15th European Signal Processing Conference (EUSIPCO 2007), Poznan, Poland, September 3-7, 2007, copyright by EURASIP

A PROBABILISTIC SHADING INVARIANT COLOR DISTANCE MEASURE

Slawo Wesolkowski[1] *and Paul Fieguth*[2]

[1]DRDC
Ottawa, Ontario, Canada
email: s.wesolkowski@ieee.org

[2]Systems Design Eng., University of Waterloo
Waterloo, Ontario, Canada
email: pfieguth@uwaterloo.ca

ABSTRACT

We develop a probabilistic color distance measure based on hypothesis testing in order to achieve shading invariance in image segmentation. We derive this new color distance measure based on the Dichromatic Reflection Model and noise statistics. We show preliminary results of using the new semi-metric in a color image segmentation task to show its effectiveness.

1. INTRODUCTION

Humans can easily, even effortlessly, distinguish between separate objects in an image scene. This has long been a key problem in computer vision where a number of steps, from low-level to high-level vision, are needed to understand an image or some portion of it. A critical step is that of image segmentation — the partition of an image into distinguishable subsets based on the premise that objects having a distinct appearance can be visually separated. Images are composed of pixels which depending on the sensors used to capture it can represent light intensity values, colors or some other electromagnetic quantities. Image segmentation requires two distinct components: pixel comparison and pixel grouping. The pixel comparison function requires the design of a pixel similarity criterion. The pixel grouping mechanism, on the other hand, aggregates the pixels with respect to this pixel similarity criterion.

In this paper, we focus on the second component of an image segmentation algorithm. Without a reliable distance measure it is not possible to determine which pixels should go together and which constitute disparate parts. The appropriate distance measure needs to reflect the kind of problem that is being solved and thus knowledge about the problem is encoded in the distance measure either implicitly (by carrying out comparisons in a feature space assumed to be Euclidean) or explicitly (carrying out comparisons in the sensor space which is not necessarily Euclidean).

Choosing an appropriate pixel distance measure for a particular application can determine whether the algorithm devised to solve the problem will be successful [2]. In the case of color image segmentation, the choice of distance measure will depend on the color space and color model being used to solve the given problem. This paper discusses statistics- and physics based derivation of a color distance semi-metric that are shading invariant in *RGB*.

Both Euclidean distance and vector angle have drawbacks:

1. Euclidean distance [2] is not an appropriate physics-based metric in *RGB* since it is highly intensity dependent [7].
2. Vector angle [9], although shading invariant in *RGB*, is not an appropriate physics-based semi-metric since distance calculations for low pixel values are unreliable.
3. Euclidean distance is shading invariant in normalized *rgb* [5], however, it is susceptible to the same problems as vector angle.

As opposed to the Euclidean distance and the vector angle, statistical distance measures depend on stochastic information encoded in the data in addition to difference computation between the pixels being compared. Since vector angle in *RGB* and Euclidean distance in normalized *rgb* give highly unreliable results for low intensity pixels, we will endeavor to develop a shading invariant metric that is also noise resistant.

There are several important issues and assumptions that need to be made under this basic premise:

1. A metric which is invariant to shading needs to factor out the illumination. We will assume that the illumination varies linearly with the pixel values (as the illumination intensity increases, so will the RGB pixel values increase). This assumption is not always valid especially for very low (dark regions) and very high (usually small highlight regions) intensity values. However, we are not dealing with perceptually correct color differences such as for example those achieved in CIE Luv [7] or CIE Lab where small differences might matter.
2. We will assume that the illumination of the color scene is white light or a white balancing step has occurred. In all cases, images studied in this paper were obtained under white light and therefore we have no reason to believe that this assumption would not hold.
3. We will assume that we are working in the *RGB* or sensor space where we can easily apply the Dichromatic Reflection Model [8]. Therefore, we need not be concerned with transformations into a different color space.
4. Noise resistance can be achieved with respect to different types of noise. Noise can frequently introduce errors in measurement and it is important to take it into account. In this paper, we consider pixel independent white noise stemming from the image capture process which we model as an additive Gaussian distributed noise:

$$\underline{x} = \underline{a} + \underline{\upsilon} \tag{1}$$

where $\underline{x}$ is the pixel vector, $\underline{a}$ is the true representation of $\underline{x}$, and $\underline{\upsilon}$ is Gaussian-distributed noise with covariance R that depends on $\underline{a}$. Therefore, the noise for each $\underline{x}$ is independent of the other noises.

However, additive Gaussian noise is not necessarily a good assumption since real CCD camera noise is strongly dependent on the image intensity level [6] and includes mainly five noise sources: fixed pattern noise, dark current noise, shot noise, amplifier noise and quantization noise. The most important noise element seems to be shot noise, the variance of which varies linearly with intensity [12]. However, finding image regions in order to estimate the noise level for different intensities and colors is generally not possible since many regions will be too small. Therefore, we will continue with our current (imperfect) assumption.

The paper is organized in the following manner. First, we will present the Dichromatic Reflection Model, what constitutes a metric and the hypothesis testing framework. Second, we will demonstrate why the vector angle and Euclidean distance are not appropriate distance measures especially when we would like to measure intensity invariant distances in highly variable and dark areas. Next, we will introduce the hypothesis testing probabilistic framework which will be followed by the development of an intensity invariant and noise resistant metric. We will show results based on well-known color images in the literature.

2. PHYSICS-BASED COLOR REFLECTION MODELS AND SPACES

Color models allow us to make certain assertions, regarding color generation and perception, about the color space that we would like

to use based on the laws of physics (optics). Physics-based color models explain how light is reflected from objects in a scene based on the physical properties of materials. They are used when algorithms need to achieve perceive objects in the real world without illumination effects such as is the case for humans.

Much work has been done on physics-based color modelling over the years [8, 10, 11]. A commonly used physics-based model is Shafer's Dichromatic Reflection Model (DRM) [8], which assumes that light reflected from objects can be separated into specular reflection and diffuse reflection. Specular reflection or highlight is characterized visually by a glossy appearance and describes light that is reflected in a mirror-like fashion from a surface. Diffuse or body reflection is the light reflected in all directions from a surface, giving a surface its colored appearance. By using the DRM, pixel difference computations can be done directly in the *RGB* space. The focus here will be on inhomogeneous dielectric materials such as plastics and painted surfaces and we will base our analysis on the work by Tominaga [10].

In the DRM, light reflected from an object surface o (called the color signal) is described as a function $c^o(\lambda,i,j)$ of wavelength λ and pixel location $\{i,j\}$:

$$\begin{aligned} c^o(\lambda,i,j) &= \text{Body Reflection} + \text{Interface Reflection} \\ &= \nu(i,j)s^o(\lambda)e(\lambda) + \eta(i,j)e(\lambda) \end{aligned} \quad (2)$$

where $e(\lambda)$ is the spectral power distribution of the light source, $s^o(\lambda)$ is the spectral-surface reflectance of object o, $\nu(i,j)$ is the shading factor and $\eta(i,j)$ is a scalar factor for the specular reflection term.

Normalized color or *rgb* [5] is obtained by dividing the RGB pixel elements by the pixel magnitude, i.e., $\underline{\bar{x}} = \frac{\underline{x}}{|\underline{x}|}$. For matte objects the color representation in the normalized *rgb* space is invariant with respect to illumination direction and intensity, as well as the viewing direction and surface orientation [3]. It is still sensitive to specular reflections, and inter-reflection. It is also not well defined for pixels with low intensives. We will focus on the *RGB* and normalized *rgb* color spaces.

3. EUCLIDEAN DISTANCE AND VECTOR ANGLE LIMITATIONS

The key to color image segmentation is to apply the appropriate color distance measure for the problem at hand. The choice of distance measure can greatly affect image segmentation or clustering results [2]; therefore, it is critical to make sure that the similarity or discontinuity measure being used is the appropriate for the assumed color space.

The most often used distance or discontinuity measure due to its mathematical properties and ease of use is the Euclidean distance [2]. However in the case of color images, where each pixel is represented as a RGB vector, the Euclidean distance is a particularly poor measure of color similarity because the RGB space is *an*-isotropic, especially when lighting effects such as specular reflection and shading are present in the image [7].

The vector angle measure or its variants has been used a few times in the literature [9]. Because the dot product between the vectors is divided by the magnitude of the vectors, vector angle will be shown to be intensity invariant with respect to the Dichromatic Reflection Model. The main problem with the vector angle is that it gives very "noisy" results for vectors with small magnitudes [9] and is undefined if $|\underline{x}| = 0$ or $|\underline{y}| = 0$.

The distance between two pixels can be computed in several different ways. Using the Euclidean distance, the distance becomes intensity-dependent and, therefore, is not applicable to assessing differences based on color. Figure 1 shows that there are high differences with respect to pixel intensity but not shading invariance for example. Note that a black square indicates that the colors are similar, while a white one shows high disagreement. Shades of gray illustrate the nuances in the color differencing results.

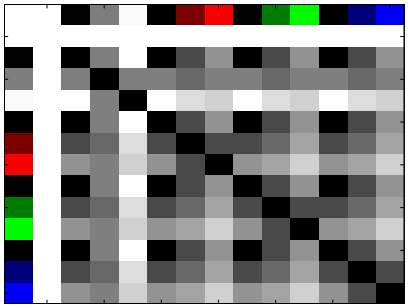

Figure 1: Grayscales show Euclidean distances between various colors in *RGB*. Lighter values imply larger distances. There is a clear pattern of having low distances between pixels of similar intensity and large distances between pixels of very different intensities all without regard to the intrinsic pixel color. The *RGB* colors correspond to the following values (from left to right and from top to bottom): $\{1,1,1\}$, $\{125,125,125\}$, $\{250,250,250\}$, $\{1,0,0\}$, $\{125,0,0\}$, $\{250,0,0\}$, $\{0,1,0\}$, $\{0,125,0\}$, $\{0,250,0\}$, $\{0,0,1\}$, $\{0,0,125\}$, and $\{0,0,250\}$.

The vector angle in *RGB* is equivalent to the Euclidean distance in the normalized color space *rgb* [5]. In addition, the normalized covariance matrix is introduced here and given by $\bar{R}_x = \frac{R}{|\underline{x}|^2}$. This is shown in Figure 2. Since $\sin\theta \approx \theta$ for similar colors, not much error is introduced. A similar outcome to vector angle in *RGB* is demonstrated in Figure 3. For very low intensity pixel values in *RGB*, the distance measure behaves erratically. That is, for small changes in low intensity pixel values, the angle can be arbitrarily different. This means that a small amount of noise will create vastly different results implying that using the vector angle the statistics break down for very dim (or low intensity) pixels. This invalidates the use of vector angle or Euclidean distance in normalized *rgb* even though the intensity invariant feature is very attractive.

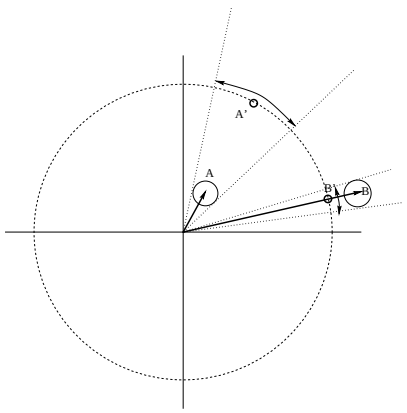

Figure 2: The normalized *rgb* color space is demonstrated through projecting *RGB* vectors onto the unit sphere (shown here in 2-D for ease of viewing). The RGB pixels A and B (thick black arrows) are projected onto the unit sphere at points A' and B' respectively (small circles on the unit circle). The dotted lines are indicative of the variance of the pixel values. Therefore, a pixel with low *RGB* intensities that is projected onto the unit sphere will have relatively higher variance on the unit sphere than pixels that have high *RGB* intensities (shown by the arcs with double arrows).

4. HYPOTHESIS TESTING FRAMEWORK

In order to create a probability-based distance measure, we will first need to introduce the concept of hypothesis testing. Then we will

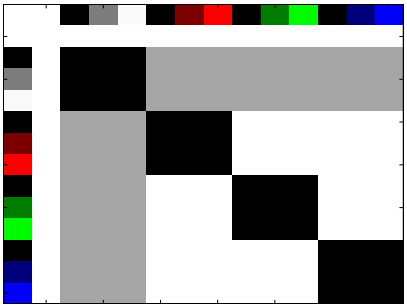

Figure 3: Shading Invariant Distances Between Various Colors in RGB (vector angle) or equivalently in normalized *rgb* (Euclidean distance). For example, notice that distances between pixels with low intensity values such as $\{1,0,0\}$ and other pixels can change drastically with the change of just one pixel value. The same results would be obtained with vector angle in *RGB*

describe a hypothesis test-based distance measures that can be used as a shading invariant method with the desired characteristics.

The most commonly used method to formulate hypothesis testing involves asking the question whether quantities (in this case pixel values) are from the same class. In essence, this corresponds to the following hypothesis tests:

$$\begin{aligned} H_0 &: \underline{x} = \underline{y} \\ H_1 &: \underline{x} \neq \underline{y} \end{aligned} \tag{3}$$

where we test the null hypothesis H_0.

We could also ask which class each pixel (of a set of adjacent pixels) belongs to. The hypothesis tests then become

$$\begin{aligned} H_{ij} : \quad & \underline{x} = \underline{a}_i \\ & \underline{y} = \underline{a}_j \end{aligned} \tag{4}$$

This formulation is seldom used since as the number of classes grows, the number of tests grows quadratically which can quickly become unmanageable. However, we can ask some questions about $\underline{a}_i$ and $\underline{a}_j$ in order to reduce the number of tests to only one. For example, instead of having formal classes we can estimate $\underline{a}_i$ and $\underline{a}_j$ based on $\underline{x}$ and $\underline{y}$ by asking whether $\underline{x}$ and $\underline{y}$ are explained by the same mean (i.e., $\underline{a}_i = \underline{a}_j$) or different means (i.e., $\underline{a}_i \neq \underline{a}_j$). We will call this the common mean hypothesis test.

We will begin with the common mean hypothesis test and see that it is related under certain conditions to the classical hypothesis test (3).

4.1 Common Mean Hypothesis Test Color Distance Measure

An original approach using hypothesis test (4) involves the estimation of $\underline{a}_i$ based on the values of $\underline{x}$ and $\underline{y}$ in order to find the $\underline{a}_i$ that maximizes the joint probability distribution $p(\underline{x},\underline{y}|\underline{a}_i)$. We can also look for the $\underline{a}_i$ that is equidistant in terms of standard deviations from $\underline{x}$ and $\underline{y}$, in other words the $\underline{a}_i$ for which $p(\underline{x}|\underline{a}_i) = p(\underline{y}|\underline{a}_i)$ is true. For simplicity of notation and without loss of generality, we will use $\underline{a} = \underline{a}_i$ for the developments in this section and the next one.

If we first assume that the means $\underline{a}$ are known, then we can estimate the likelihood of $p(\underline{x}|\underline{a})$ using

$$p(\underline{x} \mid \underline{a}) = \frac{1}{\sqrt{2\pi} \mid R \mid} e^{-\frac{1}{2}(\underline{x}-\underline{a})^T R^{-1} (\underline{x}-\underline{a})}. \tag{5}$$

(5) gives us the well-known measure of how likely it is for $\underline{x}$ to come from a distribution with mean $\underline{a}$. Note that we can transform this probability into a distance measure by using $-\ln p(\underline{x}|\underline{a})$. We now assume that the prior means $\underline{a}$ are unknown and ask the question of what is the likelihood that there is a common $\underline{a}$ which explains both $\underline{x}$ and $\underline{y}$?

Through the independence property (we can assume that $\underline{x}$ and $\underline{y}$ are conditionally independent since their noises are independent), we can state that $p(\underline{x},\underline{y} \mid \underline{a}) = p(\underline{x} \mid \underline{a}) p(\underline{y} \mid \underline{a})$. Therefore, given that we would like to find out how consistent $\underline{x}$ and $\underline{y}$ are with respect to each other, the desired likelihood is

$$p(\underline{x},\underline{y}) = \max_{\underline{a}} p(\underline{x},\underline{y} \mid \underline{a}) \tag{6}$$

where the max operator is computed over all possible $\underline{a}$. This can be a computationally expensive task depending on whether the set of reals or integers is used to represent $\underline{a}$. Making the conditional independence assumption we have

$$p(\underline{x},\underline{y}) = p(\underline{x} \mid \underline{a}) p(\underline{y} \mid \underline{a}) \tag{7}$$

which can also be written by

$$\Phi(\underline{x},\underline{y}) = -\ln[p(\underline{x} \mid \underline{a})] - \ln[p(\underline{y} \mid \underline{a})] \tag{8}$$

The distance metric can therefore be represented by

$$\Phi_C(\underline{x},\underline{y}) = (\underline{x}-\underline{a})^T R^{-1} (\underline{x}-\underline{a}) + (\underline{y}-\underline{a})^T R^{-1} (\underline{y}-\underline{a}) \tag{9}$$

This metric is very easy to calculate and has no trouble with dark pixels assuming that we can somehow find a good $\underline{a}$. However, it is dependent on intensity and will not work with non-trivial illumination-dependent error covariances. To transform (9) into an intensity invariant measure we project the vectors $\underline{x}$, $\underline{y}$ and $\underline{a}$ onto the unit sphere using normalized *rgb*. Then we obtain the following distance measure:

$$\Phi_C(\bar{\underline{x}},\bar{\underline{y}}) = (\bar{\underline{x}}-\bar{\underline{a}})^T \bar{R}_x^{-1} (\bar{\underline{x}}-\bar{\underline{a}}) + (\bar{\underline{y}}-\bar{\underline{a}})^T \bar{R}_y^{-1} (\bar{\underline{y}}-\bar{\underline{a}}) \tag{10}$$

Distance measure (10) is a semi-metric since it violates the triangle inequality due to the pixel projection on the unit sphere; however, this is not of major concern since we are only concerned with relative distances and not absolute one. This distance measure is very easy to calculate, has no trouble with dark pixels, is intensity invariant and will work with non-trivial illumination-dependent error covariances.

$\bar{\underline{a}}$ is not considered to be a region prototype that is commonly accepted in the literature [2]. Rather, the idea is to find an $\bar{\underline{a}}$ which best explains both $\bar{\underline{x}}$ and $\bar{\underline{y}}$ and to use this intermediary quantity as a means of assessing the distance between $\bar{\underline{x}}$ and $\bar{\underline{y}}$. As mentioned at the beginning of this section, there are two different ways to determine $\bar{\underline{a}}$. First, we could determine the $\bar{\underline{a}}$ which maximizes the joint conditional probability $p(\bar{\underline{x}},\bar{\underline{y}}|\bar{\underline{a}})$. Other options exist as well. In essence, we want to find the optimum $\bar{\underline{a}}$ on the unit sphere.

4.2 Finding the minimum mean

To keep the analysis tractable we will assume that $R = \sigma^2 I$ and for the intensity invariant case we have $\bar{R}_x = \bar{\sigma}_x^2 I$ and $\bar{R}_y = \bar{\sigma}_y^2 I$. In order to minimize (10), we perform a component-based differentiation with respect to $\bar{a}_j$ where j represents the jth component of the vector, and set these partial derivatives to 0, i.e., $\frac{\partial}{\partial \bar{a}_j} p(\bar{\underline{x}},\bar{\underline{y}}|\bar{\underline{a}}_j) = 0$. For each component $\bar{a}_j$ the we have

$$-2\bar{\sigma}_x^{-2}(\bar{x}_j - \bar{a}_j) - 2\bar{\sigma}_y^{-2}(\bar{y}_j - \bar{a}_j) = 0 \tag{11}$$

Rearranging terms we obtain

$$\bar{a}_j = \frac{\frac{\bar{x}_j}{\bar{\sigma}_x^2} + \frac{\bar{y}_j}{\bar{\sigma}_y^2}}{\bar{\sigma}_x^{-2} + \bar{\sigma}_y^{-2}} \tag{12}$$

which simplifies to

$$\bar{a}_j = \frac{\bar{x}_j\bar{\sigma}_y^2 + \bar{y}_j\bar{\sigma}_x^2}{\bar{\sigma}_x^2 + \bar{\sigma}_y^2} \tag{13}$$

or in vector notation

$$\underline{\bar{a}} = (\underline{\bar{x}}^T\bar{R}_y + \underline{\bar{y}}^T\bar{R}_x)(\bar{R}_x + \bar{R}_y)^{-1} \tag{14}$$

Note that if $|\underline{\bar{x}}| = 0$ or $|\underline{\bar{y}}| = 0$, $\underline{\bar{a}}$ will correspond to the null set which is the desired behavior.

Interestingly. it is easy to show that substituting (14) into (10) will yield

$$\Phi_S(\underline{\bar{x}}, \underline{\bar{y}}) = (\underline{\bar{x}} - \underline{\bar{y}})^T(\bar{R}_x + \bar{R}_y)^{-1}(\underline{\bar{x}} - \underline{\bar{y}}) \tag{15}$$

which corresponds to a derivation under hypothesis (3) in normalized *rgb*. Therefore, we no longer need to compute means $\underline{\bar{a}}$ for this probabilistic distance reducing the computational complexity of each distance computation. Note that in this formulation, R_x now varies with the magnitude of $\underline{x}$.

Note that (15) is very similar in structure to the Mahalanobis distance [2] with one crucial difference: the Mahalanobis distance assumes that $\underline{\bar{x}}$ and $\underline{\bar{y}}$ have the same noise distribution whereas we do not make that assumption. Also, note that we did not pick the Mahalanobis distance to derive the new semi-metric; instead, we started with basic principles from which a generalized Mahalanobis distance emerged.

Figure 4 shows the distance computations using the new distance measure. The distances between dark pixels are now a bit lower than distances between dark pixels and other colors with higher intensities (e.g., compare distances between "dark red" and "dark green" and between "medium red" and "dark green" in Figures 3 and 4). Furthermore, the distances between colors with high intensity values are large as they were before. This analysis suggests that low intensity pixels will most likely merge with pixels of high intensity. This is desirable in areas where a shadow falls upon an object which results in some parts of the object being very dark. However, other areas which are adjacent to this dark region might become merged.

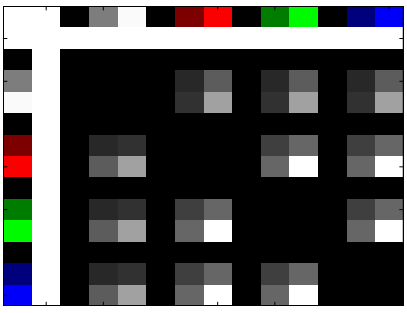

Figure 4: Probabilistic shading invariant distances between various colors in RGB. Notice that distances between pixels with low intensity values are now very small compared to distances between higher intensity pixels while remaining "0" for pixels of exactly the same color as compared with Figure 3.

5. RESULTS

We perform the segmentation using a Markov Random Field-based hierarchical image segmentation procedure with the hierarchical edge-based Potts model first detailed in [13]. Details of the exact pixel grouping algorithm are given in [14]. In this case, we use Iterated Conditional Modes as a greedy optimization method [1]. The number of labels used will be $K = 10$ and does not represent the number of regions. We will assume that the variance of the noise

Figure 5: Original *Toys* image (size: 256×256) and *Peppers* image (size: 512×512).

σ_i^2 for each color band $i \in \{R, G, B\}$ is the same and is given by the user. Original images are shown in Figure 5. For the *Toys*[1] image, we assumed $\sigma_i^2 = 4^2$ while for the *Peppers* image we assumed a $\sigma_i^2 = 8^2$.

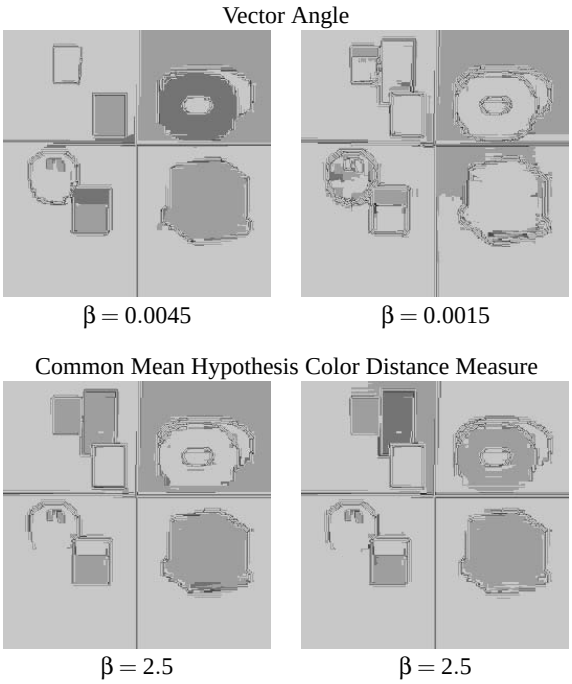

Figure 6: Results on the *Toys* image using the Common Mean Hypothesis color distance measure in *RGB*. The Potts models used were $\beta = 0.0045$ and $\beta = 0.0015$. A higher β is needed to produce less noisy vector angle-based results at the cost of one region merging into the background. Some regions of differing color are merged together with the Common Mean Hypothesis Test Distance Measure usually through the intermediary of a dark region.

Figure 6 shows the results for the *Toys* image [3]. Note that vector angle results for a lower β are generally worse especially around highlight areas whereas for a higher β the resulting image appears much cleaner, however, one object (in the upper left panel) has been completely absorbed by the background from the image. Regions of very low intensity help to separate objects from the background (e.g., the ball in the lower left panel). Regions of constant color are grouped together and separated from others. One very dark shadow in the lower left panel causes the ball object to merge with the background color (expected result). Many edge pixels are grouped sep-

[1] *Toys* image [4] publicly available at the following website `http://www.science.uva.nl/research/isla/themes/FeaturesAndColor.php`.

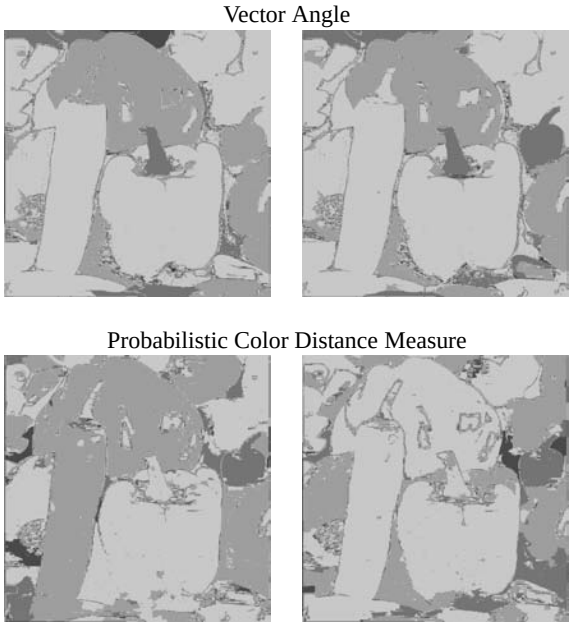

Figure 7: Results on the *Peppers* image using the probabilistic color distance measure in *RGB*. We used $\beta = 0.015$ for vector angle, and $\beta = 0.65$ for the probabilistic distance measure. Note that dark regions appear as individual regions or are merged with similarly colored lighter regions. Some regions of differing color are merged together usually through the intermediary of a dark region. Two different segmentation runs are shown.

arately from the objects due to jpeg compression artifacts in the image. Furthermore, the proliferation of small regions along the edges causes large regions such as the background and the sphere to have only a few pixels of common border which increases the likelihood of region spilling.

Figure 7 shows an example of the segmentation of the *Peppers* image. As opposed to the vector angle results where areas of dark pixels contained a proliferation of small regions, the results for the probabilistic distance measure show that pixels in dark regions can be grouped together or grouped with an adjoining region with a well defined color. However, regions of different color can still merge together if they are connected by a darker region. Also, note that highlight areas are detected as separate regions since this semi-metric is not highlight invariant.

6. CONCLUSIONS

In this paper, we have demonstrated an intensity invariant and noise resistant distance measure. The distance measure has been derived from first principles and is well grounded in statistics. The distance measure is a metric in Euclidean spaces and a semi-metric in projected spaces like normalized *rgb*. It allows the extension of vector angle and the Euclidean distance in normalized *rgb* to noise resistance. Its effectiveness has been validated by image segmentation results on several images. However, some segmentation results suffer from region-to-region spilling which could potentially be avoided by using a local region prototype-based segmentation model where between region differences are calculated based on region means and not just edge differences [14]. Furthermore, a comprehensive model which also includes specular reflection has been devised and will be presented in the near future.

7. ACKNOWLEDGEMENT

This work was done when the first author was a Ph.D. student at the University of Waterloo. Partial funding for this project from NSERC and OGS is hereby acknowledged.

REFERENCES

[1] J. Besag, "On the statistical analysis of dirty pictures," *J. Roy. Statist. Soc. B*, vol. 48, pp. 259Ű-302, 1986.

[2] R. O. Duda, P. E. Hart, and D. G. Stork, *Pattern Classification*, 2nd ed., New York, USA: John Wiley & Sons, 2001.

[3] T. Gevers and A.W.M. Smeulders, "Color-based object recognition," *Pattern Recognition*, vol. 32, pp. 453–464, 1999.

[4] T. Gevers, "Adaptive Image Segmentation by Combining Photometric Invariant Region and Edge Information," *IEEE Transactions on Pattern Analysis and Machine Intelligence*, vol. 24, no. 6, June 2002, pp. 848-852.

[5] G. Healey, and T. O. Binford, "A color metric for computer vision," *IEEE CVPR Conf.*, p. 10–17, Ann Arbor, 1988.

[6] G. E. Healey and R. Kondepudy, "Radiometric CCD camera calibration and noise estimation," *IEEE Trans. on Pattern Analysis and Machine Intelligence*, vol. 16, no. 3, pp. 267Ű-276, 1994.

[7] L. Shafarenko, M. Petrou, and J. Kittler, "Automatic watershed segmentation of randomly textured color images," *IEEE Trans. on Image Processing*, v. 6, p. 1530–1544, November 1997.

[8] S. A. Shafer, "Using Color to Separate Reflection Components," *Color Res. and Ap.*, v. 10, n. 4, p. 210–218, 1985.

[9] K.-K. Sung, "A Vector Signal Processing Approach to Color", Technical report AITR-1349, AI Lab, MIT, Jan 1992.

[10] S. Tominaga, "Surface Identification Using the Dichromatic Reflection Model," *IEEE Trans. Pattern Analysis and Machine Intelligence*, Vol. 13, No. 7, pp. 658–670, July 1991.

[11] S. Tominaga, "Dichromatic Reflection Models for a Variety of Materials," *Color Research and Application*, Vol. 19, No. 4, pp. 277–285, 1994.

[12] Y. Tsin, V. Ramesh, and T. Kanade, "Statistical calibration of CCD imaging process," *IEEE IntŠl Conf. Computer Vision*, pp. 480Ű487, 2001.

[13] S. Wesolkowski and P. Fieguth, "Hierarchical Region-Based Gibbs Random Field Image Segmentation," *International Conference on Image Analysis and Recognition*, Porto, Portugal, Sept 2004.

[14] S. Wesolkowski, *Stochastic Nested Aggregation for Images and Random Fields*, Ph.D. Thesis, University of Waterloo, Waterloo, Canada, 2007.