Error Propagation due to Color Space Transforms

Harro Stokman and Theo Gevers

Intelligent Sensory Information Systems,
University of Amsterdam,
Kruislaan 403, 1091 SJ Amsterdam, The Netherlands

ABSTRACT

It is known that the transformation of RGB color space to the normalized color space is invariant to changes in the scene geometry. The transformation to the hue color space is additionally invariant to highlights. However, due to sensor noise, the transforms become unstable at many RGB values. This effect is usually overcome by ad hoc thresholding, for example if the RGB coordinates are located near the achromatic axis then the corresponding hue value is rejected. To arrive at a principled way to deal with the unstabilities that result from these color space transforms, the contribution of this report is as follows. Uncertainties in the measured RGB values are caused by photon noise, which arises from the statistical nature of photon production. Using a theoretical camera model, we determine the number of photons required to cause a color value transition. Based on the associated uncertainty according to the Poisson distribution, we then derive theoretical models that propagate this uncertainty to the uncertainty in the transformed color coordinates. We propose a histogram construction method based on Parzen estimators that incorporates this theoretical reliability. As a result, we overcome the need for thresholding of the transformed color values.

Keywords: Color histogram, sensor noise, color space transform, error propagation, Parzen estimator

1. INTRODUCTION

The RGB values measured by a color CCD camera depend on the color of the object, the object surface orientation, the viewing angle, the direction of the incident light, and on the color of the light source. As a result, a homogeneous colored object may give rise to a broad spectrum of RGB values. For industrial machine vision tasks, it may be advantageous to recognize objects independent of the scene geometry. To that end, the RGB sensor values can be transformed into e. g. normalized or hue color values. Both color spaces are well known and are widely used in many image processing tasks. Gevers[1] proved in theory that the transform into the normalized color space results in invariance for the camera viewpoint, object pose and orientation, and for the direction and intensity of the incident light. The transform into the hue color space has the additional advantage of invariance to highlights.

However, these color space transforms bring with them serious drawbacks. Kender[2] analyzed in detail the problems that arise when RGB data is transformed to the normalized and hue color space. He showed that the transformations are singular at some points and unstable at many others. For instance, the essential singularity of normalized coordinates is at the black point ($R = G = B = 0$), whereas for the hue the singularity is at the entire achromatic axis ($R = G = B$). Both color spaces are unstable near the singularity where a small perturbation of input RGB can cause a large jump in the transformed values. Traditionally, these effects are either ignored or overcome by ad hoc thresholding. For example Ohta[3] rejects normalized color values if the sum of RGB is less than 30, and rejects hue values if the saturation times the intensity is less than the value 9. Gevers[4] discards normalized values if the intensity is less than 10 percent of the intensity range, and rejects hue values if the intensity or saturation is less than 10 percent. Healey[5] considers sensor noise with standard deviation σ and rejects all sensor measurements that lie within the sphere of radius 4σ centered at the origin in sensor space. Kender proposes to set the threshold value at that value that guarantees, in response to one unit input change, an output change of no more than one unit.

The drawback of the unstability of the transformed color spaces also occurs when the color spaces are used for construction and matching of histograms. A histogram is one of the simplest and most useful tools in image processing. Given a discrete color space, a color histogram counts how much of each color occurs in the image. Histograms are often used for image segmentation when the image properties are not sufficiently stable to use fixed threshold values. E. g. the illumination can vary or the appearance of objects can change over time. If a histogram would show two prominent peaks, one for the background color and one for the objects, then the threshold value for image segmentation would typically be placed at the minimum between the two most dominant peaks. Another application of histograms is demonstrated by Swain and Ballard.[6] The authors show that normalized color histograms of multicolored objects provide a robust, efficient cue for indexing into a large database of models. However, the unstability of the transformed colors negatively affects the performance of these histogram based methods. Moreover, Kender analyzes the effect of the computation of normalized colors and hue with respect to spurious modes and gaps that occur in the histogram distribution of the transformed coordinates. These modes and gaps occur because certain transformed color values occur much more often than other values. He shows that histograms are negatively affected by these spurious modes and gaps. Kender proposes to overcome the problem by adding random noise in the range of 0 to 1 to the RGB values prior to the color space transforms.

With regard to these problems, the contribution of this report is as follows. Uncertainties in the measured RGB values are caused by photon noise, which arises from the statistical nature of photon production. Using a theoretical camera model, we determine the number of photons required to cause a color value transition. Based on the associated uncertainty according to the Poisson distribution, we then derive theoretical models that propagate this sensor uncertainty to the uncertainty in the transformed coordinates. We propose a principled way for histogram construction based on Parzen estimators that incorporates this theoretical reliability. As a result, we overcome the need for thresholding of the transformed color values. Additionally, we show that the new proposed method is less affected by spurious modes and gaps.

This paper is organized as follows: In section 2 the theoretical models that predict the uncertainty in the transformed color values are derived. In section 3 these models are incorporated in the proposed method for histogram construction. In section 4 the predicted uncertainty in transformed color values is empirically compared to the actual uncertainty, and the proposed methods for histogram construction are compared conventional methods. Section 5 finishes the report with a discussion.

2. ERROR PROPAGATION

It is common practice to state the result of a series of measurements of a quantity x as

$$(\text{measured value of } x) = x_{\text{best}} \pm \sigma_x \tag{1}$$

where x_{best} is the best estimate for the quantity x (the average value) and σ_x is the uncertainty or error in the measurement of x (the standard deviation).

Suppose that $x, \cdots, z$ are measured with corresponding uncertainties $\sigma_x, \cdots, \sigma_z$, and the measured values are used to compute the function $q(x, \cdots, z)$. If the uncertainties in $x, \cdots, z$ are independent and random, then the predicted uncertainty in q is[7]

$$\sigma_q = \sqrt{\left(\frac{\partial q}{\partial x}\sigma_x\right)^2 + \cdots + \left(\frac{\partial q}{\partial z}\sigma_z\right)^2} \tag{2}$$

where $\partial q/\partial x$ and $\partial q/\partial z$ are the partial derivatives of q with respect to x and z.

Based on the measured RGB values, the normalized color values are computed as

$$\begin{aligned} r &= R/(R+G+B) \\ g &= G/(R+G+B) \end{aligned} \tag{3}$$

and the hue as

$$\theta = \arctan\left(\frac{\sqrt{3}(G-B)}{(R-G)+(R-B)}\right) \tag{4}$$

Note that the color space hue orders colors in a circular mode. As a consequence, calculations involving hue require special care. In this paper, the best estimate of a series of measured hue values is computed as

$$\overline{\theta} = \arctan\left(\frac{\sum_i \cos(\theta_i)}{\sum_i \sin(\theta_i)}\right) \tag{5}$$

The standard deviation of the hue σ_θ is computed as

$$\sigma_\theta = \frac{1}{N-1} \sum_{i=N} (\overline{\theta} - \theta_i)^2 \tag{6}$$

where

$$\theta_i - \theta_j = \pi - |\{\pi - |\theta_i - \theta_j|\}| \tag{7}$$

2.1. Model for the prediction of sensor noise

Modern CCD cameras are sensitive enough to be able to count individual photons. Photon noise arises from the fundamentally statistical nature of photon production because a different number of photons will be counted for two different recordings during the same time interval T. The probability distribution for counting ρ photons during T seconds is known to be Poisson. For a Poisson distribution, the average number of photons measured by an image sensing element at position (x, y) is given as

$$f(x, y) = \rho T \pm \sqrt{\rho T} \tag{8}$$

where ρ is the rate of photons measured per second. Our aim is to analyze how the uncertainty $\sigma_f = \sqrt{\rho T}$ propagates to the uncertainty in transformed color coordinates.

The following camera model describes the relation between a physical input signal (e. g. photons) and the output signal (e. g. voltage) as

$$g(x, y) = \text{gain} \cdot f(x, y) + dc(x, y) \tag{9}$$

where $f(x, y)$ is the input signal at position (x, y), $g(x, y)$ is the output signal and $dc(x, y)$ is the dark current To determine the uncertainty associated with a measured gray value, the number of electrons, or sensitivity S, necessary to change from one brightness level to the next is given by Young, Gerbrands and van Vliet[8] as

$$S = 1/\text{gain} \tag{10}$$

The sensitivity can be measured as follows: After compensating for the dark-current offset, the camera output is given by

$$g(x, y) = \text{gain} \cdot f(x, y) \tag{11}$$

Substitution of the measured number of photons and the associated uncertainty of Eq. 8 in Eq. 11 gives

$$g(x, y) = \text{gain} \cdot (\rho T \pm \sqrt{\rho T}) \tag{12}$$

Using the mean $\overline{g} = \text{gain} \cdot \rho T$ and standard deviation $\sigma_g = \text{gain} \cdot \sqrt{\rho T}$ of a region with homogeneous pixel values, the sensitivity S for a particular CCD camera is derived as[8]

$$S = \frac{\overline{g}}{(\sigma_g)^2} \tag{13}$$

From Eq. 12 follows that the uncertainty in the number of photons measured at an arbitrary pixel $g(x, y)$ is given by

$$\sigma_g(x, y) = \sqrt{\frac{g(x, y)}{S}} \tag{14}$$

Under the assumption that the dark-current noise σ_{dc} and photon noise σ_g are independent, the total amount of noise σ_n for a certain pixel value is

$$\sigma_n(x, y) = \sqrt{\sigma_{dc}^2 + \sigma_g^2} \tag{15}$$

2.2. Propagation of sensor noise to the hue and normalized colors

Substitution of Eq. 4 in Eq. 2 gives as the associated uncertainty for the hue

$$\sigma_\theta = \frac{1}{2}\sqrt{3}\left(\frac{-2BR\sigma_G^2 + R^2\sigma_B^2 + \sigma_G^2 + G^2(\sigma_B^2 + \sigma_R^2 + B^2(\sigma_G^2 + \sigma_R^2) - 2G(R\sigma_B^2 + B\sigma_R^2)}{B^2 + G^2 - GR + R^2 - B(G+R)^2}\right) \tag{16}$$

where σ_R, σ_G and σ_B denote the uncertainty according to Eq. 15 given the dark current and the measured value for the red, green and blue channel, respectively. Similarly, substitution of Eq. 3 in Eq. 2 gives as the associated uncertainty for the normalized coordinates

$$\sigma_r = \sqrt{\frac{R^2(\sigma_B^2 + \sigma_G^2) + (B+G)^2\sigma_R^2}{(B+G+R)^4}}$$

$$\sigma_g = \sqrt{\frac{G^2(\sigma_B^2 + \sigma_R^2) + (B+R)^2\sigma_G^2}{(B+G+R)^4}} \tag{17}$$

In conclusion, once the sensitivity of a color CCD camera is established by Eq. 13, it is possible using Eq. 14 to predict for the measured RGB values the associated uncertainty. In combination with the standard deviation of the dark current, the overall uncertainty is derived using Eq. 15. Substitution of the RGB values and this associated uncertainty into Eq. 16 and Eq. 17 gives, in theory, the uncertainty in the color space coordinates. In the next section, methods are be developed which incorporate this uncertainty in the construction of the transformed color histograms.

3. PARZEN-ESTIMATOR APPROACH FOR INCORPORATING UNCERTAINTIES IN THE CONSTRUCTION OF HISTOGRAMS

The histogram of an image is a plot of the color value versus the number of pixels at that value. The histogram can be used as a model of the probability distribution of the colors. The (one-dimensional) histogram probability $P(i)$ is defined as

$$P(i) = H(i)/M \tag{18}$$

where M is the number of pixels in the image that contribute to the histogram, and $H(i)$ is the number of pixels at index i. As with any probability distribution, all the values for $P(i)$ are less than or equal to 1, and the sum of all the $P(i)$ values is equal to 1.

In most experiments, if the number of measurements is increased, then the histogram begins to take on some definite shape. As the number of measurements approaches infinity, the histogram distribution approaches some definite, continuous curve, called the limiting distribution. Different types of measurements have different limiting distributions. The limiting distribution for counting ρ photons during T seconds (Eq. 8) is known to be Poisson. If the average number of counts is large, then the Poisson distribution is well-approximated by the Gauss distribution. Thus, if the histogram of a homogenously colored image would be computed, a unimodal Gaussian distribution (have a single local maximum) would be derived.

However, many images encountered in practice involve multimodal histogram distributions. The Parzen approach[9] to estimate a probability distribution is to let a measured value x' located close to the real value x contribute more to the probability distribution $P(i)$ than a measured value x'' located at a greater distance. To achieve this, the contribution of the measured values is weighted by a kernel function $K(x, x')$. Summation of the contributions of N measurements gives the estimate for the histogram probability

$$\tilde{P}(i) = \frac{1}{N}\sum_{i=j}^{N} K(x, x'_j) \tag{19}$$

The Poisson distribution is well-approximated by the Gauss distribution. Therefore, we use the Gaussian kernel for the Parzen-estimator:

$$K(x, x') = \frac{1}{\sigma\sqrt{2\pi}}\exp(-(x - x')/(2\sigma^2)) \tag{20}$$

where σ is a free parameter which affects both the amplitude and the width of the kernel. If σ is very large, $\tilde{P}(i)$ is the superposition of N broad, slowly changing functions, and is a very smooth estimate for $P(i)$. On the other hand, if σ is very small, $\tilde{P}(i)$ is the superposition of N sharp pulses and is a "noisy" estimate of $P(i)$. Based on the theoretical results of the previous subsection, both the transformed color space coordinate and the uncertainty associated with the value are known. We therefore use kernels with variable σ to estimate the histogram probability. Again, this parameter depends on the theoretical estimate of the uncertainty of the transformed color.

3.1. Hue Histogram Update

The value of hue ranges from 0 to 2π. If histogram indices for the hue need to be updated, the index i is obtained by multiplication of hue θ by a scale factor and rounded:

$$i = \text{round}(\text{scale} \cdot \theta) \bmod \text{round}(\text{scale} \cdot 2\pi) \tag{21}$$

Two conventional and one novel histogram construction methods are investigated for the hue color space: The first method which will be denoted "Achromatic" updates all histogram indices unless the RGB values are located on the achromatic axis:

$$H(i) = \begin{cases} 0 & \text{if } R = G = B \\ 1 & \text{otherwise} \end{cases} \tag{22}$$

Under the assumption that RGB range from zero to 255, let the intensity be scaled between zero and one as

$$\text{intensity} = (R + G + B)/(3 \cdot 255) \tag{23}$$

and let saturation be defined as

$$\text{saturation} = 1 - \frac{\min\{R, G, B\}}{R + G + B} \tag{24}$$

The second method, which will be denoted as "Cut-off" updates the hue histogram unless the RGB values are close to the black point or close to the achromatic axis:

$$H(i) = \begin{cases} 0 & \text{if intensity} < 0.1 \ \text{ or saturation} < 0.1 \\ 1 & \text{otherwise} \end{cases} \tag{25}$$

The new proposed method which will be denoted as "Parzen" updates the hue histogram with the Gaussian approximation of the Poisson distribution. Compared to the previous methods, for each contributing hue value, all indices of the histogram are updated instead of a single index. If N is the number of bins in an intermediate histogram G, and j ranges from bin 1 to bin N, then

$$G(j) = \frac{1}{\sigma_\theta \sqrt{2\pi}} \exp(-(j - \pi)/(2\sigma_\theta^2)) \tag{26}$$

where σ_θ is derived according to Eq. 16. The equation uses π as the mean hue value due to the circularity of the hue. The actual hue histogram H is then updated as

$$\begin{aligned} H(0 \cdots \theta) &= G(2\pi - \theta \cdots 2\pi) \\ H(\theta \cdots 2\pi) &= G(0 \cdots 2\pi - \theta) \end{aligned} \tag{27}$$

3.2. Normalized Color Histogram Update

If the histogram of normalized coordinates need to be updated, the index i is obtained by multiplication of e. g. the normalized coordinates r by a scale factor and rounded:

$$i = \text{round}(\text{scale} \cdot r) \tag{28}$$

Note that this scale factor depends on the transformation and need not be the same as for the hue.

Two conventional and one novel histogram construction methods are investigated for the normalized color space: The "Achromatic" method updates all histogram indices unless the RGB values lie on the black point:

$$H(i, j) = \begin{cases} 0 & \text{if } R = G = B = 0 \\ 1 & \text{otherwise} \end{cases} \tag{29}$$

patch	$\overline{f}$	$S = \overline{f}/\sigma^2$
24	8 ± 3	1
23	24 ± 2	10
22	53 ± 2	18
21	96 ± 3	13
20	154 ± 3	18
19	229 ± 4	15

Table 1. Experiment: Establishing sensitivity parameter S. In the middle column, the average intensity value for the obtained images is given. The right column shows the computed sensitivity. The average value S is 15 ± 3

The "Cut-off" method updates the normalized histogram unless the RGB values are close to the black point:

$$H(i,j) = \begin{cases} 0 & \text{if intensity} < 0.1 \\ 1 & \text{otherwise} \end{cases} \tag{30}$$

The new proposed "Parzen" method updates the hue histogram with the Gaussian approximation of the Poisson distribution. Compared to the previous two methods, for each contributing normalized color pair, all indices of the histogram are updated instead of a single index. If N is the number of bins for the normalized red coordinate, then let ρ range from bin number 1 to bin N. Similarly, if M the number of bins for the normalized green coordinate, then let γ range from bin 1 to bin M. Now

$$H(\rho,\gamma) = \frac{1}{\sigma_r\sqrt{2\pi}} \exp(-(\rho - r)/(2\sigma_r^2)) * \frac{1}{\sigma_g\sqrt{2\pi}} \exp(-(\gamma - g)/(2\sigma_g^2)) \tag{31}$$

where $*$ denotes convolution and where σ_r and σ_g are derived according to Eq. 17.

4. EXPERIMENTS

In this section, the theoretically predicted uncertainty in transformed color coordinates will empirically be compared to the actual uncertainty, and the proposed Parzen method for histogram construction will be compared to conventional methods. All experiments are performed using a Hewlett & Packer ScanJet 4P flatbed scanner. The color patches of the MacBeth ColorChecker are used. This ColorChecker is a checkerboard array of 24 colored squares in a wide range of colors. Many of these squares represent natural objects of special interest, such as human skin, foliage, and blue sky. Color patches numbers 19 to 24 represent gray colors ranging from white to black. For the experiments, the scale factor for the hue (Eq. 21) is set to $360/(2\pi)$, and the scale factor for normalized colors (Eq. 28) is set to 255.

4.1. Establishing Sensitivity Parameter S

For the experiment, the achromatic patches numbers 19 to 24 of the MacBeth ColorChecker are taken. Goal of the experiment is to determine the sensitivity S of Eq. 10 of the scanner. To that end, the average intensity value and the standard deviation of each patch are determined, which are given in Table 1. The experiment shows that the sensitivity of the camera approximates 15 ± 3. The first value of patch 24 is considered an outlier and is further ignored.

4.2. Propagation of Uncertainties

For the experiment, all patches of the MacBeth ColorChecker are taken. Theoretical models were proposed in Eq. 16 17 to predict the propagation of uncertainties from the combination of photon and thermal noise to uncertainty in hue and normalized color values. The goal of this experiment is to analyze the validity of the developed formulae.

In the previous experiment, the sensitivity parameter S is determined. In the current experiment, first the predicted uncertainty in the hue is evaluated. For each RGB value of a pixel in a patch of the ColorChecker, the uncertainty in the hue σ_ϕ is predicted according to Eq. 16. These predicted uncertainties are then averaged over all the pixels of the patch. The actual uncertainty is derived as the standard deviation of the hue values of the color patch. These predicted and actual uncertainties of the hue are shown in Fig. 1 for the color patches. For patches

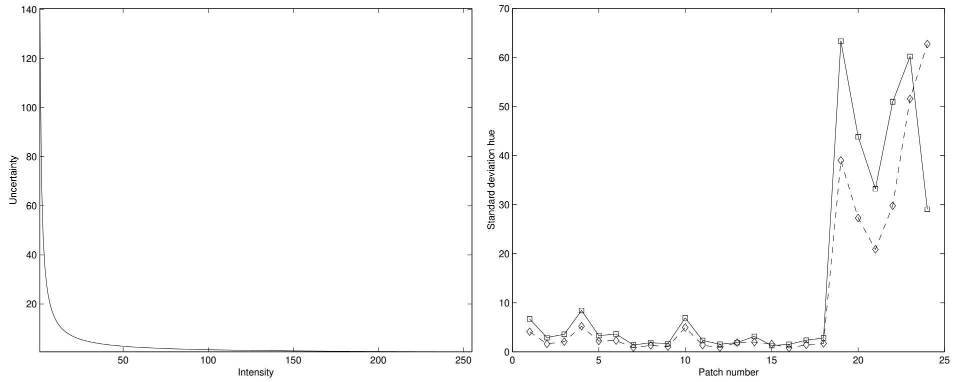

Figure 1. Left: Plot of the uncertainty in the hue according to Eq. 16. R ranges from 0 to 255 and $G = B = 0$. The uncertainty in the hue value increases as the intensity decreases. The same figure would be obtained if the uncertainty in the normalized coordinates of Eq. 17 would be visualized. Right: Analysis of the theoretically predicted uncertainty versus the actual uncertainty of the hue. The dashed line corresponds to actual standard deviations of the hue, the solid line corresponds to the predicted standard deviations.

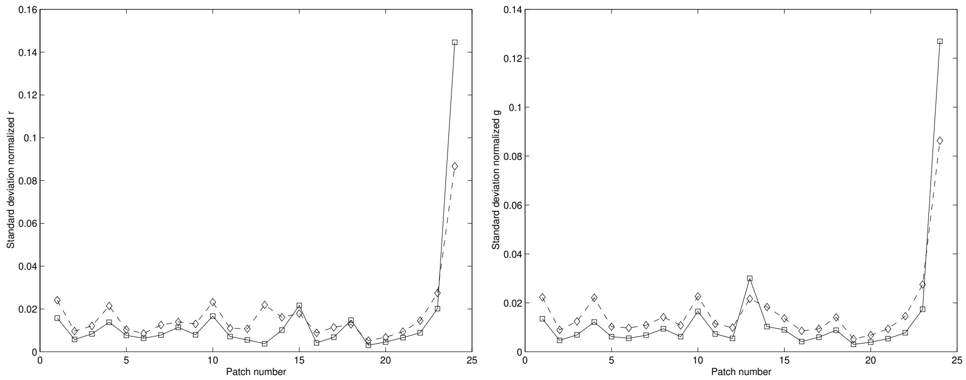

Figure 2. Experiment: Analysis of the theoretically predicted uncertainty versus the actual uncertainty of the normalized red coordinate (left) and normalized green coordinate (right). The dashed line corresponds to actual standard deviations, the solid line corresponds to the predicted standard deviations.

$1 \cdots 18$, the absolute difference between the predicted and actual values is 1.1 ± 0.8 degrees hue. If the achromatic patches 19 to 24 are included, whose pixel values are near or on the singularity of the hue, then the difference is 6 ± 9 degrees hue. The experiment shows that the proposed theoretical model for prediction of uncertainty due to photon noise to the uncertainty in the hue values gives reasonable predictions of the actual uncertainties.

The experiment is repeated for normalized colors in a similar fashion. The results are shown in Fig. 2. For patches $1 \cdots 23$, the absolute difference between the predicted and actual values is 0.005 ± 0.003 for normalized red and 0.005 ± 0.002 for normalized green. If the black patch 24 is included, whose pixel values are near or on the singularity, than the difference is 0.007 ± 0.011 for normalized red and 0.007 ± 0.008. Again, the experiment shows that the proposed theoretical model for prediction of uncertainty due to photon noise to the uncertainty in the normalized red and green values gives reasonable predictions of the actual uncertainties.

4.3. Sensitivity of different histogram methods for photon noise

For the experiment, ten identical images are taken from the MacBeth ColorChecker . The images are recorded a few seconds apart. Thus, the only difference between these images is the image noise. From the images only the

	Achromatic (Eq. 22)	Cut-off (Eq. 25)	Parzen (Eq. 26)
Number of Local Maxima	44 ± 7	27 ± 4	14.2 ± 0.6
χ^2 Difference	500 ± 100	340 ± 70	15 ± 7

Table 2. Experiment: Sensitivity of different histogram methods in hue color space for photon noise. A sub-image of the MacBeth ColorChecker is used containing 18 differently colored patches. Ideally, 18 local maxima should be detected and the χ^2 difference between histograms computed for an image of a single object, recorded at different time intervals, should be zero.

	Achromatic (Eq. 29)	Cut-off (Eq. 30)	Parzen (Eq. 31)
Number of Local Maxima	88 ± 8	54 ± 6	17.6 ± 0.5
χ^2 Difference	700 ± 90	450 ± 40	3 ± 2

Table 3. Similar experiment as shown at table 2 but for normalized color space

sub-image containing patches 1 to 18 are used. Goal of the experiment is to compare the new proposed Parzen methods for histogram construction to the conventional Achromatic and Cut-off methods.

First consider the histograms computed for the hue color space. For each of the 10 sub images, a hue histogram is computed. The resulting 10 histograms are then compared by the χ^2 method. If no image noise were present, then it is expected that the histograms are identical, and that the χ^2 difference is equal to zero. Due to photon noise however, slight variations in the recorded RGB values occur. The histogram construction method which has the smallest χ^2 difference is the least sensitive to photon noise. Note that a method where all histogram indices would be updated with equal values, regardless of the hue value, would be the most invariant to photon noise, and the χ^2 difference between the histograms would always be zero. Thus, to ensure that color discriminability is maintained, the number of local maxima are therefore determined in histogram H at index i according to the following scheme:

$$\text{local maximum}(i) = \begin{cases} 1 & \text{if } H(i-1) < H(i) < H(i+1) \\ 0 & \text{otherwise} \end{cases} \tag{32}$$

The results are given in Table 2. The new proposed method outperforms the classical methods both with regard to the sensitivity to photon noise as with regard to the number of detected local maxima. First, the proposed method has the smallest χ^2 difference between the ten recordings. Second, the number of detected local maxima (14) is close to the ideally expected number of maxima (18). Here, the patches 1, 2 and 7 (dark and light skin, orange) and patches 3, 8 and 13 (blue sky, purplish blue and blue) are indistinguishable from one another because the hue values of the different patches have approximately equal values. Example histograms are shown in Fig. 3

The experiment is repeated in similar fashion for normalized color. For determining the number of local maxima, the method of Eq. 32 is extended to a 2 dimensional histogram. The results are given in Table 3. Again, the new proposed method outperforms the classical methods both with regard to the sensitivity to photon noise as with regard to the number of detected local maxima.

4.4. Comparison of performance of different histogram construction methods for dark and bright colors

The experiment is performed on patches nr. 14, 15 (green and blue) in combination with a variable density filter. The filter attenuates an incident beam of radiation without altering its spectral distribution. The density varies linearly across the filter's length. As a result, one part of the image is bright, the other part is dark. The images are manually divided in a dark and light part. The new proposed histogram construction methods are tailored to handle uncertain (dark) color values. Goal of this experiment is therefore to evaluate the performance of the different histogram construction methods for dark and bright colors. To that end, first hue histograms are computed for the dark and bright images, and the local maxima are counted. The results are given in Table 4.

For bright colors, there is no difference in the result between the achromatic and cut-off method, since the bright image does not contain any unreliable pixels. Still, the advantage of the Parzen method over the conventional methods is that the histogram is somewhat smoothed, and as a result, only a single local maximum is found. Due to photon noise, in the conventional methods several local maxima are found. The average and standard deviation of the hues are comparable for the three methods.

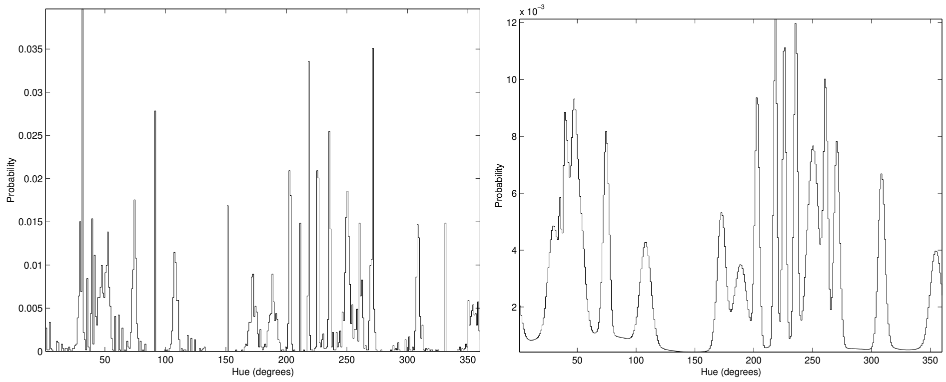

Figure 3. Experiment: Comparison of two different hue histogram construction methods. For both histograms, the same sub-image of the MacBeth ColorChecker is used which contains 18 different colored patches. Left is the probability distribution resulting from the "Achromatic" histogram construction method, right the probability distribution resulting from the "Parzen" method. The latter method appears to be less sensitive to noise.

	Achromatic (Eq. 22)		Cut-off (Eq. 25)		Parzen (Eq. 26)	
Color	Hue	#Max.	Hue	#Max.	Hue	#Max.
Bright green	162 ± 2	3	162 ± 3	3	163 ± 4	1
Dark green	150 ± 30	32	161 ± 3	1	150 ± 20	1
Bright blue	39 ± 1	2	39 ± 1	2	40 ± 2	1
Dark blue	40 ± 20	17	None	None	40 ± 10	1

Table 4. Experiment: Comparison of performance of different histogram construction methods for dark and light colors in hue color space. Ideally, all methods would compute the same average hue value with small standard deviations, while simultaneously one local maximum would be detected.

For dark colors, the achromatic version and the Parzen method are comparable for the average and standard deviation of the hue values. The difference is the detected number of local maxima: Due to the instability of the hue, a large number of maxima is found for the achromatic version. The cutoff method performs better than the Parzen method in the case of dark green pixels. This can be explained as follows. A small number of dark green pixels satisfies the reliability. Therefore, the same output is generated as in the case of bright colors. In contrast, in the case of dark blue pixels, none of the pixels satisfies the reliability threshold. As a result, no output is generated. From the experiment, it is concluded that the proposed Parzen method outperforms the traditional methods for both dark as for bright images.

4.5. Spurious modes and gaps

For the experiment, the hue values are generated for all possible RGB combinations. Two hue histograms are constructed for these hue values by the Achromatic and Parzen method. The computation of hue and normalized color leads to spurious modes and gaps that occur in the histogram distribution of the transformed coordinates. The essential problem is the division in Eq. 3 4. Because the input is digital, certain output quotients are "favored" over others, creating false modes in the histograms. The cause is that certain fractions can have more representatives

	Achromatic (Eq. 29)		Cut-off (Eq. 30)		Parzen (Eq. 31)	
Color	(r, g)	#Max.	(r, g)	#Max.	(r, g)	#Max.
Bright green	$(64 \pm 2, 152 \pm 2)$	12	$(64 \pm 2, 152 \pm 2)$	12	$(64 \pm 2, 152 \pm 2)$	1
Dark green	$(50 \pm 30, 150 \pm 20)$	113	$(64 \pm 3, 148 \pm 3)$	7	$(50 \pm 30, 150 \pm 20)$	1
Bright blue	$(0 \pm 2, 36 \pm 3)$	12	$(0 \pm 2, 36 \pm 3)$	13	$(0 \pm 2, 36 \pm 3)$	1
Dark blue	$(10 \pm 20, 30 \pm 30)$	62	None	None	$(10 \pm 20, 30 \pm 30)$	1

Table 5. Similar experiment as in table 4 but for normalized color space.

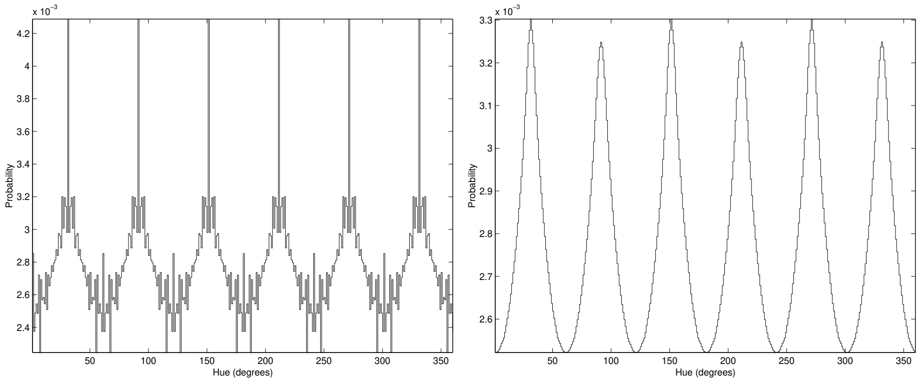

Figure 4. Experiment: Sensitivity to spurious modes. Left is the probability distribution shown for the "Achromatic" histogram construction method. Clearly, spurious modes and gaps are present. The right hand side graph shows the probability distribution for the "Parzen" method. The method is insensitive to spurious modes. The peaks correspond to the primary colors red, yellow, green, cyan, blue and magenta.

than others. For example, the normalized r value of $1/2$ can be attained through various fractions by many more input triples than the value of $1/511$. Typically the value $1/2$ represents a spurious mode, whereas $1/511$ represents a spurious gap. The distributions are shown in Fig. 4. Spurious modes and gaps are clearly visible for the achromatic method. For the Parzen histogram, spurious modes still occur at the colors red, yellow, green, cyan, blue and magenta, though at a much lesser extend. The experiment shows that the proposed method is less sensitive to spurious modes and gaps.

5. DISCUSSION

Using a theoretical camera model, the number of photons required to cause a color value transition is determined. Based on the associated uncertainty according to the Poisson distribution, theoretical models are derived that propagate the sensor uncertainty to the uncertainty in the transformed coordinates. A principled way for histogram construction based on Parzen estimators has been proposed that incorporates this theoretical reliability. Drawbacks are that the sensitivity of the camera needs to be known, and that the method is computational more expensive than conventional histogram construction methods. The advantages are that the need for thresholding of transformed color values is overcome and that the new proposed method is less sensitive to spurious modes and gaps.

REFERENCES

1. T. Gevers and A. W. M. Smeulders, "Color based object recognition," *Pattern Recognition* **32**, pp. 453–464, March 1999.
2. J. R. Kender, "Saturation, hue and normalized color: Calculation, digitation effects, and use," tech. rep., Carnegie-Mellon University, November 1976.
3. Y.-I. Ohta, T. Kanade, and T. Sakai, "Color information for region segmentation," *Computer Graphics and Image Processing* **13**, pp. 222–241, 1980.
4. T. Gevers, *Color Image Invariant Segmentation and Retrieval.* PhD thesis, University of Amsterdam, 1996.
5. G. Healey, "Segmenting images using normalized color," *IEEE Transactions on Systems, Man and Cybernetics* **22**(1), pp. 64–73, 1992.
6. M. J. Swain and D. H. Ballard, "Color indexing," *International Journal of Computer Vision* **7**(1), pp. 11–32, 1991.
7. J. R. Taylor, *An Introduction to Error Analysis*, University Science Books, 1982.
8. I. T. Young, J. J. Gerbrands, and L. J. van Vliet, *Fundamentals of Image Processing*, Delft University of Technology, 1995.
9. R. O. Duda and P. E. Hart, *Pattern Classification and Scene Analysis*, John Wiley & Sons, 1973.