

Sensor sharpening for computational color constancy

Kobus Barnard

Computer Division, University of California, Berkeley, Berkeley, California 94720-1776

Florian Ciurea and Brian Funt

School of Computing Science, Simon Fraser University, Burnaby, BC, Canada V5A 1S6

Received May 7, 2001; accepted May 16, 2001

Sensor sharpening [J. Opt. Soc. Am. A **11**, 1553 (1994)] has been proposed as a method for improving computational color constancy, but it has not been thoroughly tested in practice with existing color constancy algorithms. In this paper we study sensor sharpening in the context of viable color constancy processing, both theoretically and empirically, and on four different cameras. Our experimental findings lead us to propose a new sharpening method that optimizes an objective function that includes terms that minimize negative sensor responses as well as the sharpening error for multiple illuminants instead of a single illuminant. Further experiments suggest that this method is more effective for use with several known color constancy algorithms.

© 2001 Optical Society of America

OCIS codes: 330.1690, 330.1720, 100.2000, 100.2960.

1. INTRODUCTION

Sensor sharpening refers to using a linear transformation of camera responses [red–green–blue (RGB)] to improve computational color constancy methods based on independently scaling the three sensor channels.¹ Here the scaling, or diagonal map, is applied to all the sensor responses of a scene under one illuminant to estimate the corresponding responses under a standard, canonical illuminant. Using the appropriate linear transformation before the scaling and its inverse after the scaling reduces the rms mapping error. The original work¹ thus suggested that some color constancy algorithms could, in theory, benefit by sensor sharpening. However, the practicality of sharpening remains unclear because determining the sharpening transform requires knowledge of the scene illuminant (database sharpening), strong conditions on scene illuminant and scene surfaces (perfect sharpening), or optimizations that are not directly related to color constancy error (sensor-based sharpening). Furthermore, the error in scene illuminant parameter estimation is often much larger than the reduction in error that is possible to achieve by sharpening, even assuming that the illumination is known. Finally, color constancy processing is not necessarily independent of sharpening. For example, sharpening may help color constancy processing above and beyond the possible reduction in mapping error by improving the estimation of the required mapping. Alternatively, sharpening may work against the color constancy algorithm.

In preliminary work we applied sharpening to several algorithms² and found that sharpening could improve performance in a few cases. The difficulties we encountered led us to propose a new sharpening method (multiple-illuminant-with-positivity sharpening), which we describe here. We also analyze the relationship be-

tween sharpening and color constancy error for several color constancy algorithms. We then provide detailed results using an improved experimental protocol.

The efficacy of sensor sharpening is expected to be dependent on the camera sensors. Thus we provide results for four different cameras: a Sony DXC-930 CCD video camera,³ a Kodak DCS-460 digital camera, a Kodak DCS-200 digital camera,⁴ and a Kodak DCS-420 digital camera.⁴ Our general conclusion is that when the sensors are already relatively sharp (e.g., the Sony camera), further sensor sharpening is not helpful and can have a detrimental effect. However, when the sensors are not sharp (e.g., the DCS-460, the DCS-200, and the DCS-420), sensor sharpening can have a substantial positive effect, depending on the algorithm, especially when the proposed new method is used.

2. SENSOR-SHARPENING FORMULATION

We begin with an explanation of sensor sharpening.¹ The motivation for sensor sharpening is the observation that many of the most effective current color constancy algorithms employ a diagonal model of illumination change. To understand this model, consider a white patch under two different illuminants. Suppose that under the first illuminant the color is (r, g, b) and under the second illuminant the color is (r', g', b') . It is possible to map the color of white under the first illuminant to the color under the second by postmultiplication by a diagonal matrix:

$$(r', g', b') = (r, g, b) \text{diag}(r'/r, g'/g, b'/b).$$

If the same diagonal matrix transforms the (r, g, b) of all surfaces (not just the white ones) to a good approximation, then we say that we have a diagonal model of illu-

mination change. It turns out that the accuracy of the approximation is a function of the vision system's sensors.

The idea of sensor sharpening is to map the data by a linear transform T into a new space where the diagonal model holds more faithfully. Color constancy algorithms that rely on the diagonal model can then proceed more effectively. The final result is then mapped back to the original RGB space with the inverse transformation T^{-1} . Working in the transform space is like having new sensors that are linear transformations of the old ones. Generally speaking, the sensitivity functions of sensors that support the diagonal model tend to look sharper in the sense that they have narrower peaks than ones that do not. For the limiting case—that of Dirac delta functions as sensors—the diagonal model holds exactly. From these observations, we get the name “sensor sharpening.”

The main technical result in sensor sharpening is how to find the transformation T . Finlayson *et al.*¹ propose three methods for finding T : “sensor-based sharpening,” “database sharpening,” and “perfect sharpening.” In database sharpening, RGBs are generated by using a database of reflectance spectra, together with an illuminant spectrum and the sensors. This is done for two separate illuminants. Let A be the matrix of RGB for the first illuminant and B be the matrix for the second, with the RGBs placed rowwise. In the sharpening paradigm, we map from B to A with a sharpening transform, followed by a diagonal map, followed by the inverse transform. If we express each transform by postmultiplication by a matrix we get $A \approx BTDT^{-1}$. In database sharpening the matrix T (and implicitly D) is found that minimizes the rms error, $\|A - BTDT^{-1}\|_2$. The sharpening transform gives exactly the same error as the best linear transform M . In fact, T is found by diagonalizing M , where M minimizes $\|A - BM\|_2$.

Sensor-based sharpening finds the transform T by using an optimization designed to capture the intuitive notion of sharpness. This optimization requires specifying in advance a sharpening interval and the value of a Lagrange multiplier. For this work we chose database sharpening over sensor-based sharpening because of the unambiguous correspondence between the sharpening method and a color constancy error measure, as well as the difficulties in relying on user-specified parameters.

Perfect sharpening computes the sharpening transform on the assumption that surface reflectances can be fitted by using a three-dimensional linear model and, more critically, that illuminant spectra can be fitted by using a two-dimensional linear model. Although one of the strengths of this method is that it applies to a range of illuminants, it did not work well for us because our test illuminant set did not meet the key requirement of being two dimensional, partly owing to the inclusion of fluorescent lights.

Regarding database sharpening, two implementation issues should be noted. First, the above procedure can lead to complex numbers in T . Fortunately, this does not occur too often, and when it does occur, the imaginary components tend to be small. We set all imaginary components to zero.

A second implementation issue is as follows: The re-

sult of the diagonalization is ambiguous up to scaling and swapping of the columns of T . As is standard, we use columns of norm 1. Furthermore, we put the element of T of largest absolute value on the diagonal by swapping columns and ensure that it is positive by multiplying the column by -1 if necessary. Then in a similar way we attempt to make the other diagonal elements as large as possible. This procedure is used to reduce the number of negative components of sharpened data.

In this paper we view color constancy as the transformation of an input image of a scene taken under an unknown illuminant to an output image representing the same scene as though it were taken under a known, “canonical,” illuminant.⁵ *A priori*, the nature of the transformation is open, but most algorithms find a diagonal transform, and it is these algorithms that interest us here. Of course, the best linear transformation, if it could be found, would by definition always do as well as or better than the best diagonal transformation, but clearly the generalized diagonal transform TDT^{-1} gives us a chance of doing as well with a diagonal model.⁶

It could be argued that the need to improve diagonal color constancy through sharpening indicates that the diagonal model itself should be reconsidered. In short, why not find a linear transform? We argue that sharpening is important for several reasons. First, a number of promising algorithms are defined in terms of diagonal models, and effective counterparts for linear transforms are not available. Sharpening allows us to use these algorithms even when the diagonal model is not a very good approximation. Second, we have found empirically that when a camera with sharp sensors is used, the error in finding the best diagonal transform far exceeds that due to the subsequent mapping. Put differently, if we could consistently find the best diagonal transform with such a camera, the color constancy problem would essentially be solved (assuming spatially uniform illumination). This indicates that the simplifications afforded by reducing the number of parameters to be estimated can easily justify the less accurate model.

3. SENSOR SHARPENING IN PRACTICE

Database sharpening transform is defined in terms of an illuminant pair. In the context of color constancy, this presents a problem since we know only one of the illuminants (the canonical). The other is the scene illuminant, which is unknown and is precisely what needs to be found by any color constancy method. Therefore, for sharpening to be useful in color constancy, we need to find a substitute for the unknown illuminant.

One simple approach, introduced here to provide a baseline for later comparison, is to use the average of a database of common illuminant spectra^{7,8} to represent the second illuminant. The illuminant spectra must be normalized with respect to overall magnitude so that the different illuminants make comparable contributions to the average spectra. We will refer to this method as average-illuminant sharpening (“ave” sharpening).

Another approach is to find a smaller set of feasible illuminants and then average their spectra to derive a representative illuminant for sharpening. The feasible illu-

minants are found by using gamut-mapping color constancy.^{5,9} The estimates can then be refined by repeating this process: The average illuminant is used to recompute the sharpening transform, which is then used to recompute the feasible set of illuminants. Although this iterative method has some intuitive appeal, we found that the results were not much better than using the simple average, and thus we do not include this method in the experiments reported here.

To determine the upper bound on color constancy performance by using sharpening, we also computed the best database sharpening transform \mathcal{T} on the basis of correctly guessing the illuminant spectrum. Since \mathcal{T} minimizes possible rms mapping error, we will call this method “optimal.” It should be noted, however, that this method is not necessarily truly optimal once it is used in conjunction with a specific color constancy algorithm. It is possible that illuminant estimation may be best performed in some other space, with the effect that the other transform gives lower error with the chosen color constancy method. This is more likely to be the case when the initial sensors are already sharp, because a smaller portion of the error is due to the lack of sharpness. Thus with sharp sensors, using this “optimal” sharpening method can lead to worse results than not using sharpening at all. This emphasizes one of the main points of this paper: Since sharpening is not defined as minimizing the error of any particular practical computational color constancy method, existing sharpening methods cannot be assumed to improve color constancy. In this paper we provide some theoretical analysis and empirical results to help choose and develop sharpening methods on an algorithm-by-algorithm basis.

An empirical approach is useful because the interaction between sharpening and the color constancy algorithms is complex. One complication is that sharpening can lead to negative camera responses (in the sharpened space), which can be a problem for the gamut-mapping^{5,9–12} color constancy algorithms. The gamut-mapping algorithms are otherwise good candidates for sharpening since they are very effective^{12–14} and they rely completely on the diagonal model. Not only can negative sensor responses arise in the input data, but they can also occur in the algorithms’ calibration sets when they are mapped into the sharpened space. In fact, negative input data can be tolerated by the three-dimensional gamut-mapping algorithms, and the problems with negative sensor responses are confined to some of the calibration sets. In general, depending on the variation used, negative components in calibration data can present severe difficulties.

If we are to add sharpening to gamut-mapping algorithms, we must decide what to do with negative components. One strategy is to set the negative components to zero or to a small positive number if zero is also unacceptable. Unfortunately, this often leads to poor performance, or at least worse performance than without sharpening. Simply discarding input with negative components can also lead to worse performance because the additional data may help the nonsharp version. In both cases the problem is aggravated by the fact that the more extreme data (data that affects hull boundaries, including those of the calibration sets) is the most critical

for the gamut-mapping algorithms, and thus either changing or ignoring such data is problematic. We therefore take the view that when these algorithms are faced with invalid data, the best action is to revert to processing in a nonsharp space.

This leads to the following general strategy. If an algorithm is faced with invalid data due to sharpening, then sharpening is simply not used. Instead, the standard result without sharpening is computed. We report the number of times this occurred for each combination of algorithm and sharpening method tested. This strategy exposes the overall benefits (or lack thereof) of sharpening. For example, an algorithm that gives better results with sharpening half of the time, but cannot use sharpening because of negative values the other half of the time, still is improved overall by sharpening. Our testing strategy thus measures the effect of sharpening when it can be applied without modifying the specification of any algorithm. We note that a few sensor/algorithm/sharpening combinations are always incompatible. The resulting performance in these cases is exactly the same as it would be without sharpening.

Reverting back to nonsharp computation in case of difficulties is preferred to the approach used in preliminary work,² where we simply excluded from further consideration all data that gave any algorithm trouble. All algorithms were then tested on the same reduced data set. This approach is biased because the excluded data are not random, and its exclusion can favor some color constancy algorithms.¹⁵

So far we have described performing all color constancy processing in the sharpened space. A second possibility is to estimate the RGB of the illuminant in the original sensor space and then use a sharpening to correct the image on the basis of that estimate. Thus sharpening could improve the image correction results for algorithms tuned to estimate the illuminant RGB on the basis of nonsharp sensor responses. This strategy, which we will refer to as “sharp correction,” circumvents the negative-component problem. By definition, it does not affect the illumination estimation performance but only the mapping performance. Thus sharp correction cannot help with the illumination parameter estimation part of color constancy, which, at least in the case of gamut-mapping color constancy, is an expected benefit of sharpening. Further analysis of the relationship between sharpening and the color constancy algorithms tested in this work are provided in Section 7.

4. NEW APPROACH TO SHARPENING

We now introduce a new method for sharpening that deals with some of the problems inherent in database sharpening. We begin by changing the objective function of the database sharpening minimization problem. In its original form, this objective function minimizes the diagonal mapping error between the responses of a set of reflectances under two known illuminants. However, since in the context of color constancy one of these illuminants is not known, we instead minimize the average of the mapping error over a representative set of possible illuminants. Next, we include a term that encourages posi-

tive sensor response values. Finally, we include a term that encourages the transform to be norm 1. Since the resulting function is difficult to minimize directly, we look for suitable local minima by using gradient descent. We will refer to this method as multiple-illuminant-with-positivity sharpening (“mip” sharpening).

Computing the objective function requires a database of reflectance spectra that model the surfaces of the world and a database of illuminants that is a representative set of the illuminants in the world. The illuminant set should be normalized so that their magnitudes are identical. Then, for a given illuminant i , all the responses to all the reflectances can be computed by using the camera sensitivity functions. We put the responses into the rows of a matrix \mathcal{A}_i . We also form a similar matrix for the canonical illuminant \mathcal{B} . Then, given a transform \mathcal{T} , we compute the sharpened responses $\mathcal{A}_i\mathcal{T}$ and $\mathcal{B}\mathcal{T}$. We then compute the best diagonal map, in the least-squares sense, between these two matrices. This is done by averaging the rows of each matrix, dividing the two averages elementwise, and using these three ratios as the elements of a diagonal matrix. For each illuminant we thus obtain a diagonal map, \mathcal{D}_i . Finally, we can express the overall mapping error by

$$\mathbf{E} = \sum_i \|\mathcal{A}_i\mathcal{T}\mathcal{D}_i\mathcal{T}^{-1} - \mathcal{B}\|_{\mathbb{F}}, \quad (1)$$

where $\|\cdot\|_{\mathbb{F}}$ is the Frobenius matrix norm (the square root of the sum of the squares of all matrix elements).

To encourage positive sensor responses, we compute a penalty function from all the current sharpened responses, $\mathcal{A}_i\mathcal{T}$. Since some algorithms cannot cope with zeros, it helps for the responses to be at least a small offset larger than zero (we used a value of 0.1 in the experiments). The penalty function is

$$f(x) = \begin{cases} (x - \text{offset})^2 & \text{if } (x < \text{offset}) \\ 0 & \text{otherwise} \end{cases}. \quad (2)$$

The total penalty, P , is the sum of $f(x)$ over all sensor responses $\mathcal{A}_i\mathcal{T}$ over all illuminants, i . Finally, we add a term to keep \mathcal{T} near norm one. Without such a term, the positivity term could be reduced by simply scaling \mathcal{T} downward. For this we use

$$N = [\text{trace}(\mathcal{T}'\mathcal{T}) - 3]^2. \quad (3)$$

The overall objective function is therefore

$$E + \lambda_P P + \lambda_N N, \quad (4)$$

where λ_P and λ_N are weights that control the relative importance of the three parts of the objective function. Finding a suitable value of λ_N is straightforward because we are essentially using it as a convenient way to enforce a constraint, and thus we can use a very large value. We note that the magnitude of \mathcal{T} has no effect on the mapping error, because both \mathcal{T} and \mathcal{T}^{-1} occur in the expression. Once λ_N is sufficiently large, increasing it further will have very little additional impact on the resulting \mathcal{T} . On the other hand, increasing λ_P will lead to larger overall mapping error, and thus we would like to use the smallest value possible. This is discussed further below.

Thus given some \mathcal{T} , we can compute the error, and, given a change in any of the nine components of \mathcal{T} , we can compute the change in the error and thus the gradient. We then change \mathcal{T} in the direction of the gradient. If this change increases the error, we reduce the amount of change until the error does decrease. We carry on the process until \mathcal{T} does not change more than a small threshold.

This leaves the determination of the initial \mathcal{T} . Since gradient descent does not necessarily converge to the global minimum, different starting points can give different results. We tried two different strategies. First, we tried starting with the sharpening matrix obtained by using the average illuminant, as described above. Second, we tried starting with the identity matrix. Both methods gave good results, but over the three camera sensors used in preliminary work, the identity matrix starting point proved to be the best choice, and thus we use it here. The preference of the identity starting point makes sense in terms of the results presented below, which indicate that ignoring positivity can lead to much larger error than simply missing out on a possible sharpening benefit. Sharpening by using the average illuminant can lead to small or negative responses, and moving significantly away from such a transform by gradient descent will not always be possible.

Finlayson *et al.* added positivity to sensor-based sharpening by using two different criteria.¹⁶ The first criterion constrains the coefficients of the sharpening matrix to be positive. The second criterion, which is weaker and thus more useful, takes a more direct approach and simply insists that the sensors themselves be positive. The conditions on positivity imposed by Eq. (2) are even weaker in two ways. First, Eq. (2) only penalizes negative sensor responses to the set of input spectra; the sensor functions themselves can have small negative components. Second, our penalty function allows for the possibility of an occasional small negative sensors response. In fact, this is how we set the weight for the positivity term. We use a value that yields “positive-enough” responses for the algorithms under consideration. Thus conceptually, the weight is a function of the algorithms(s) to be improved. Given our strategy of reverting back to the standard form of the algorithm if necessary, we can tolerate the occasional negative sensor response. However, since we wished to use the same sharpening transform for all the algorithms under consideration, we found it convenient to choose a value for the weight that yielded very few negative sensor responses.

Finally, we note that the method developed in this section also sidesteps having to deal with the occurrence of complex elements in the transform matrix. As mentioned above, such elements sometimes occur when standard database sharpening is used.

5. COLOR CONSTANCY METHODS

We will now discuss briefly the color constancy algorithms that we test here in the context of sensor sharpening. Further details on these algorithms are available.^{12,13} The first two algorithms are based on the assumption that the average RGB of a scene under a given illuminant

is the same as that of “gray.” The exact definition of “gray” distinguishes the algorithms. One possibility is simply true gray, specifically, a 50% uniform reflectance. This leads to the algorithm labeled GW in the results. A second choice is to use the average of the reflectance database. This method has an unfair advantage when tested on synthetic data generated from the same reflectance database. Its results are guaranteed to be excellent if a large number of surfaces are used. In the case of real images, however, the actual average surface reflectance is not known, and thus this method is expected to fare relatively less well.^{12,17} We denote this algorithm DB-GW (database-gray world) in the results. The implied diagonal map is the ratio of the known RGB of gray under the canonical illuminant to the RGB of gray under the test illuminant, estimated by the average of the scene RGB.

The second algorithm is loosely based on the Retinex model of human vision.^{18–20} The result is computed by using the maximum in each channel as an estimate of the color of white under the test illuminant. Similar to the gray-world algorithm, the implied diagonal map is the ratio of the known color of white under the canonical illuminant to the estimate. This algorithm is labeled SCALE-BY-MAX in the results.

The gamut-mapping approach introduced by Forsyth⁵ estimates the diagonal map directly from a set of possible maps. The possible maps can be constrained by considering the scene RGB (Ref. 5) and insisting on plausible illuminants.⁹ The two gamut-mapping algorithms described below use both sources of constraints.

Given a set of possible diagonal maps, a solution needs to be chosen from this set. One method for doing so is maximizing the volume of the convex hull of the mapped RGBs (Ref. 5) that is used in the algorithm ECRULE-MV. A second method is to use the average of the possible maps.^{10–12} When Finalyson’s illumination constraint is used, the set of possible maps is nonconvex, and thus in the algorithm labeled ECRULE-ICA in the results, the average is estimated numerically.

We do not report any results for the chromaticity versions of gamut mapping, for three reasons. First, since these methods do not consider overall illuminant brightness, their results cannot be used in conjunction with the RGB mapping error. Second, we have found the three-dimensional methods more effective overall.^{10,12–14,17} Third, we have found that the chromaticity gamut-mapping methods are by far the most sensitive to negative sensor responses.

6. ERROR MEASURES

We consider color constancy processing as transforming an image of a scene taken under an unknown illuminant to an image of the same scene as if it were taken under a standard, canonical illuminant.⁵ In the context of image correction, the canonical illuminant would normally be one for which the camera was balanced. Achieving this transformation implicitly characterizes the illuminant. Many color constancy algorithms first explicitly characterize the illuminant, often by the RGB of pure white, and then correct the image on that basis. Regardless, it is

natural to consider the correction error as the measure of color constancy performance.

In this work we compute the rms RGB difference between the corrected image and the target image, where the target is defined as the same scene imaged under the canonical illuminant. This measure is similar to the one that sensor sharpening attempts to minimize. There is one ambiguity worth further discussion. Either we can measure the transformation error for the specific scene as described above, or we might consider the transformation error for the set of known surfaces as reported in preliminary work.² The second error is more precisely what is minimized when the sharpening transform is computed but is less indicative of how well images are being corrected and is further removed from our main goal. Therefore in this work we report the scene-oriented mapping error.

To compute this error, we start with a diagonal transformation that takes sharpened responses under the unknown illuminant to sharpened responses under the canonical illuminant. If the algorithm is being applied to sharpened input data, then this transformation is immediately available as output. If sharpening is being used only to correct the image on the basis of an estimate of the RGB of the illuminant, then that estimate is transformed to an estimate of the illuminant in the sharpened space. The ratio of the RGB of the canonical illuminant in sharpened space to this estimate is then used as the diagonal transform in sharpened space. Regardless of which approach is used to find the transform, it is then applied to the sharpened input data to obtain an estimate of the scene under the canonical illuminant in the sharpened space. We then apply the inverse sharpening transform to this mapped set to obtain the estimate of the responses under the canonical illuminant without sharpening. This estimate is compared with the target by using the rms of the RGB difference between them.

We also report a measure of the error in the estimated illuminant chromaticity. This estimate is implicit in the gamut-mapping algorithms and is essentially explicit in the other algorithms studied here that estimate the RGB of the illuminant. It is reasonable to consider chromaticity because we are often less interested in the overall brightness of the illuminant. For example, if we are correcting a properly exposed image taken with incorrect camera balance, then the brightness of the illuminant is implicit, and we need only to correct for the illuminant chromaticity. In fact, a number of algorithms work only with chromaticity.^{9,21,22} We do not consider any of these algorithms here, but a good counterpoint to the rms RGB mapping error measure is a chromaticity error measure. The one that we use here is the angle between the illuminant RGB and the estimate thereof, considered as vectors in RGB space.

7. COLOR CONSTANCY WITH SHARPENING

Each color constancy algorithm is affected differently by sharpened data. We consider the gray-world algorithms first. The illuminant estimate provided by these two algorithms is given by

$$\mathbf{w} = \left(\frac{1}{n} \sum_i^n \mathbf{r}_i \right) \text{diag}(\mathbf{w}_c ./ \mathbf{g}_c), \quad (5)$$

where the \mathbf{r}_i are the observed sensor responses, \mathbf{w}_c is the RGB of white under the canonical illuminant, \mathbf{g}_c is the RGB of gray under the canonical illuminant, $./$ is used to denote elementwise division, and $\text{diag}(\cdot)$ transforms a vector to a diagonal matrix. When sharpening is used, Eq. (5) becomes

$$\mathbf{w} = \left\{ \left(\frac{1}{n} \sum_i^n \mathbf{r}_i T \right) \text{diag}[(\mathbf{w}_c T) ./ (\mathbf{g}_c T)] \right\} T^{-1}. \quad (6)$$

As discussed above, we have two different definitions of gray, giving two different values of \mathbf{g}_c . In the case of the GW algorithm, the reflectance spectrum of gray is assumed to be half that of perfect white. Under this condition, \mathbf{w}_c is twice \mathbf{g}_c and $\text{diag}(\mathbf{w}_c ./ \mathbf{g}_c)$ can be replaced by 2. Similarly, $\mathbf{w}_c T = (2\mathbf{g}_c)T = 2(\mathbf{g}_c T)$, and $\text{diag}[(\mathbf{w}_c T) ./ (\mathbf{g}_c T)]$ can also be replaced by 2. It is then straightforward to see that Eqs. (5) and (6) are the same. Hence we expect sharpening to have absolutely no effect on the estimate of the illuminant when the GW algorithm is used.

It should also be clear that when \mathbf{g}_c is not a simple scalar times \mathbf{w}_c , as is the case with DB-GW, that there will be some effect. Intuitively, the effect should be small, as we expect \mathbf{g}_c to be not too far from a scalar times \mathbf{w}_c . However, given the nature of the expression, it is possible for the effect to be large. For example, the illuminant estimate becomes unstable for sharpening transforms T that map a component of \mathbf{g}_c close to zero.

The effects of sharpening in terms of the mapping error are more interesting. In the nonsharp case, the gray-world algorithms estimate the diagonal mapping \mathcal{D} , which maps RGB taken under the unknown illuminant to the corresponding RGB taken under the canonical illuminant by

$$\mathcal{D} = \text{diag} \left[\mathbf{g}_c ./ \left(\frac{1}{n} \sum_i^n \mathbf{r}_i \right) \right]. \quad (7)$$

In the sharp case, the estimate is for the analogous mapping, \mathcal{D}^\sharp , applied in the sharpened space. \mathcal{D}^\sharp is computed similarly to \mathcal{D} , but with the sharpened quantities $\mathbf{g}_c T$ and $\mathbf{r}_i T$ used in place of \mathbf{g}_c and \mathbf{r}_i :

$$\mathcal{D}^\sharp = \text{diag} \left[(\mathbf{g}_c T) ./ \left(\frac{1}{n} \sum_i^n \mathbf{r}_i T \right) \right]. \quad (8)$$

Now, if the sharpening transform produces negative components, it is entirely possible that the sum in Eq. (8) might be close to zero. In this case, the error in \mathcal{D}^\sharp can become arbitrarily large. Thus even in the case of the GW algorithm where the illuminant estimate is identical to that obtained without sharpening, the RGB mapping error can be much larger than that obtained without sharpening. This may seem contradictory, but it even applies when sharpening is used only for correcting the

image. Consider computing the mapping using the estimate of \mathbf{w} obtained in Eq. (5). In the nonsharp case, we have

$$\mathcal{D} = \text{diag}(\mathbf{w}_c ./ \mathbf{w}) \quad (9)$$

In the sharp case, we have:

$$\mathcal{D}^\sharp = \text{diag}[(\mathbf{w}_c T) ./ (\mathbf{w} T)]. \quad (10)$$

If the illuminant estimate has components close to zero in the sharp space, then the error in \mathcal{D}^\sharp , and therefore the mapped RGB, will likely be large.

This analysis indicates that this problem should be much less severe with use of the multiple-illuminant-with-positivity sharpening method introduced in Section 4 since it is designed to yield at most a few small negative components. However, since there is no indication that the result should be better—only less worse—we conclude that sharpening does not make sense with gray-world algorithms.

The SCALE-BY-MAX algorithm is not likely to have the same problem with near-zero values since its estimate is based on the maximum from each channel, which are likely to be significantly greater than zero for any reasonable sharpening method. This argument applies with less force in the case of sharp correction. Here the space used for the maximum is different from the mapping space, and the problem illustrated by Eq. (10) still exists.

If we apply the SCALE-BY-MAX algorithm entirely in a sharpened space, we need to consider whether the maxima will still provide a reliable estimate of the illuminant. Without sharpening, the maximum of each channel is justified, since as the number of surfaces in the scene increases, eventually there will either be a white one or at least one each with maximal R, G, and B values. This is not necessarily the case when sharpening is used. For example, suppose that

$$T = \begin{bmatrix} 1 & 0 & -\frac{1}{3} \\ 0 & 1 & 0 \\ 0 & 0 & 1 \end{bmatrix}. \quad (11)$$

Further, suppose that the RGB corresponding to white is (1, 1, 1). Then the RGB corresponding to white in the sharpened space is (1, 1, 1)* $T = (1, 1, 2/3)$. In general, the RGB in the sharpened space are given by [R, G, B - (1/3)R]. It is quite possible that there are valid sensor responses that are larger in the third coordinate than the one corresponding to white. For example, for RGB = (1/2, 1, 1), then the maximum of the third channel after sharpening could reach 5/6, which is greater than the value of 2/3 for white.

For the above example to work, the sharpened sensors must have significant negative responses. Thus sharpening with positivity should largely solve this problem. Furthermore, there may be some benefit to sharpening with SCALE-BY-MAX during illumination estimation since it estimates each channel of the illuminant RGB separately, and these channels become less correlated with sharpening. Consider the red channel. If the sensor has broad

support, then for the response to be close to the maximum possible under the given illuminant, it needs a broad input spectrum. Thus it needs a correspondingly whiter surface in the image for the same performance. Since the existence of a near-white surface does not degrade performance in the sharp case, the overall effect should be improved performance in the sharp case, all else being equal. Since sharpening can also improve results by reducing the mapping error, we investigate these effects by considering both the full-sharpening results and the sharp-correction results, using both the mapping-error measure and illuminant-chromaticity error measure.

With the gamut-mapping methods we estimate the diagonal maps first and then, if necessary, use the maps to compute illuminant estimates. We remind the reader that the estimation of the maps consists of two parts. First, we compute the set of possible maps using applicable constraints. Second, we choose a map from this constraint set. The first part is highly dependent on the diagonal model, and thus gamut-mapping algorithms are likely candidates for improvement through sharpening.

It is less clear what the effect of sharpening will be on the various methods for choosing the solution from the constraint set. We first consider using the average of the possible maps to choose the solution. Averaging as used in ECRULE-ICA was originally proposed to reduce the error using the RGB mapping measure,¹⁰ and further work supports the notion that this method is the best choice for this measure.¹²⁻¹⁴ Furthermore, we expect that the arguments that justify this choice in the nonsharp case are not overly eroded in the sharpened case. Specifically, even if sharpening is used, we still expect that the solutions will be distributed throughout the constraint set without overly large bias and that the computation of the corresponding mapping error will not add too much additional bias.

On the other hand, the maximum volume heuristic used in ECRULE-MV may be sensitive to problems similar to those outlined above for the SCALE-BY-MAX method. This heuristic is not fully understood in the standard case, and things become even more complex when sharpening is added. Nonetheless, we expect that encouraging positivity will likely alleviate most problems.

In summary, the effect of sharpening on the various algorithms is hard to predict. This is especially the case for the gamut-mapping algorithms, which are the conjunction of two different parts, each of which interacts with sharpening differently. Thus we believe that the most effective way to further understand the role of sharpening is to look at empirical results, which we do next.

8. EXPERIMENTS WITH SYNTHETIC DATA

We investigated sensor sharpening in the case of a Sony DXC-930 CCD video camera and a Kodak DCS-460 camera (both calibrated as described by Barnard⁷), a Kodak DCS-200 digital camera,⁴ and a Kodak DCS-420 digital camera.⁴ The sensors for each camera together with the sharpened versions for two sharpening methods are shown in Fig. 1. We provide results for each camera in

turn, first numerically in Tables 1-4 and then graphically with respect to the RGB mapping error measure in Fig. 2.

The algorithms were first calibrated on a set of 87 illuminants constructed to cover that portion of chromaticity space [$r = R/(R + G + B)$, $g = G/(R + G + B)$] which includes most common illuminants in a roughly uniform manner. For testing, we used a similar set of illuminants, but we reduced the spacing in (r, g) space by a factor of 2 in order to obtain a larger illuminant set (287 illuminants) that had similar scope in (r, g) space. Additional description of these illuminant sets is available,^{8,12} and the data itself are on line.⁷ For surface reflectances we used a set of 1995 spectra compiled from several sources (also available on line⁷). These surfaces included the 24 Macbeth ColorChecker® patches, 1269 Munsell chips, 120 Dupont paint chips,²³ 170 natural objects,²³ the 350 surfaces in the Krinov data set,²⁴ and 57 additional surfaces measured by ourselves.

The results are the rms average results for 3000 synthetically generated scenes, each using a randomly selected illuminant and eight randomly chosen surfaces. For each algorithm we computed the results both with and without sharpening. There were three sharpening methods as described above: (i) the optimal database method, which uses the actual illuminant (not normally available); (ii) database sharpening using the average training set of the illuminants; and (iii) multiple illuminant with positivity sharpening. We also used these sharpening transforms to correct the images on the basis of an estimate of the illuminant RGB computed by using the unmodified (nonsharp) input (sharp correction).

For the optimal sharpening we recorded the number of eigenvectors having imaginary components above a small threshold, and we averaged their magnitudes for the 3000 generated scenes. Since there are only 287 test illuminants, and the optimal sharpening is only a function of the illuminant, there are many repeats. Each optimal sharpening yields 3 eigenvectors of unit length. For the DXC-930 camera, 1086 eigenvectors had nonnegligible imaginary components with average magnitude 0.021. For the DCS-460, there were 1248 with average magnitude 0.039; for the DCS-200, there were 1692 with average magnitude 0.046; and for the DCS-420, there were 1344 with average magnitude 0.028. These relatively small numbers support the conclusion that imaginary components do not create a significant problem.

We include the results for several comparison "algorithms." The first is the minimum-error result obtained by finding the best linear fit between the RGB of the reflectance database under the canonical and the test illuminants (BEST-LINEAR). This result is invariant to matrix multiplication and thus is the same for all sharpening methods. The second comparison method is the best-diagonal map computed over the entire dataset (BEST-DIAGONAL). The third comparison method is the RGB of white under the test illuminant (ACTUAL). The error with this method is zero, with use of illuminant-based error measures, and would also be zero in the case of mapping error if the diagonal model held perfectly. This comparison method provides an additional insight into whether the diagonal model is much improved, without the complexities introduced by the interactions with the algo-

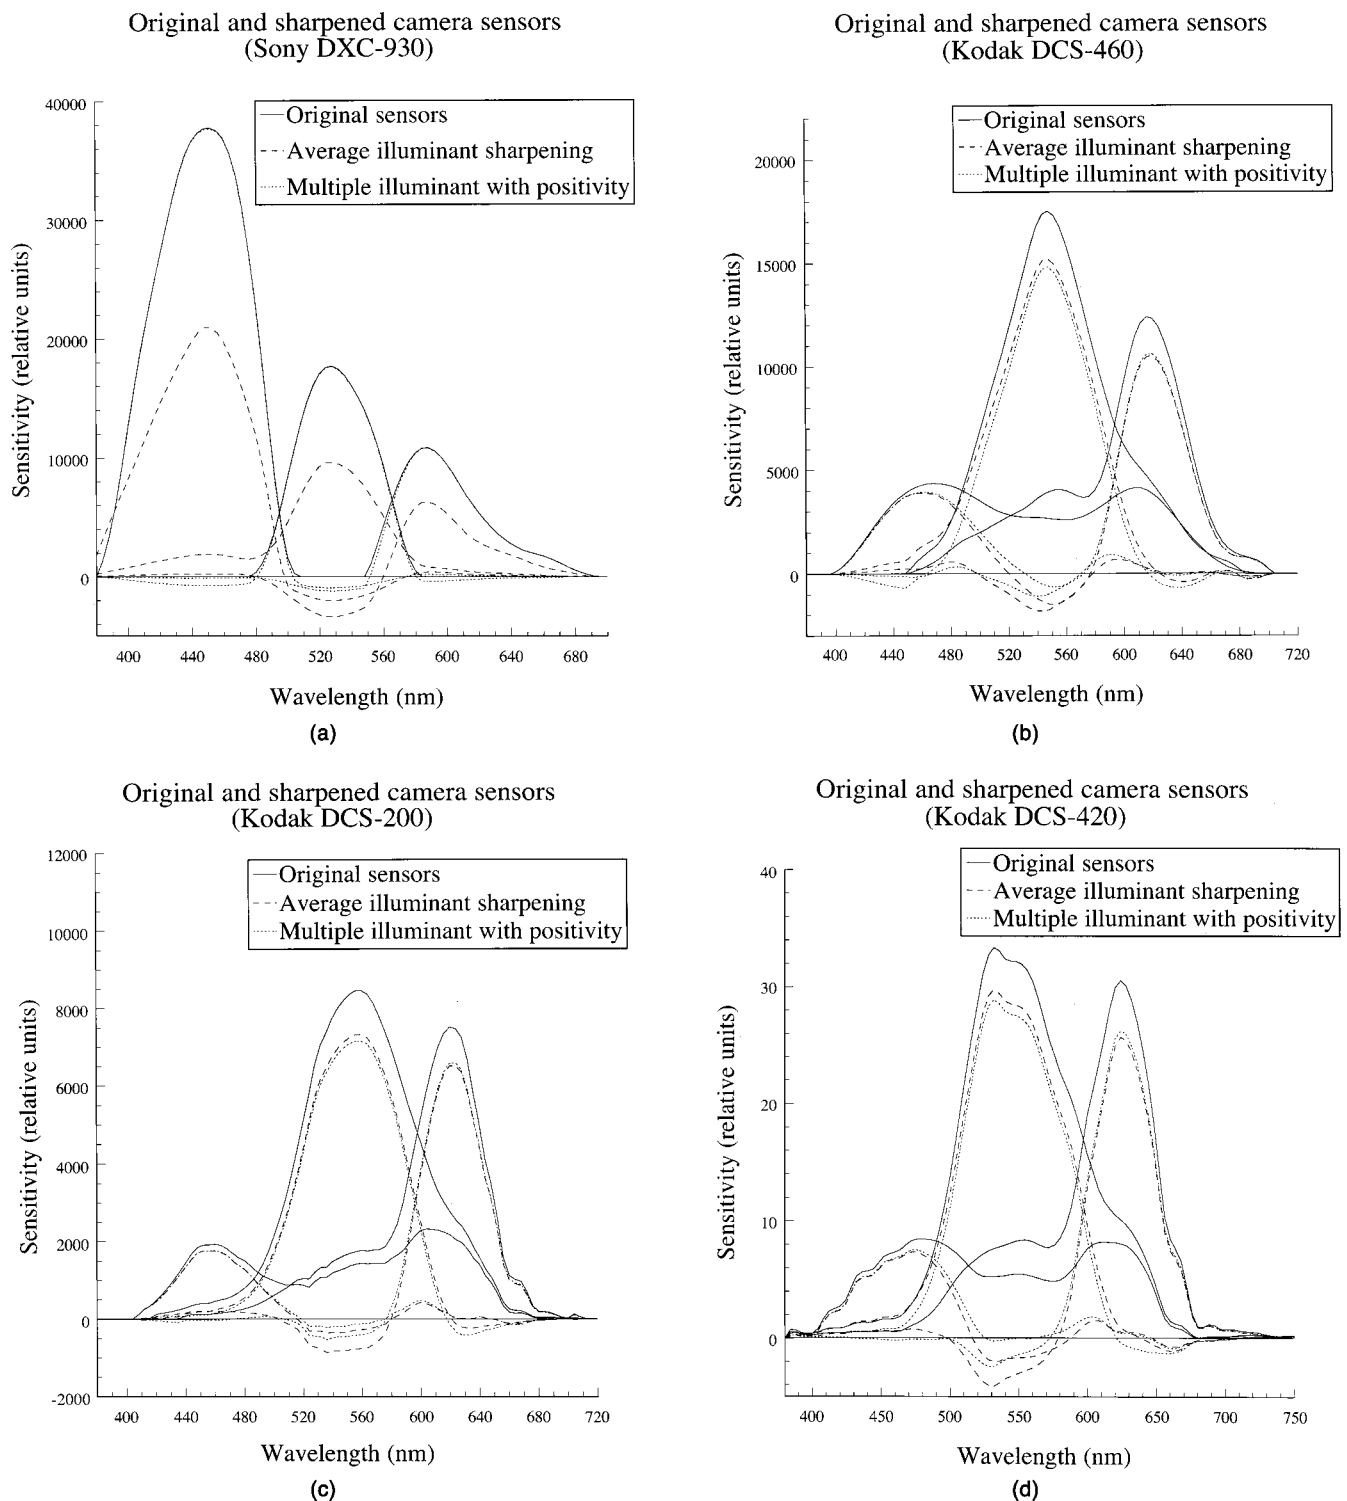


Fig. 1. Original camera sensors, sensors corresponding to sharpening based on the average of the database of illuminants, and sensors corresponding to the multiple illuminant with positivity sharpening method introduced in this paper for (a) the Sony DXC-930, (b) the Kodak DCS-460, (c) the Kodak DCS-200, and (d) the Kodak DCS-420. Note that for the DXC-930 [(a)], the sensors sharpened with the multiple-illuminant-with-positivity method are very close to the original ones, and the two sets of sensor curves are blended for much of the wavelength range. Since these sensors are already “sharp,” this is encouraging. The three Kodak cameras have blue sensors that have secondary peaks in the red region of the spectrum as well as significant responses in the green region. The red sensors also have significant response in the green region. Sharpening essentially removes these characteristics and also yields narrower green sensors. Thus for these three cameras, the term “sharpening” is very appropriate.

rithms. Finally, we provide the result of doing nothing, in other words taking the input as the output (NOTHING). This result is valid only for the illumination-based error

measure. We do not compute its RGB mapping error because the data sets were generated with use of illuminants normalized by magnitude, and so NOTHING gives an

excellent estimate of illuminant magnitude simply as a consequence of our testing paradigm. None of the other algorithms benefit since they all attempt to compute the illuminant magnitude.

9. RESULTS WITH SYNTHETIC DATA

As expected, the results are quite camera dependent. We define the degree of sharpness of the camera sensors as the ratio of the BEST-LINEAR result to the BEST-DIAGONAL result. By this measure, the Sony DXC-930 camera is quite sharp (0.70), and the Kodak DCS-460, DCS-200,

and the DCS-420 are less sharp (0.33, 0.33, and 0.34 respectively). The two comparison algorithms BEST-DIAGONAL and ACTUAL give some idea as to the scope of improvement that is possible as a result of reducing mapping error. Especially in the case of ACTUAL, sharpening plays no role in the illumination estimation part of color constancy. With this “algorithm” and the DCS-460 and DCS-200 sensors, the average-illuminant (“ave”) sharpening method produced sufficient components close to zero to destabilize the mapping process and give a large mapping error. This is also true to a lesser extent with the DXC-930 and the DCS-400. Optimal (“opt”) sharpening has the same problem in the case of the DCS-460 but oth-

Table 1. Sharpening Results for the Sony DXC-930 Video Camera on Synthetic Data^a

Algorithm	Sharpening Method	rms RGB Difference between Mapped Image and Target Image	rms RGB Difference between Mapped Image and Target Image (Sharp Correction)	rms Angle between Illumination RGB and Estimate Thereof	Count of Times Standard Results Replaced Full-Sharp Result Due to Problems (Usually Due to Negative Components)
BEST-LINEAR	none	2.93	2.93	*	0
BEST-DIAGONAL	none	4.25	4.25	0.29	0
BEST-DIAGONAL	ave	9.04	27.66	0.96	0
BEST-DIAGONAL	opt	2.98	21.58	0.31	0
BEST-DIAGONAL	mip	4.09	4.12	0.32	0
ACTUAL	none	4.38	4.38	0.00	0
ACTUAL	ave	21.91	21.91	=	0
ACTUAL	opt	4.92	4.92	=	0
ACTUAL	mip	4.20	4.20	=	0
NOTHING	none	*	*	16.30	0
GW	none	136.57	136.57	8.12	0
GW	ave	379.95	379.95	=	0
GW	opt	1951.78	1951.78	=	0
GW	mip	136.70	136.70	=	0
DB-GW	none	33.66	33.66	6.54	0
DB-GW	ave	203.32	167.33	6.54	0
DB-GW	opt	576.75	349.75	11.39	0
DB-GW	mip	33.70	33.75	6.53	0
SCALE-BY-MAX	none	105.57	105.57	9.43	0
SCALE-BY-MAX	ave	105.09	1014.44	9.56	0
SCALE-BY-MAX	opt	829.38	393.21	10.56	0
SCALE-BY-MAX	mip	103.98	106.12	9.35	0
ECRULE-MV	none	50.28	50.28	5.99	0
ECRULE-MV	ave	50.74	63.18	6.22	1
ECRULE-MV	opt	49.87	261.00	5.98	998
ECRULE-MV	mip	49.98	50.59	6.00	0
ECRULE-ICA	none	37.15	37.15	7.01	0
ECRULE-ICA	ave	39.01	37.76	6.93	0
ECRULE-ICA	opt	36.62	283.50	6.79	997
ECRULE-ICA	mip	37.47	37.05	7.10	8

^a Assuming that the algorithms provide estimates that are normally distributed around the target values, the uncertainty in these numbers is roughly 1%. An asterisk is used for values that are not relevant or appropriate. An equal sign is used for results which must be the same as the nonsharp result. The results from the optimal database transform (assuming the illuminant is known) are indicated by “opt,” the results from the database transform with the average illuminant by “ave,” and the results from the multiple-illuminant-with-positivity sharpening method by “mip.” Results from using the two error measures are provided. In the case of the rms RGB mapping error, we provide results both for using sharpening throughout the processing (full sharp), and for using sharpening only to correct the image based on an illuminant estimate found in nonsharp space (sharp correction). Sharp correction has no effect on the estimation of the illumination RGB, and the angular error is exactly the same as for no sharpening and thus is not listed. The sharp-correction mapping errors are the same as their full-sharp counterparts in the case of ACTUAL and GW. Here we simply repeat the numbers.

Table 2. Sharpening Results for the Kodak DCS-460 Digital Camera on Synthetic Data^a

Algorithm	Sharpening Method	rms RGB Difference between Mapped Image and Target Image	rms RGB Difference between Mapped Image (Sharp Correction)	rms Angle between Illumination RGB and Estimate Thereof	Count of Times Standard Results Replaced Full-Sharp Result Due to Problems (Usually Due to Negative Components)
BEST-LINEAR	none	1.78	1.78	*	0
BEST-DIAGONAL	none	5.34	5.34	0.45	0
BEST-DIAGONAL	ave	6.56	13.97	0.98	0
BEST-DIAGONAL	opt	1.80	3.62	0.08	0
BEST-DIAGONAL	mip	2.53	3.09	0.14	0
ACTUAL	none	5.51	5.51	0.00	0
ACTUAL	ave	58.68	58.68	=	0
ACTUAL	opt	1.85	1.85	=	0
ACTUAL	mip	2.58	2.58	=	0
NOTHING	none	*	*	6.34	0
GW	none	108.77	108.77	2.81	0
GW	ave	113.52	113.52	=	0
GW	opt	117.96	117.96	=	0
GW	mip	108.96	108.96	=	0
DB-GW	none	29.78	29.78	2.45	0
DB-GW	ave	34.55	82.78	2.33	0
DB-GW	opt	35.77	45.42	2.45	0
DB-GW	mip	29.68	29.88	2.43	0
SCALE-BY-MAX	none	86.06	86.06	3.93	0
SCALE-BY-MAX	ave	83.98	744.56	4.79	0
SCALE-BY-MAX	opt	4079.01	465.65	11.16	0
SCALE-BY-MAX	mip	79.49	103.72	3.45	0
ECRULE-MV	none	42.95	42.95	3.41	0
ECRULE-MV	ave	=	44.02	=	3000 (all)
ECRULE-MV	opt	42.19	74.20	2.47	973
ECRULE-MV	mip	42.79	43.33	2.38	0
ECRULE-ICA	none	49.38	49.38	2.85	0
ECRULE-ICA	ave	=	49.37	=	3000 (all)
ECRULE-ICA	opt	32.91	51.30	2.31	975
ECRULE-ICA	mip	33.63	49.40	2.27	12

^aSee notes for Table 1.

erwise gives the best mapping results. Multiple-illuminant-with-positivity (“mip”) sharpening always improved the mapping when used with ACTUAL and BEST-DIAGONAL.

We now move onto the real color constancy algorithms. In the case of the already sharp camera, we found that there was very little to be gained by sharpening. Interestingly, the ave and opt sharpening methods often made the results substantially worse. In the case of opt sharpening applied to GW, DB-GW, and SCALE-BY-MAX, the results were extremely poor owing to the inversion of numbers near zero. By contrast mip sharpening did not substantially increase (or decrease) the error of any of the algorithms.

For the Kodak cameras, the opt sharpening method was again often unstable with GW, DB-GW, and SCALE-BY-MAX. Even the generally better-behaved ave sharpening method often made matters worse; mip sharpening fared

better. Its error was comparable (within 1%) to no sharpening for GW and DB-GW, and it was modestly better than no sharpening for SCALE-BY-MAX. This supports the notion in Section 7 that sharpening can improve SCALE-BY-MAX provided that the negative components are not a problem. Furthermore, since the illumination-chromaticity error was generally decreased at least as much (relatively) as the mapping error, our results indicate that sharpening helps SCALE-BY-MAX to estimate illumination parameters. Moving to sharp correction, we found that the results (only the mapping results are applicable) are always substantially worse. The negative impact is consistently greater with ave and opt compared with mip, but even when positivity was encouraged, sharp correction gives poor results. This is not surprising in the light of the discussion in Section 7. When the correction is done in the same space used for the maximum, then, provided that positivity is encouraged, the de-

nominator in Eq. (10) should not be close to zero. On the other hand, if the correction is done in a space other than that used for the maximum, then there is no such guarantee.

With the ECRULE-ICA algorithm, we found that sharpening has a large positive effect as gauged by both error measures in the case of opt and mip, which had comparable results. With the Kodak DCS-460 and DCS-200 cameras, ave sharpening could not be used with this algorithm owing to the problem with negative components, and thus it gives the same result as without sharpening. In the case of the Kodak DCS-420 camera, the ave sharpening method was viable, but it gave a poor result. In summary, for the ECRULE-ICA algorithm, mip sharpening

offers the same significant benefits as opt sharpening, but at the same time it is realizable in practice.

The results for the ECRULE-MV algorithm are a little more complex. The results for mip again parallel those for opt; however, with this algorithm, sharpening has only a very small positive effect with the mapping measure but a very significant positive effect with the illuminant-based chromaticity measure.

10. EXPERIMENTS WITH IMAGE DATA

Testing the effect of sharpening algorithms on real data is difficult because ideally one would like to report the RGB

Table 3. Sharpening Results for the Kodak DCS-200 Digital Camera on Synthetic Data^a

Algorithm	Sharpening Method	rms RGB Difference between Mapped Image and Target Image	rms RGB Difference between Mapped Image and Target Image (Sharp Correction)	rms Angle between Illumination RGB and Estimate Thereof	Count of Times Standard Results Replaced Full-Sharp Result Due to Problems (Usually Due to Negative Components)
BEST-LINEAR	none	1.61	1.61	*	0
BEST-DIAGONAL	none	4.86	4.86	0.41	0
BEST-DIAGONAL	ave	7.05	49.68	1.43	0
BEST-DIAGONAL	opt	1.65	2.68	0.07	0
BEST-DIAGONAL	mip	2.29	2.77	0.14	0
ACTUAL	none	5.06	5.06	0.00	0
ACTUAL	ave	15.09	15.09	=	0
ACTUAL	opt	1.73	1.73	=	0
ACTUAL	mip	2.33	2.33	=	0
NOTHING	none	*	*	6.57	0
GW	none	106.71	106.71	2.70	0
GW	ave	307.67	307.67	=	0
GW	opt	125.27	125.27	=	0
GW	mip	106.94	106.94	=	0
DB-GW	none	29.86	29.86	2.34	0
DB-GW	ave	169.96	99.96	2.22	0
DB-GW	opt	37.50	125.23	2.39	0
DB-GW	mip	29.79	29.98	2.32	0
SCALE-BY-MAX	none	83.26	83.26	3.77	0
SCALE-BY-MAX	ave	116.24	199.27	4.76	0
SCALE-BY-MAX	opt	7766.85	152.31	7.67	0
SCALE-BY-MAX	mip	77.84	88.11	3.35	0
ECRULE-MV	none	44.90	44.90	3.53	0
ECRULE-MV	ave	=	45.97	=	3000 (all)
ECRULE-MV	opt	42.88	56.73	2.35	799
ECRULE-MV	mip	43.50	45.24	2.32	0
ECRULE-ICA	none	50.82	50.82	2.86	0
ECRULE-ICA	ave	=	50.85	=	3000 (all)
ECRULE-ICA	opt	31.54	61.36	2.20	801
ECRULE-ICA	mip	32.61	50.85	2.23	1

^aSee notes for Table 1.

Table 4. Sharpening Results for the Kodak DCS-420 Digital Camera on Synthetic Data^a

Algorithm	Sharpening Method	rms RGB Difference between Mapped Image and Target Image	rms RGB Difference between Mapped Image and Target Image (Sharp Correction)	rms Angle between Illumination RGB and Estimate Thereof	Count of Times Standard Results Replaced Full-Sharp Result Due to Problems (Usually Due to Negative Components)
BEST-LINEAR	none	1.99	1.99	*	0
BEST-DIAGONAL	none	5.93	5.93	0.43	0
BEST-DIAGONAL	ave	6.44	17.82	0.81	0
BEST-DIAGONAL	opt	2.03	8.32	0.08	0
BEST-DIAGONAL	mip	2.82	3.50	0.16	0
ACTUAL	none	6.16	6.16	0.00	0
ACTUAL	ave	9.50	9.50	=	0
ACTUAL	opt	26.30	26.30	=	0
ACTUAL	mip	2.87	2.87	=	0
NOTHING	none	*	*	7.14	0
GW	none	111.76	111.76	2.82	0
GW	ave	114.95	114.95	=	0
GW	opt	129.41	129.41	=	0
GW	mip	112.06	112.06	=	0
DB-GW	none	31.07	31.07	2.46	0
DB-GW	ave	33.94	37.29	2.35	0
DB-GW	opt	79.71	63.07	5.09	0
DB-GW	mip	30.94	31.17	2.43	0
SCALE-BY-MAX	none	88.51	88.51	3.94	0
SCALE-BY-MAX	ave	86.30	134.32	4.72	0
SCALE-BY-MAX	opt	6663.70	130.11	16.75	0
SCALE-BY-MAX	mip	82.29	92.25	3.54	0
ECRULE-MV	none	45.61	45.61	3.79	0
ECRULE-MV	ave	42.33	46.42	4.07	0
ECRULE-MV	opt	44.52	248.62	2.59	850
ECRULE-MV	mip	45.09	46.15	2.53	0
ECRULE-ICA	none	51.24	51.24	3.06	0
ECRULE-ICA	ave	54.62	51.22	3.88	3
ECRULE-ICA	opt	33.18	53.35	2.32	850
ECRULE-ICA	mip	34.20	51.23	2.41	3

^a See notes for Table 1.

mapping error. However, this would require images of the same scene taken under different illuminants with a minimal change in the illumination geometry and no change in the scene geometry. The only easy way to acquire such images is by using filters to change the illumination, but this conflicts with the goal of testing over the range of common lighting.¹²⁻¹⁴ Furthermore, we would prefer to test the effect of sharpening on a camera with relatively nonsharp sensors. These two criteria rule out reliance on the existing image set taken with the Sony DCX-930.¹²⁻¹⁴

Thus to test sharpening in practice, we collected a new color constancy data set using the Kodak DCS-460. The illuminants and the experimental protocol were similar to those used previously,^{7,8,12-14} except that the three fluorescent lights were excluded, leaving eight lights that could be changed without affecting the illumination ge-


ometry. These eight illuminants consisted of four incandescent lights each with and without a blue filter, chosen to cover the chromaticity range of daylight. The four light sources were affixed to a trolley that was rolled back and forth behind a single aperture. The light passed through the aperture and then alternately through a filter or no filter. We took images of 18 scenes for a total of 144 images. The results were computed by using the same canonical illuminant as in previous work,^{7,8,12-14} namely, the Sylvania 50MR16Q, whose spectrum is similar to that of a household incandescent light.

The ave and the mip sharpening transforms were the same as the ones used for the synthetic experiments. In analogy to the synthetic experiments, the opt sharpening transform was recomputed for each scene on the basis of the measured spectrum of the illumination. It is important to note that the relationship of the sharpening trans-

forms to the scene reflectance and test illuminant statistics is different from that in the synthetic experiments. In the synthetic experiments, the illuminant and reflectance sets used for finding the ave and the mip sharpening transforms were exactly the same as the ones used for


generating test scenes. On the other hand, the illuminants used for the image data were the eight described above, and the statistical occurrence of the scene surfaces is likely quite different from those in the reflectance database. We assume that if we made the conditions of the

RMS RGB mapping error as a function of algorithms and sharpening methods (Sony DXC-930)



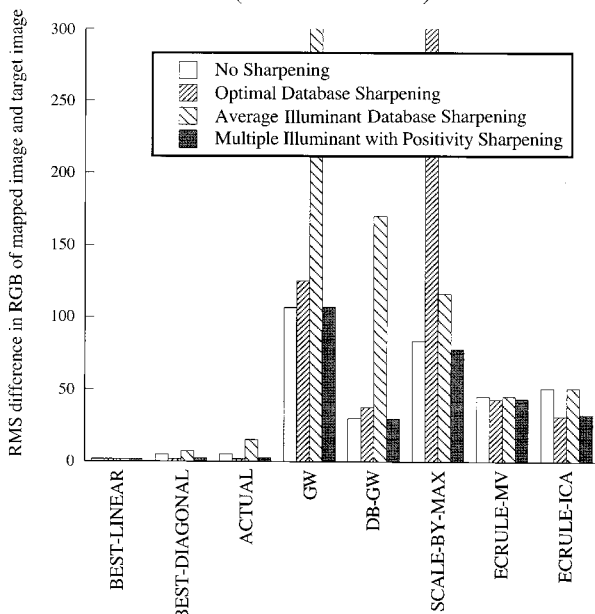
(a)

RMS RGB mapping error as a function of algorithms and sharpening methods (Kodak DCS-460)




(b)

RMS RGB mapping error as a function of algorithms and sharpening methods (Kodak DCS-200)



(c)

RMS RGB mapping error as a function of algorithms and sharpening methods (Kodak DCS-420)



(d)

Fig. 2. Rms RGB mapping error between the corrected image and the target image by algorithm-sharpening method combination for (a) the Sony DXC-930, (b) the Kodak DCS-460, (c) the Kodak DCS-200, and (d) the Kodak DCS-420. Note that the scale would have to be significantly larger to illustrate fully the extent of all the bars [7× in (a), 14× in (b), 30× in (c), and 20× in (d)]. These extreme errors are due to instabilities in the computation that occur with small or negative components. Their exact values depend on the test set. Therefore the results should be taken in a qualitative sense. RMS corresponds to rms in the text.

Table 5. Sharpening Results for the Kodak DCS-460 Digital Camera on Image Data for the Incandescent Illuminant Group^a

Algorithm	Sharpening Method	rms RGB Difference between Mapped Image and Target Image	rms RGB Difference between Mapped Image and Target Image (Sharp Correction)	rms Angle between Illumination RGB and Estimate Thereof	Count of Times Standard Results Replaced Full-Sharp Result Due to Problems (Usually Due to Negative Components)
BEST-LINEAR	none	18.8	18.8	*	0
BEST-DIAGONAL	none	21.2	21.2	1.75	0
BEST-DIAGONAL	ave	20.8	21.8	1.16	0
BEST-DIAGONAL	opt	20.8	21.6	1.41	0
BEST-DIAGONAL	mip	20.7	21.7	1.32	0
ACTUAL	none	25.2	25.2	0.00	0
ACTUAL	ave	25.6	25.6	=	0
ACTUAL	opt	26.1	26.1	=	0
ACTUAL	mip	25.8	25.8	=	0
NOTHING	none	83.7	83.7	7.92	0
GW	none	101.8	101.8	5.72	0
GW	ave	102.3	102.3	=	0
GW	opt	105.2	105.2	=	0
GW	mip	102.8	102.8	=	0
DB-GW	none	77.8	77.8	5.12	0
DB-GW	ave	77.9	77.5	5.16	0
DB-GW	opt	78.5	77.8	5.18	0
DB-GW	mip	78.0	77.5	5.15	0
SCALE-BY-MAX	none	48.1	48.1	2.33	0
SCALE-BY-MAX	ave	45.2	123.2	2.35	0
SCALE-BY-MAX	opt	44.0	81.7	1.91	0
SCALE-BY-MAX	mip	44.6	49.9	2.17	0
ECRULE-MV	none	85.3	85.3	5.55	0
ECRULE-MV	ave	38.2	85.5	2.45	0
ECRULE-MV	opt	38.8	84.8	3.02	0
ECRULE-MV	mip	52.8	85.2	3.93	0
ECRULE-ICA	none	102.7	102.7	5.69	0
ECRULE-ICA	ave	68.7	107.6	2.90	0
ECRULE-ICA	opt	68.8	100.5	3.69	0
ECRULE-ICA	mip	84.0	102.4	4.14	0

^aSee notes for Table 1. The results are the rms of 144 results. Assuming that the variation of the errors for a given algorithm is Gaussian, the error in the results is roughly 5%.

real image experiment sufficiently similar to those in the synthetic case (or vice versa), then the results would also be similar. We chose instead to apply the existing machinery to the new conditions in order to complement the synthetic experiments as well as possible.

11. RESULTS WITH IMAGE DATA

We provide numerical results for the image data in Table 5. In general, the results confirm that sharpening can provide modest benefit to the SCALE-BY-MAX algorithms and very significant benefit to the gamut-mapping algorithms. Interestingly, the improvement with average-illuminant sharpening is significantly better than with multiple-illuminant-with-positivity sharpening in some

cases. However, the multiple-illuminant-with-positivity method does very well in general, significantly improving computational color constancy performance in every case where sharpening can help. These two observations expose the trade-offs of multiple-illuminant-with-positivity sharpening: Peak possible performance is exchanged for robust performance under a wider variety of circumstances. The results based on the average illuminant illustrate the other side, which is that in specific circumstances, another method can do better.

12. CONCLUSIONS

We have investigated applying sensor sharpening to a variety of computational color constancy algorithms. We

found that doing so with existing methods leads to a number of problems, and as a result, using these methods is attractive only in specific circumstances. These difficulties led us to propose a new sharpening method (multiple-illuminant-with-positivity sharpening) which is less ambitious in terms of theoretical gains but addresses the needs of the color constancy algorithms studied here. The resulting sharpening transform substantially improves color constancy performance in many cases and only rarely has a small negative effect. Thus this work validates the original¹ intuition that sharper sensors should help facilitate color constancy.

In more detail, we found that sensor sharpening can be beneficial, sometimes significantly so, but only for some combinations of cameras, algorithms, and sharpening methods. To summarize the main points:

1. Some cameras, such as the Sony DXC-930, have quite sharp sensors that already accommodate diagonal color constancy. Getting better results by further sensor sharpening is difficult at best.

2. For the algorithms studied, if sharpening is used, it is almost always better to do all the color constancy calculations in the sharpened space. This works much better than using sharpening only for mapping results based on an illuminant estimate found in nonsharp space ("sharp correction").

3. The gray-world algorithm cannot be significantly improved by sharpening. Furthermore, depending on the variant and error measure, sharpening with gray world may reduce color constancy performance. Sharpening with gray-world algorithms is not recommended.

4. With initially nonsharp camera sensors and multiple-illuminant-with-positivity sharpening, a modest benefit can be expected with the SCALE-BY-MAX algorithm.

5. Significant benefit can be obtained in conjunction with gamut-mapping algorithms. Since these algorithms rely on the diagonal model, we expected the most benefit here, but the large degree of improvement often exceeded our expectations.

6. Multiple-illuminant-with-positivity sharpening worked very well in general, but in the experiments with real images the method based on the average illuminant performed even better. This suggests that additional characterization of the relationships between sharpening transforms and color constancy algorithm performance should be helpful.

7. Although sharpening was originally introduced from the perspective of improving image correction, our results show that with some camera/algorithm/error combinations, sharpening can also improve illumination parameter estimation. Since such estimation is the more difficult part of color constancy processing, this is a significant result. It is supported by two observations. First, the mapping errors can be reduced substantially beyond that indicated by the pure diagonal error. Second, the illumination-chromaticity estimation error can be substantially reduced. The separation of color constancy into these two parts is somewhat artificial in the case of the gamut-mapping algorithms, but it is precisely the tight integration of these parts that makes sharpening attractive when gamut mapping is used.

ACKNOWLEDGMENTS

We are grateful for the support of Hewlett-Packard Corporation and the Natural Sciences and Engineering Council of Canada. In addition, we acknowledge the efforts of Lindsay Martin who helped with the data collection.

REFERENCES AND NOTES

1. G. D. Finlayson, M. S. Drew, and B. V. Funt, "Spectral sharpening: sensor transformations for improved color constancy," *J. Opt. Soc. Am. A* **11**, 1553–1563 (1994).
2. K. Barnard and B. Funt, "Experiments in sensor sharpening for color constancy," in *Proceedings of the IS&T/SID Sixth Color Imaging Conference: Color Science, Systems and Applications* (The Society for Imaging Science and Technology, Springfield, Va., 1998), pp. 43–46.
3. K. Barnard and B. Funt, "Camera calibration for color research," *Color Res. Appl.* (to be published).
4. P. L. Vora, J. E. Farrell, J. D. Tietz, and D. H. Brainard, "Digital color cameras. 2. Spectral response," Hewlett-Packard Laboratory Tech. Rep. HPL-97-54, available from <http://www.hpl.hp.com/techreports/97/HPL-97-54.html> (1997).
5. D. Forsyth, "A novel algorithm for color constancy," *Int. J. Comput. Vision* **5**, 5–36 (1990).
6. G. D. Finlayson, M. S. Drew, and B. V. Funt, "Color constancy: generalized diagonal transforms suffice," *J. Opt. Soc. Am. A* **11**, 3011–3020 (1994).
7. Available from <http://www.cs.sfu.ca/~colour/data>.
8. K. Barnard, L. Martin, B. Funt, and A. Coath, "Data for colour research," *Color Res. Appl.* (to be published).
9. G. D. Finlayson, "Color in perspective," *IEEE Trans. Pattern Anal. Mach. Intell.* **18**, 1034–1038 (1996).
10. K. Barnard, "Computational colour constancy: taking theory into practice," M.Sc. thesis (Simon Fraser University, School of Computing Burnaby, B.C., Canada), available from <ftp://fas.sfu.ca/pub/cs/theses/1995/KobusBarnardMSc.ps.gz> (1995).
11. G. Finlayson and S. Hordley, "A theory of selection for gamut mapping colour constancy," *Image Vision Comput.* **17**, 545–588 (1999).
12. K. Barnard, "Practical colour constancy," (Ph.D. thesis Simon Fraser University, School of Computing, Burnaby, B.C., Canada), available from <ftp://fas.sfu.ca/pub/cs/theses/1999/KobusBarnardPhD.ps.gz> (1999).
13. K. Barnard, V. Cardei, and B. Funt, "A comparison of computational color constancy algorithms. Part one. Methodology and experiments with synthesized data," available from http://www.cs.berkeley.edu/~kobus/research/publications/comparison_1.
14. K. Barnard, L. Martin, A. Coath, and B. Funt, "A comparison of color constancy algorithms. Part two. Experiments with image data," available from http://www.cs.berkeley.edu/~kobus/research/publications/comparison_2.
15. For example, in the preliminary study, the results indicated that the two sharpening methods tested ("ave" and "opt") yield a modest benefit in conjunction with the SCALE-BY-MAX algorithm (labeled "Retinex" in that paper). However, the SCALE-BY-MAX algorithm was helped by the removal of data rejected by the gamut-mapping algorithms. Without this help, the conclusion is reversed in several circumstances.
16. G. Finlayson and M. Drew, "Positive Bradford curves through sharpening," in *Proceedings of the IS&T/SID Seventh Color Imaging Conference: Color Science, Systems and Applications* (The Society for Imaging Science and Technology, Springfield, Va., 1999), pp. 227–232.
17. B. Funt, K. Barnard, and L. Martin, "Is colour constancy good enough?" in *Proceedings of the 5th European Conference on Computer Vision* (Springer, Berlin, 1998), pp. I:445–459.

18. J. J. McCann, S. P. McKee, and T. H. Taylor, "Quantitative studies in Retinex theory," *Vision Res.* **16**, 445–458 (1976).
19. E. H. Land, "The Retinex theory of color vision," *Sci. Am.* **237**, 108–129 (1977).
20. B. K. P. Horn, "Determining lightness from an image," *Comput. Vision Graph. Image Process.* **3**, 277–299 (1974).
21. B. Funt, V. Cardei, and K. Barnard, "Learning color constancy," in *Proceedings of the IS&T/SID Fourth Color Imaging Conference: Color Science, Systems and Applications* (The Society for Imaging Science and Technology, Springfield, Va., 1996), pp. 58–60.
22. G. D. Finlayson, P. H. Hubel, and S. Hordley, "Color by correlation," in *Proceedings of the IS&T/SID Fifth Color Imaging Conference: Color Science, Systems and Applications* (The Society for Imaging Science and Technology, Springfield, Va., 1997), pp. 6–11.
23. M. J. Vrhel, R. Gershon, and L. S. Iwan, "Measurement and analysis of object reflectance spectra," *Color Res. Appl.* **19**, 4–9 (1994).
24. E. L. Krinov, *Spectral Reflectance Properties of Natural Formations* (National Research Council of Canada, Ottawa, 1947).