et al.'s method [11], two methods that manipulate the training examples to construct diverse trees, and with random forests [24, 6], a method that manipulates both the training examples and the input features. However, in each case the alternative led to a decrease in accuracy.

Given a new image, the individual predictions of the trees in the ensemble are combined into a single estimated illuminant $\hat{e}$ for the image as follows (see Figure 3). Let $S = \{\hat{e}_1, \ldots, \hat{e}_k\}$ be the set of predictions from the individual trees. The final estimated illuminant is the normalized RGB vector $\hat{e} = (\hat{r}, \hat{g}, \hat{b})$ that minimizes $\sum_{i \in S} dist(\hat{e}, \hat{e}_i)$; i.e., the same minimization problem as the one that is solved in partitioning the training examples at a node and in labeling a leaf when constructing the trees (Equation 2).

Fitting multivariate regression trees. In fitting a multivariate regression tree to a set of labeled training examples, the idea is to repeatedly, in a greedy and top-down manner, find the best feature j and partition p that results in the largest drop in the total error, as measured by a given distance measure (see Equation 2). The minimization in Equation 2 is a constrained, non-linear optimization problem. The problem must be solved many times when learning the ensemble of trees (approximately 40,000 calls to the minimization routine is typical for constructing a single tree) and also once every time the ensemble is applied to a new image and the individual predictions of the trees are combined to obtain a final estimated illuminant. Note that learning the trees is an offline process that occurs once while applying the ensemble to a new image is an online process that will occur many times, as it occurs each time an image is captured by the camera.

The minimization problem can be solved using *exact* but sophisticated and computationally expensive numerical optimization routines such as those based on interior-point or sequential quadratic programming methods. Exact methods are suitable for offline learning of a single ensemble of multivariate trees, but unfortunately are too slow
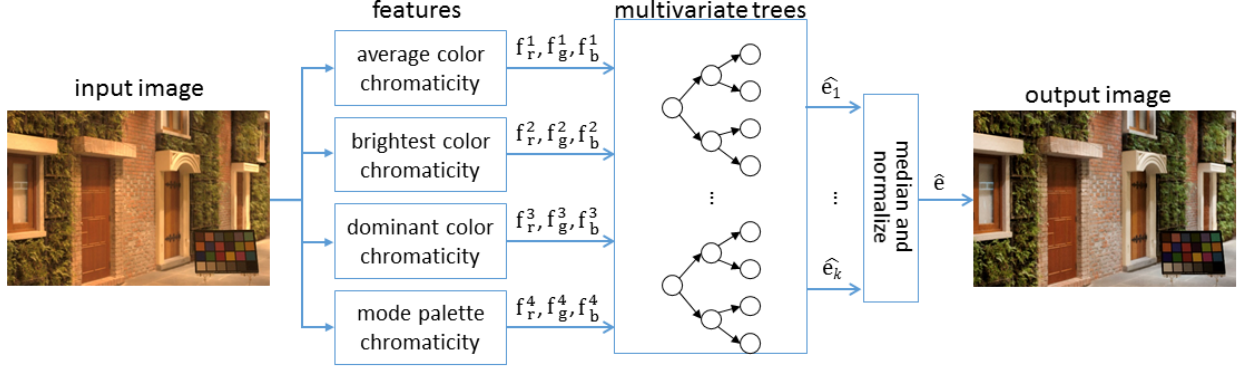
Figure 3: Our adaptation of Cheng et al.'s [11] method. Given an input image, four triples of feature values are calculated, the feature values are used by each tree in the ensemble to predict the illumination, and the individual predictions $\{\hat{\mathbf{e}}_1, \ldots, \hat{\mathbf{e}}_k\}$ are combined into a single estimate $\hat{\mathbf{e}}$ of the illuminant that is used to white balance the image.

for experimental evaluation where cross-validation and repeated trials are necessary to obtain accurate estimates of performance. This impediment also arises in univariate trees when using loss functions other than squared error (see, e.g., [23, p. 342]). As well, sophisticated exact methods are unlikely to be suitable for online application of the ensemble onboard the camera, given the real-time and space considerations.

Instead, we propose to solve the minimization problem *approximately* by taking the median of the ground truth illuminants of the training examples at a node and normalizing so that the RGB values sum to one (see Example 2). The method is simple and fast. As well, the method often finds the exact solution and otherwise finds good quality approximate solutions for distance measures that arise in white balancing (see supplemental material). Further, solving the minimization problem approximately does not appear to have a significant impact on the accuracy of the multivariate regression trees, perhaps because the tree construction itself is a greedy, and hence approximate, process.

Example 2. *Consider the follow three training examples, where the ground truth illuminants are shown but feature values are not shown, and let* $dist_2$ *be the distance measure.*

example	r	g	b
$\mathbf{e}_1$	0.4285	0.4468	0.1247
$\mathbf{e}_2$	0.4221	0.4473	0.1306
$\mathbf{e}_3$	0.4098	0.4682	0.1220

The median of the training examples gives (0.4221, 0.4473, 0.1247) and normalizing so that the RGB values sum to one results in the illuminant estimate (0.4246, 0.4500, 0.1254) with an associated cost of 3.5134. The mean is (0.42012, 0.45411, 0.12577) with an associated cost of 3.7619. The illuminant estimate that exactly minimizes the sum of the distance measure over the training examples is (0.4228, 0.4487, 0.1285) with an associated cost of 3.4049.

4 Experimental Evaluation

We compare a MATLAB implementation of our multivariate regression tree method to Cheng et al.'s [11] univariate regression tree method and Shi et al.'s [32] CNN method[1].

We used the following image sets in our experiments. The *SFU Laboratory image set* consists of images of objects in a laboratory setting captured by a video camera under 11 different illuminants [1]. Following previous work, we use a common subset of 321 images (the minimal specularities and non-negligible dielectric specularities subset). The *Gehler-Shi image set* consists of 568 indoor and outdoor images captured by two different cameras [18, 31]. As done

---

[1]The software is available at: To be filled in.

in previous work, we used the reprocessed version of the image set which starts from a camera RAW image file that has been minimally processed by the camera and creates a linearly processed, lossless 12-bit PNG file using the well-known dcraw program. The *NUS 8-camera image set* consists of 1736 images captured by eight different cameras [9], where in most cases each camera has photographed the same scene. Each image is captured as a minimally processed camera RAW image file and is linearly processed to create a lossless 16-bit PNG file. The *NUS-Laboratory 8-camera image set* is a complementary laboratory set of 840 images captured with the same cameras. These additional images correct for a bias towards images captured outdoors in daylight [10].

4.1 Experimental methodology

To assess our approach, we followed Cheng et al.'s [11] original experimental setup closely. In particular, we used the following methodology in our experimental evaluation.

Training data generation. Cheng et al. [11] use four pairs of simple features in their univariate trees (see the description in Section 3.1) and in our main set of experiments, we used the same features. In addition, to each of these pairs of features we add a blue component f_b^i, $i = 1, \ldots 4$. Although redundant, these additional features come for free and significantly improve the accuracy of our multivariate trees. In both approaches, training and testing is done on each camera separately; i.e., a predictive model is built for a particular model of camera rather than a generic model that can be used by any camera. To construct the machine learning data for a camera, each image taken by the camera in an image set is processed by (i) normalizing the image to a [0, 1] image, using the saturation level and darkness level of the camera, (ii) masking out the saturated pixels and the color checker, if present, (iii) determining the feature values for the image, and (iv) labeling the example using the ground truth illuminant. The result is a set of labeled training examples $\{(\mathbf{x}_1, \mathbf{y}_1), \ldots, (\mathbf{x}_n, \mathbf{y}_n)\}$, where $\mathbf{x}_i = (f_r^1, f_g^1, f_b^1, \ldots, f_r^4, f_g^4, f_b^4)$ is a vector of feature values and $\mathbf{y}_i = (r_i, g_i, b_i)$ is the ground truth illuminant.

Parameter selection. In both approaches, parameters were set on 1/3 of the Gehler-Shi image set using $dist_\theta$ as the distance measure. The parameters were then fixed for all other cameras, image sets, and distance measures. Cheng et al. [11] set the following parameters: (i) the number of trees in an ensemble ($30 \times 4 \times 2 = 240$ trees, where each feature pair is used to build a tree for predicting r and a tree for predicting g, and this is repeated 30 times), (ii) the amount of overlap and the number of slices of the training data used in constructing the ensemble, and (iii) a threshold value for determining the consensus of the ensemble. In our method, we set the following parameters: (i) the number of trees in an ensemble (30 trees), (ii) the amount of randomization used in constructing the ensemble (10%), and (iii) a threshold value, where a node is partitioned only if the average error at the node is greater than the threshold (0.5).

As well, Cheng et al. [11] use the MATLAB routine `fitrtree` for fitting a tree using the squared-error loss function. The routine has two parameters that were used at their default values: each branch node in the tree has at least "MinParentSize = 10" examples and each leaf has at least "MinLeafSize = 1" examples per tree leaf. These values are used as a stopping criteria when building the tree. We used the same stopping criteria in our implementation for fitting a tree using a chosen white balancing distance measure.

Performance evaluation. We used standard k-fold cross validation, to evaluate and compare the accuracy of our method and Cheng et al.'s [11] method. In k-fold cross validation, the image set is randomly partitioned into k approximately equal folds and each of the folds is, in turn, used as a testing set and the remaining $k - 1$ folds are used as a training set. In our experiments we used 10-fold cross validation, as 10-fold is the most widely recommended, especially for our setting where the amount of data per camera is limited in some regions of the output space [26, 2, 23]. To reduce variance, the statistics we report are the result of performing 30 runs of 10-fold cross validation with different random seeds. Both methods share the same random seeds so that the partitions into training and test are the same for each algorithm for each experiment. Although for space considerations we only directly compare against Cheng et al.'s [11] method, we note that results comparing [11] to many other algorithms can be found in their original paper (the results reported here differ somewhat to those reported in [11] on common image sets as that paper uses 3-fold cross validation and reports results based on only a single run).

4.2 Experimental results

We compared the approaches on accuracy, simplicity, and speed. All experiments were performed on a PC with an Intel i7-6700K, 4GHz running MATLAB R2016a.

Accuracy. Figure 1 shows a comparison of the accuracy of our method of ensembles of multivariate regression trees against Cheng et al.'s [11] method on the five distance measures discussed in Section 2. On these image sets, our method gives comparable accuracy as measured by the mean of the worst 25% errors and a significant improvement in accuracy as measured by the mean, median, and mean of the best 25% errors. The current best-performing CNN approach [32] reports percentage improvements over Cheng et al. [11] of 5.1% for the mean, 8.2% for the median, 2.0% for the mean of the best 25%, and 4.7% for the mean of the worst 25%, where the improvement is measured over the geometric means of the eight cameras on the NUS 8-camera image set. As a point of comparison, on the combined NUS and NUS-Laboratory 8-camera image set, we achieve an improvement of 3.5% for the mean, 11.1% for the median, 24.7% for the mean of the best 25%, and -0.6% for the mean of the worst 25%, again as measured over the geometric means. (Adding one more feature approaches the CNN's accuracy for the mean of the worst 25%; see supplemental material.)

Simplicity. The only source of complexity of trees is their size and number. For the Canon 1Ds Mark III camera, we recorded the size of each tree and the size of each ensemble of trees, as measured by the number of nodes. Our method leads to ensembles that are an order of magnitude smaller. For example, averaging over 30 runs, Cheng et al.'s method has approximately 207 nodes per tree and 49,575 nodes per ensemble compared to 143 nodes and 4,317 nodes for our method, when using the $dist_\theta$ distance measure.

Speed. For the processing of an image, our method and Cheng et al.'s method are for practical purposes identical, ours being negligibly faster due to having fewer and shallower trees. As noted by Cheng et al. [11], processing an image takes 0.25 seconds or less and such speeds are comparable to simple gray-world algorithms, as used in many commercial cameras [13]. In contrast, Shi et al. [32] state that the CNN processing of an image takes approximately 3 seconds *on a GPU*.

5 Conclusion

Multivariate regression trees, where each tree predicts multiple responses, can be used to effectively estimate an illuminant for white balancing an image. In our method a multivariate regression tree is fit by directly minimizing a distance measure of interest. We show empirically that overall our method leads to improved performance on diverse image sets. Our ensembles of multivariate regression trees are more accurate and simpler, while inheriting the fast run-time of previous work.

References

[1] K. Barnard, L. Martin, A. Coath, and B. Funt. A comparison of computational color constancy algorithms II: Experiments with image data. *IEEE Trans. on Image Processing*, 11:985–996, 2002.

[2] Y. Bengio and Y. Grandvalet. No unbiased estimator of the variance of k-fold cross-validation. *Journal of Machine Learning Research*, 5:1089–1105, 2004.

[3] S. Bianco, G. Ciocca, C. Cusano, and R. Schettini. Automatic color constancy algorithm selection and combination. *Pattern Recognition*, 43:695–705, 2010.

[4] S. Bianco, C. Cusano, and R. Schettini. Color constancy using CNNs. *CoRR*, abs/1504.04548, 2015.

[5] L. Breiman. Bagging predictors. *Machine Learning*, 24:123–140, 1996.

[6] L. Breiman. Random forests. *Machine Learning*, 45:5–32, 2001.

[7] L. Breiman, J. Friedman, C.J. Stone, and R.A. Olshen. *Classification and Regression Trees*. CRC press, 1984.

[8] V. C. Cardei, B. Funt, and K. Barnard. Estimating the scene illumination chromaticity by using a neural network. *JOSA A*, 19:2374–2386, 2002.

[9] D. Cheng, D. K. Prasad, and M. S. Brown. Illuminant estimation for color constancy: Why spatial-domain methods work and the role of the color distribution. *JOSA A*, 31:1049–1058, 2014.

[10] D. Cheng, B. Price, S. Cohen, and M. S. Brown. Beyond white: Ground truth colors for color constancy correction. In *Proceedings of the IEEE International Conference on Computer Vision*, pages 298–306, 2015.

[11] D. Cheng, B. Price, S. Cohen, and M. S. Brown. Effective learning-based illuminant estimation using simple features. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, pages 1000–1008, 2015.

[12] G. De'Ath. Multivariate regression trees: A new technique for modeling species-environment relationships. *Ecology*, 83(4):1105–1117, 2002.

[13] Z. Deng, A. Gijsenij, and J. Zhang. Source camera identification using auto-white balance approximation. In *Proceedings of the IEEE International Conference on Computer Vision*, pages 57–64, 2011.

[14] T. G. Dietterich. Ensemble methods in machine learning. In *Proceedings of the International Workshop on Multiple Classifier Systems*, pages 1–15, 2000. Available as: Springer Lecture Notes in Computer Science 1857.

[15] G. D. Finlayson. Corrected-moment illuminant estimation. In *Proceedings of the IEEE International Conference on Computer Vision*, pages 1904–1911, 2013.

[16] G. D. Finlayson and E. Trezzi. Shades of gray and colour constancy. In *Proceedings of the Color and Imaging Conference*, 2004.

[17] G. D. Finlayson and R. Zakizadeh. Reproduction angular error: An improved performance metric for illuminant estimation. In *Proceedings of the British Machine Vision Conference*, 2014.

[18] P. V. Gehler, C. Rother, A. Blake, T. Minka, and T. Sharp. Bayesian color constancy revisited. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, pages 1–8, 2008.

[19] A. Gijsenij and T. Gevers. Color constancy using natural image statistics and scene semantics. *IEEE Trans. on Pattern Analysis and Machine Intelligence*, 33:687–698, 2011.

[20] A. Gijsenij, T. Gevers, and M. P. Lucassen. Perceptual analysis of distance measures for color constancy algorithms. *JOSA A*, 26:2243–2256, 2009.

[21] A. Gijsenij, T. Gevers, and J. van de Weijer. Generalized gamut mapping using image derivative structures for color constancy. *IJCV*, 86:127–139, 2010.

[22] A. Gijsenij, T. Gevers, and J. van de Weijer. Computational color constancy: Survey and experiments. *IEEE Trans. on Image Processing*, 20:2475–2489, 2011.

[23] T. Hastie, R. Tibshirani, and J. Friedman. *The Elements of Statistical Learning: Data mining, Inference and Prediction*. Springer, 2nd edition, 2009.

[24] T. K. Ho. Random decision forests. In *Proceedings of the 3rd International Conference on Document Analysis and Recognition*, pages 278–282, 1995.

[25] S. D. Hordley and G. D. Finlayson. Reevaluation of color constancy algorithm performance. *JOSA A*, pages 1008–1020, 2006.

[26] R. Kohavi. A study of cross-validation and bootstrap for accuracy estimation and model selection. In *Proceedings of the International Joint Conference on Artificial Intelligence*, pages 1137–1145, 1995.

[27] D. R. Larsen and P. L. Speckman. Multivariate regression trees for analysis of abundance data. *Biometrics*, 60(2):543–549, 2004.

[28] W.-Y. Loh and W. Zheng. Regression trees for longtitudinal and multiresponse data. *The Annals of Applied Statistics*, 7(1):495–522, 2013.

[29] J. R. Quinlan. Induction of decision trees. *Machine learning*, 1(1):81–106, 1986.

[30] M. R. Segal. Tree structured methods for longitudinal data. *J. American Statistical Association*, 87:407–418, 1992.

[31] L. Shi and B. Funt. Re-processed version of the Gehler color constancy dataset of 568 images. `http://www.cs.sfu.ca/~colour/data`.

[32] Wu Shi, Chen Change Loy, and Xiaoou Tang. Deep specialized network for illuminant estimation. In *Proceedings of the European Conference on Computer Vision*, pages 371–387, 2016.

[33] H. Vaezi Joze and M. Drew. Exemplar-based colour constancy and multiple illumination. *IEEE Trans. on Pattern Analysis and Machine Intelligence*, 36:860–873, 2014.

[34] W. Xiong and B. Funt. Estimating illumination chromaticity via support vector regression. *Journal of Imaging Science and Technology*, 50(4):341–348, 2006.

- Appendix A: Approximation error

- Appendix B: Detailed experimental results (I)

- Appendix C: Detailed experimental results (II)

A   Approximation error

Our proposed method is to solve *approximately* the minimization problem posed in Equation 2 for fitting multivariate trees (see Section 3.4). In this appendix, we report on experiments investigating the quality of the approximate solutions.
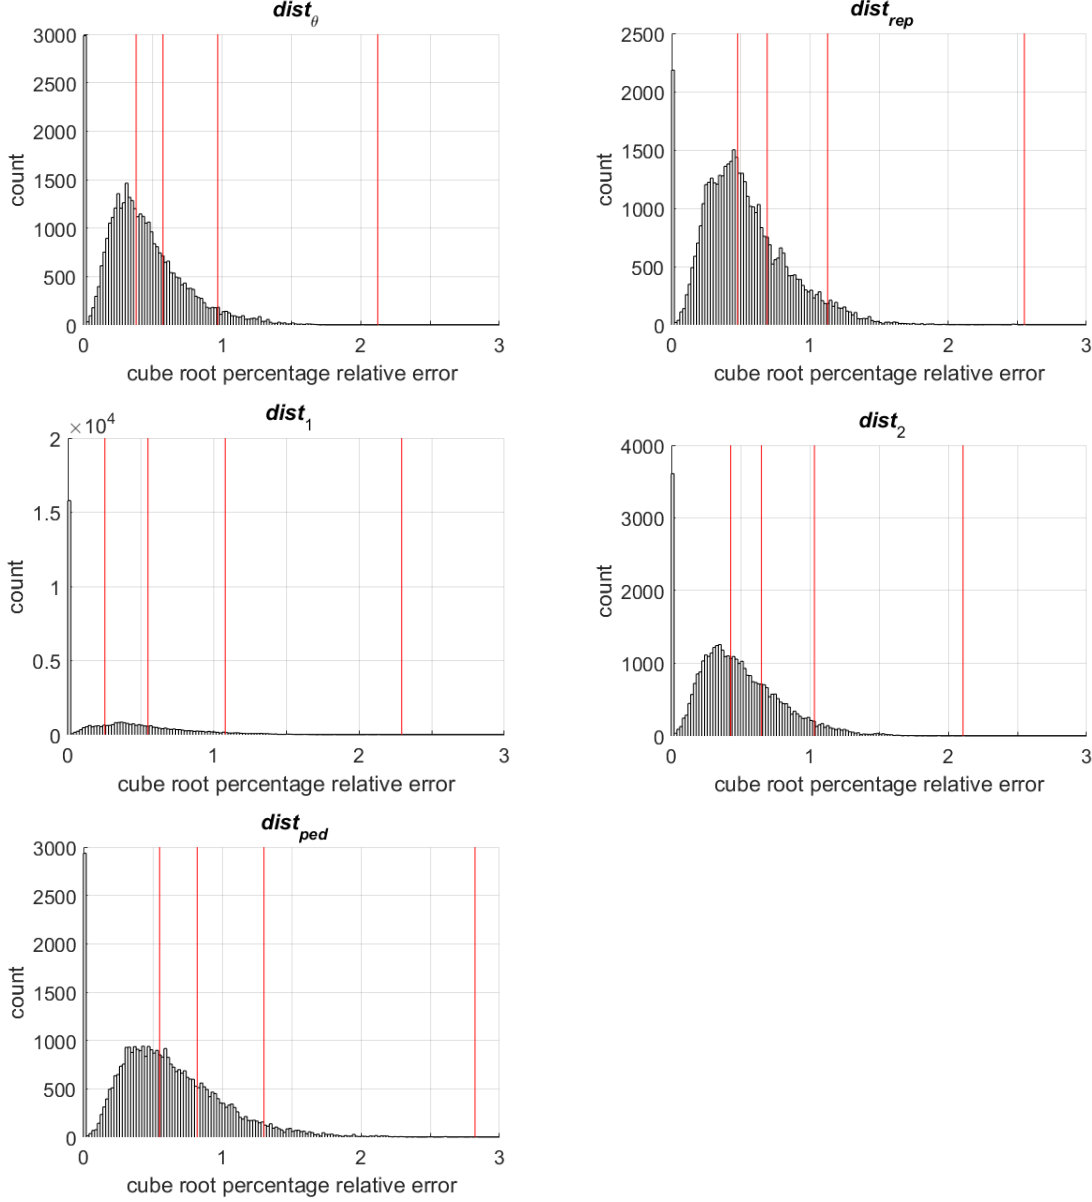
Figure 4: Percent relative approximation error on each call to the minimization routine when building one tree for the Canon EOS-1Ds Mark III camera. The red bars show the median, 75th percentile, 95th percentile, and the maximum. For visualization purposes, the results were transformed by taking the cube root of the percentage relative error.

For determining the exact solution to the minimization problems we used the MATLAB routine `fmincon`. Figure 4 shows how close to optimal the approximate solutions were for the various distance measures. The majority of the approximate solutions are within 1% of optimal.

B Detailed experimental results (I)

Due to space restrictions, we summarize our experimental results in the main body of the paper by providing histograms of the percentage improvements over previous work. In this appendix, we provide detailed results comparing the accuracy of our method to Cheng et al.'s method [11], for the three benchmark suites—with a total of eleven different cameras—and the five distance measures. For uniformity of presentation, the results for $dist_1$, $dist_2$, and $dist_{ped}$ are multiplied by 100. A shaded box indicates that one method was more accurate than the other. If viewing in colour, cyan indicates our method is more accurate and magenta indicates Cheng et al. [11] is more accurate.

Table 1: SFU Laboratory [1] image set.

meas.	meth.	mean	med.	tri-mean	best 25%	worst 25%
Sony DXC-930, 321 images						
$dist_\theta$	Cheng	4.24	2.24	2.85	0.35	11.29
	Ours	3.88	1.79	2.45	0.24	10.72
$dist_{rep}$	Cheng	4.62	2.68	3.24	0.43	11.95
	Ours	4.31	2.27	2.77	0.29	11.74
$dist_1 \times 100$	Cheng	7.48	4.01	5.03	0.63	19.89
	Ours	6.90	3.40	4.33	0.42	19.17
$dist_2 \times 100$	Cheng	4.76	2.57	3.21	0.40	12.63
	Ours	4.38	2.17	2.75	0.27	12.13
$dist_{ped} \times 100$	Cheng	2.07	1.18	1.45	0.19	5.35
	Ours	1.95	0.99	1.29	0.13	5.29

Table 2: Gehler-Shi [31, 18] image set.

meas.	meth.	mean	med.	tri-mean	best 25%	worst 25%
Canon 1D, 86 images						
$dist_\theta$	Cheng	3.74	2.83	2.91	0.74	8.30
	Ours	3.58	2.50	2.79	0.78	8.08
$dist_{rep}$	Cheng	4.62	3.59	3.73	0.89	10.10
	Ours	4.45	3.32	3.60	0.90	9.72
$dist_1 \times 100$	Cheng	5.97	4.66	4.78	1.21	13.03
	Ours	5.72	4.05	4.50	1.21	12.83
$dist_2 \times 100$	Cheng	3.94	3.03	3.11	0.78	8.70
	Ours	3.79	2.71	2.99	0.83	8.45
$dist_{ped} \times 100$	Cheng	1.58	1.23	1.27	0.36	3.40
	Ours	1.53	1.15	1.22	0.36	3.37
Canon 5D, 482 images						
$dist_\theta$	Cheng	2.18	1.37	1.54	0.35	5.42
	Ours	1.99	1.18	1.36	0.28	5.05
$dist_{rep}$	Cheng	2.85	1.77	2.00	0.42	7.13
	Ours	2.59	1.53	1.76	0.35	6.64
$dist_1 \times 100$	Cheng	3.57	2.30	2.56	0.58	8.74
	Ours	3.27	1.97	2.25	0.46	8.20
$dist_2 \times 100$	Cheng	2.33	1.47	1.64	0.37	5.77
	Ours	2.12	1.26	1.45	0.30	5.37
$dist_{ped} \times 100$	Cheng	0.96	0.63	0.70	0.17	2.30
	Ours	0.89	0.56	0.63	0.15	2.19

Table 3: *Part I:* Combined NUS [9] and NUS-Laboratory 8-camera [10] image sets.

meas.	meth.	mean	med.	tri-mean	best 25%	worst 25%
Canon EOS-1Ds Mark III, 364 images						
$dist_\theta$	Cheng	2.13	1.45	1.61	0.36	5.06
	Ours	2.00	1.30	1.46	0.26	4.98
$dist_{rep}$	Cheng	2.81	1.89	2.14	0.46	6.71
	Ours	2.57	1.57	1.85	0.35	6.42
$dist_1 \times 100$	Cheng	3.56	2.41	2.68	0.61	8.42
	Ours	3.34	2.08	2.40	0.45	8.28
$dist_2 \times 100$	Cheng	2.29	1.56	1.73	0.39	5.42
	Ours	2.13	1.36	1.55	0.29	5.29
$dist_{ped} \times 100$	Cheng	1.00	0.70	0.76	0.19	2.32
	Ours	0.96	0.62	0.70	0.14	2.34
Canon EOS 600D, 305 images						
$dist_\theta$	Cheng	2.21	1.55	1.67	0.45	5.12
	Ours	2.21	1.51	1.67	0.32	5.25
$dist_{rep}$	Cheng	2.92	2.04	2.23	0.57	6.76
	Ours	2.90	1.95	2.19	0.43	6.87
$dist_1 \times 100$	Cheng	3.65	2.55	2.79	0.76	8.44
	Ours	3.67	2.55	2.78	0.55	8.66
$dist_2 \times 100$	Cheng	2.37	1.67	1.80	0.49	5.48
	Ours	2.37	1.67	1.81	0.36	5.59
$dist_{ped} \times 100$	Cheng	1.01	0.72	0.78	0.23	2.32
	Ours	1.00	0.70	0.76	0.16	2.36
Fujifilm X-M1, 301 images						
$dist_\theta$	Cheng	2.17	1.35	1.54	0.31	5.39
	Ours	2.13	1.21	1.44	0.23	5.54
$dist_{rep}$	Cheng	2.94	1.77	2.07	0.40	7.37
	Ours	2.86	1.60	1.95	0.31	7.45
$dist_1 \times 100$	Cheng	3.57	2.22	2.55	0.52	8.85
	Ours	3.52	2.00	2.41	0.39	9.13
$dist_2 \times 100$	Cheng	2.34	1.45	1.66	0.34	5.81
	Ours	2.30	1.30	1.57	0.25	5.99
$dist_{ped} \times 100$	Cheng	1.04	0.66	0.75	0.17	2.54
	Ours	1.01	0.58	0.69	0.12	2.64
Nikon D5200, 305 images						
$dist_\theta$	Cheng	2.07	1.30	1.47	0.39	5.02
	Ours	2.02	1.16	1.36	0.30	5.09
$dist_{rep}$	Cheng	2.92	1.89	2.11	0.49	7.11
	Ours	2.85	1.67	1.96	0.40	7.26
$dist_1 \times 100$	Cheng	3.51	2.21	2.52	0.67	8.47
	Ours	3.44	2.02	2.36	0.53	8.64
$dist_2 \times 100$	Cheng	2.27	1.44	1.62	0.43	5.48
	Ours	2.21	1.30	1.52	0.33	5.56
$dist_{ped} \times 100$	Cheng	1.03	0.66	0.76	0.23	2.43
	Ours	0.99	0.59	0.70	0.16	2.45

Table 4: *Part II:* Combined NUS [9] and NUS-Laboratory 8-camera [10] image sets.

meas.	meth.	mean	med.	tri-mean	best 25%	worst 25%
Olympus E-PL6, 313 images						
$dist_\theta$	Cheng	1.77	1.11	1.25	0.26	4.38
	Ours	1.79	1.02	1.18	0.23	4.62
$dist_{rep}$	Cheng	2.45	1.49	1.70	0.36	6.15
	Ours	2.50	1.37	1.61	0.32	6.54
$dist_1 \times 100$	Cheng	2.96	1.86	2.10	0.45	7.29
	Ours	3.01	1.73	1.99	0.39	7.76
$dist_2 \times 100$	Cheng	1.92	1.20	1.36	0.29	4.75
	Ours	1.96	1.11	1.28	0.26	5.06
$dist_{ped} \times 100$	Cheng	0.87	0.56	0.63	0.15	2.11
	Ours	0.88	0.50	0.58	0.12	2.25
Panasonic Lumix DMC-GX1, 308 images						
$dist_\theta$	Cheng	2.22	1.41	1.57	0.36	5.47
	Ours	2.02	1.12	1.32	0.23	5.30
$dist_{rep}$	Cheng	3.02	1.92	2.17	0.45	7.44
	Ours	2.75	1.48	1.82	0.29	7.22
$dist_1 \times 100$	Cheng	3.65	2.37	2.63	0.62	8.89
	Ours	3.30	1.90	2.22	0.38	8.52
$dist_2 \times 100$	Cheng	2.39	1.53	1.70	0.40	5.88
	Ours	2.17	1.23	1.44	0.25	5.64
$dist_{ped} \times 100$	Cheng	1.03	0.68	0.75	0.21	2.44
	Ours	0.94	0.58	0.65	0.13	2.39
Samsung NX2000, 307 images						
$dist_\theta$	Cheng	2.16	1.38	1.52	0.39	5.33
	Ours	2.12	1.30	1.47	0.33	5.32
$dist_{rep}$	Cheng	2.96	1.85	2.07	0.49	7.38
	Ours	2.87	1.72	1.99	0.41	7.27
$dist_1 \times 100$	Cheng	3.69	2.39	2.63	0.68	9.03
	Ours	3.62	2.22	2.50	0.56	9.08
$dist_2 \times 100$	Cheng	2.36	1.52	1.68	0.44	5.79
	Ours	2.31	1.42	1.62	0.37	5.76
$dist_{ped} \times 100$	Cheng	1.08	0.73	0.80	0.23	2.55
	Ours	1.06	0.71	0.76	0.19	2.54
Sony SLT-A57, 373 images						
$dist_\theta$	Cheng	1.99	1.37	1.49	0.38	4.68
	Ours	1.84	1.11	1.30	0.28	4.58
$dist_{rep}$	Cheng	2.76	1.88	2.05	0.49	6.55
	Ours	2.53	1.50	1.78	0.37	6.31
$dist_1 \times 100$	Cheng	3.41	2.35	2.58	0.66	7.99
	Ours	3.17	1.89	2.24	0.49	7.86
$dist_2 \times 100$	Cheng	2.18	1.50	1.64	0.42	5.12
	Ours	2.02	1.20	1.42	0.31	5.01
$dist_{ped} \times 100$	Cheng	0.98	0.68	0.75	0.21	2.26
	Ours	0.93	0.55	0.66	0.16	2.27

C Detailed experimental results (II)

As can be seen in Figure 1 of the main paper, our method significantly improves the accuracy as measured by the mean, the median, and the mean of the best 25% errors—the statistics that are arguably the most important in white balancing. However, an examination of why our method was comparable but did not improve the mean of the worst 25% errors suggested that our method was running out of predictive features when building the multivariate trees. We thus pursued a secondary set of experiments where we added one additional triple of features: $(f_r^{sg}, f_g^{sg}, f_b^{sg})$, the color chromaticity as provided by the shades-of-gray algorithm [16]. The shades-of-gray algorithm is based on the l_p-norm, for $p = 5$. Figure 5 summarizes the results and it can be seen that the addition of the feature improved the performance of our method on the worst cases, albeit at the expense of reducing slightly the improvement on the median and slowing down the processing time of each image.


Figure 5: Percentage improvement of our method over Cheng et al. [11] when adding the feature shades-of-gray for (a) mean, (b) median, (c) mean of the best 25%, and (c) mean of the worst 25% of the distances measures for an image set, where each image set is captured by a single model of a camera, and there are 11 image sets and 5 distance measures for each image set.

The current best-performing CNN approach [32] reports percentage improvements over Cheng et al. [11] of 5.1% for the mean, 8.2% for the median, 2.0% for the mean of the best 25%, and 4.7% for the mean of the worst 25%, where the improvement is measured over the geometric means of the eight cameras on the NUS 8-camera image set. As a point of comparison, on the combined NUS and NUS-Laboratory 8-camera image set, with the added feature we achieve an improvement of 5.5% for the mean, 9.9% for the median, 23.4% for the mean of the best 25%, and 2.9% for the mean of the worst 25%.