Improving Perceptual Uniformity of Sampling in Color Look-up Tables

Rohit A. Patil and Maxim W. Derhak; ONYX Graphics Inc; Salt Lake City, Utah, USA

Abstract

Color look-up tables (CLUTs) are frequently used to embed pre-computed transformations between a colorimetric space (e.g. CIELAB) and a device space (e.g. CMYK). The sampling method used to populate a CLUT has a significant impact on its accuracy and its error distribution. This paper discusses sampling methods of CLUTs which convert CIELAB space to device space. Based on the discussion, a new sampling method to achieve better perceptual uniformity of sampling between grid points is proposed. The method uses CIEDE2000 color difference equation to sample the grid points along a and b* dimensions in a way that the distance between the grid points along these dimensions is more perceptually equal. Using such a strategy improves the interpolation accuracy of neutral colors while keeping the CLUT size unchanged.*

Introduction

Color look-up tables (CLUTs) are frequently used to transform data from one space to a different space. The focus of this paper is on CLUTs used to convert CIELAB colors to device space (e.g. CMYK) often referred to as an inverse CLUT. Such CLUTs are commonly used in ICC profiles [1] for output devices such as printers. The computation of CLUT entries involves a set of transforms such as device models and gamut mapping. These transforms are rarely directly used to convert images in real time due the complex nature of the transforms and the amount of time it may take to convert large numbers of pixels [2]. Assuming that such transforms capture the device behavior accurately, the onus of conveying this information with the same accuracy then falls on the methods used to create the CLUT.

The creation of the CLUT and then its subsequent usage to transform colors can be divided into three parts as described by Kang [3]. The first part is the ‘sampling’ used to create the CLUT and is the topic of further discussion of this paper. The other two are ‘look-up’ and ‘interpolation’ parts which are the operations that convert a particular pixel to the destination space using the CLUT. The later two are not the focus of this paper and interested readers are referred to books by Sharma [2] and Kang [3] for detailed discussion of these topics.

Sampling of a CLUT can be defined as the number of grid points in the CLUT and the spacing between these grid points. A simple sampling method commonly used is to divide the source color space into equal number of steps in each dimension with equal spacing between the grid points (referred to as “uniform sampling method” henceforth). Using this method the accuracy of the CLUT depends directly on the number of grid points. The more the number of grid points the greater will be the accuracy. This will also decrease the error amplitude [3] - defined as the largest of all the errors calculated between the grid points. But, there are some drawbacks to using a uniform sampling strategy. As the number of grid points increases the storage size increases as well as increasing the time to populate the CLUT. Using the uniform sampling method for creating an inverse CLUT has another caveat that will be discussed next.

Perceptual non-uniformity of uniform sampling method in CIELAB space

When equal spacing is used between grid points of a CLUT with its source space as a colorimetric space, the perceptual differences between neighboring grid points should ideally be equal. When the source space used is CIELAB, the preceding statement would be true if all Euclidean distances in CIELAB (calculated using the ΔE_{ab} color difference equation [4]) correlated well with perceptual differences throughout the CIELAB space. As documented in prior color literature [4][5][6], this correlation is much less than perfect. It has been noted that perceptually equal color differences have smaller Euclidean distances for neutral colors compared to the Euclidean distances for high chroma colors [5] in the CIELAB space. This resulted in CIE recommending an optimized color difference equation referred to as the CIEDE2000 color difference formula; that can be used to more accurately predict perceptual color differences using the CIELAB space [5].

Figure 1 illustrates the distribution of maximum CIEDE2000 color differences in a 33x33 a*-b* plane at constant L* of 50. For the CIELAB address of each grid point in the plane, the CIEDE2000 color differences compared against the CIELAB grid point addresses of its eight neighbors were calculated and the maximum of these was stored at that point. These maximum differences were then plotted as a surface plot shown in Figure 1. The color blue indicates the lowest CIEDE2000 value of about 3.0 and the color brown indicates the highest CIEDE2000 value of about 12.9.

As can be seen in Figure 1, the distribution of color differences is not uniform across the CIELAB grid constructed using the uniform sampling method. The differences are higher for neutral colors while they are lower for high chroma colors.

In the CIEDE2000 formula [5], the lightness term is independent of the chroma and hue terms. This means that the color difference maps (as shown in Figure 1) of a*-b* planes at different lightness levels are exactly the same. Ideally if there was complete perceptual uniformity of differences between grid points there would be only one color map with a fixed CIEDE2000 color difference between all the grid points. As this is not the case, it can be concluded that a uniform sampling of the CIELAB color space does not equate to perceptual uniformity of differences between the grid points as calculated using the CIEDE2000 color difference equation.

Having perceptual uniformity of grid point address sampling of a CLUT is desirable to reduce possible perceptual interpolation errors for neutral colors. As can be seen from Figure 1, when uniform sampling is used, the perceptual differences of addresses in neutral colors are larger than for high chroma colors indicating that the sampling for the neutral colors is perceptually farther apart than for the high chroma colors. Bigger perceptual distances between addresses of the grid points for neutral colors could potentially result in reduced accuracy for colors that lie in between these grid points when they are used for interpolation. It is hypothesized that, improving perceptual uniformity of the grid point addressing would result in more grid points in the neutral region thereby better capturing potential non-linear behavior and improving the accuracy of these colors.

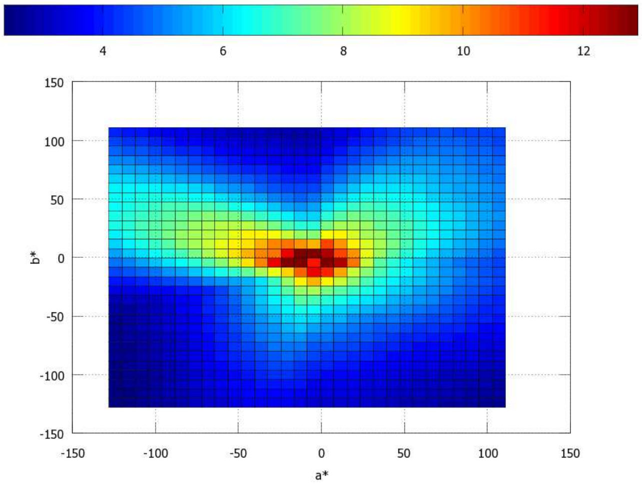

***Figure 1**. Maximum CIEDE2000 color difference distribution across the a*-b* plane with a granularity of 33x33 at constant L* of 50*

Improving perceptual uniformity

Since we are using the CIEDE2000 color difference equation to gauge the perceptual uniformity of sampling of a CLUT, we propose a method which makes use of this equation to come up with a sampling method that has improved perceptual uniformity over the uniform sampling method discussed above.

Urban et al [6] previously presented a method to convert CIELAB space into a perceptually uniform color space in which all Euclidean distances correlate better with perceptual differences as defined by color differencing equations. Such a color space is often referred to as a Euclidean color space. They used the CIEDE2000 formula to computationally determine a transform to do the conversion from CIELAB to a nearly Euclidean space. Since the lightness term is independent of a* and b* terms in the CIEDE2000 equation, the transform proposed by Urban et al can be implemented as a one dimensional LUT for converting lightness and a two dimensional LUT for converting a* and b*. These LUTs were constructed using a computational technique and were then used to convert CIELAB space to a Euclidean color space. Ideally one would want to encode such a Euclidean space directly into an inverse CLUT. An inverse CLUT encoded using this color space would have much improved perceptual uniformity even when using a uniform sampling method.

The aim of the current paper is to improve the perceptual uniformity of a CLUT encoded using the CIELAB space. This constraint is imposed due to the fact that ICC profiles only allow use of either CIEXYZ or CIELAB as connection spaces [1] and the authors' intention is to be able to use these CLUTs in the inverse tags of ICC profiles. Some of the other properties of the inverse tags of ICC profiles are listed below.

- Sampling between the grid points of a CLUT must be equal.
- The number of grid points for all dimensions of a CLUT must be the same if building a version 2 ICC profile. This restriction was removed in the version 4 specification.
- One dimensional LUTs can be used prior to the CLUT to preprocess incoming data for each dimension independently before passing it on to the CLUT as well as after the CLUT as shown in Figure 2.

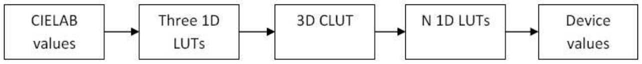

***Figure 2**. Order of transforms in an inverse tag of a typical ICC profile*

Figure 2 illustrates a typical embodiment of the order of operations in an inverse tag of an ICC profile [1]. The transform step immediately before the 3D CLUT provides a means for modifying the addressing of values in the CLUT with the first curve being applied to L* values, the second curve being applied to a* values, and the third curve being applied to b* values. If the architecture of an ICC profile allowed for a transform that included a single 1D LUT for L* and a 2D LUT for the a*-b* plane prior to the CLUT, then the method proposed by Urban et al [6] could directly be encoded into the profile to improve perceptual uniformity.

Recently, the ICC approved a new tag type called the multiProcessElementsType [7] which allows for flexible ordering of a sequence of different transform elements in a tag along with the ability to encode data using floating points. Using this type of tag it is possible to directly encode a version of the method proposed by Urban et al [6] as shown in Figure 3.

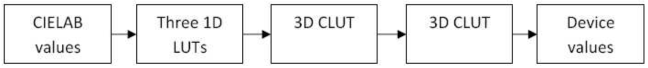

***Figure 3**. Order of transforms in an inverse multiProcessElementsType tag of an ICC profile encoding the method proposed by Urban et al [6]*

Within a multiProcessElementsType tag two processing elements can be placed before the 3D LUT associated with device value generation. The first element contains three curve elements with the first curve performing the L* perceptual equalization and the other two curves providing an identity transform on a* and b* values. Ideally the first 3D CLUT element would only provide 2-dimensional highly sampled interpolation of a* and b* to get a perceptually uniform translation of a* and b* values. However, in a CLUT in a multiProcessElementsType tag, all dimensions participate in the interpolation. An encoding of the Urban et al's method can be accomplished by using a 3D CLUT with 2xMxM grid points where M is the sampling of the a*-b* plane, and identical values for a* and b* entries are placed in both L* planes (resulting in only a doubling of the encoding overhead). In this manner the value of L* going into the first 3D CLUT does not affect the values of a* and b* coming out. With these two transforms in place, the CIELAB values going into the second CLUT will then be nearly perceptually uniform. The second 3D CLUT can then use a uniform sampling method with improved perceptual uniformity. One difficulty with this approach is that tags of the multiProcessElementsType are optional for version 4.3 profiles and they cannot be used in version 2 profiles. Hence a further simplification of this method is proposed that can more readily be applied to the transform sequence shown in Figure 2.

This simplified method makes use of 1D LUTs before the 3D CLUT (shown in Figure 2) to readdress the grid points of the succeeding CLUT. As can be seen from Figure 2, for a CLUT with CIELAB as the source space, we have three 1D LUTs before it. The first 1D LUT is for the L* dimension, the second is for the a* dimension and the third is for the b* dimension. A linear 1D LUT is used for the L* dimension, so the grid points in the CLUT are not readdressed along the L* dimension. The L* 1D LUT could also be constructed using the method proposed by Urban, but was not used in this case. We use the CIEDE2000 equation to construct 1D LUTs for the a* and b* dimensions such that the grid points in the CLUT along these dimensions are readdressed and are perceptually uniform.

Proposed method details

The aim of the proposed method is to construct 1D LUTs for a* and b* such that the CIEDE2000 color differences between the neighboring points is closer to being equal. However, since only one dimension is being modified at a time the approximation to achieving perceptual uniformity will not be as complete as that proposed using the approach shown in Figure 3.

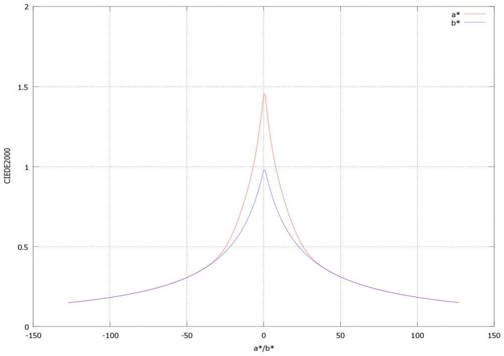

***Figure 4**. CIEDE2000 color differences between neighboring points of a linear a*/b* 1D LUT*

Figure 4 shows the CIEDE2000 color difference between neighboring points of linear a* and b* 1D LUTs. The 1D LUTs for both a* and b* range from -128 to 127 in one unit steps. This is also the range of values allowed for a* and b* in an ICC profile. The CIEDE2000 values for the a* curve were calculated between the neighboring points at a constant b* of 0 and L* of 50. Similarly the values for the b* curve were calculated between the neighboring points at a constant a* of 0 and L* of 50. Figure 4 could be interpreted as independent one dimensional representations of Figure 1 along the a* and b* axes.

Since the Euclidean distance between the neighboring points on these curves is equal, the desire is to have the CIEDE2000 color difference between these points to also be equal. But as can be clearly seen from Figure 4, this is not the case. This desired behavior can be achieved by using a computational technique described next.

Let us denote the a* 1D LUT by (X_{1-n}, Y_{1-n}) pairs, where X values are the input values and Y values are the corresponding output values and n is the number of points in the 1D LUT. The first step is to select a range of X values such that they extend a little beyond -128 and 127. Set the lowest value as X_1 and the same value as Y_1. Then search for a value which has a CIEDE2000 color difference of one unit from the previous value Y_1 and set this value as Y_2. The color difference is calculated by setting the first L*a*b* as ($50,Y_1,0$) and the second L*a*b* as ($50,Y_2,0$) . The value X_2 for this point is X_1 plus a Euclidean distance of one unit. Repeat the search until we have populated all the (X_{1-n}, Y_{1-n}) pairs. Thus we have a 1D LUT wherein the X values have a Euclidean distance of one unit between neighboring points and the corresponding Y values have a CIEDE2000 color difference of about one unit between neighboring points. Since the X values extend beyond -128 and 127, the Y values might not necessarily range between -128 and 127. First we find the Y values for a range of X values between -128 and 127 in increments of one Euclidean distance, by doing a simple one dimensional linear interpolation on the 1D LUT created in the first step. Now we have X values ranging between -128 and 127, but the Y values could still possibly be beyond the desired range. Since the Y values are perceptually uniformly spaced as gauged by the CIEDE2000 color difference equation, we can simply rescale these Y values to range between -128 and 127. After doing this, the color difference between the neighboring Y values may not necessarily be one unit anymore, but they will all be equal. The b* 1D LUT can be constructed similarly. Figure 5 shows the a* and b* 1D LUTs constructed using the method described above.

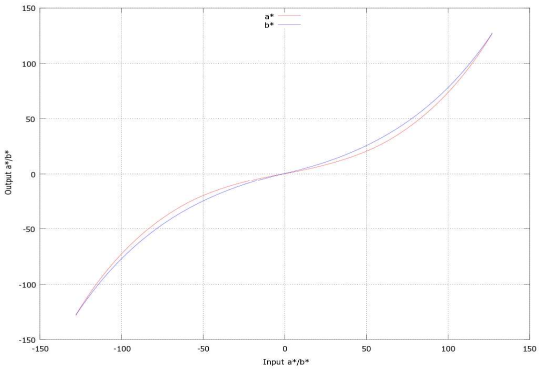

***Figure 5**. 1D LUTs for the a* and b* dimensions created using the new method*

Results and Analysis

As can be seen from Figure 5, the output values are compressed near the neutrals and expanded near the high chroma colors. This would suggest that more a* and b* values are being mapped near the neutral region compared to using a linear 1D LUT. This can also be verified by analyzing the histogram plot of grid point sampling as shown in Figure 6.

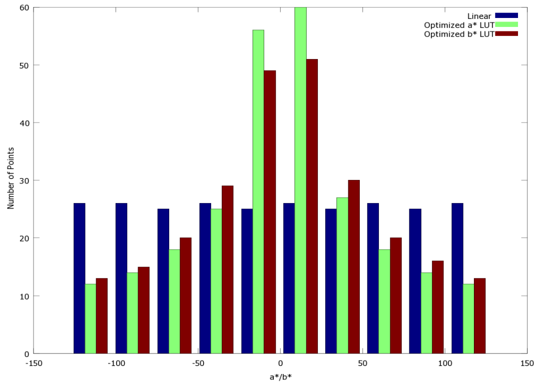

***Figure 6**. Histogram illustrating distribution of 1D LUT output values*

Figure 6 is a plot of the number of output values in the 1D LUTs using two different methods, linear and optimized. As expected, the "Linear" 1D LUT distributes the output values evenly across the output range. Using the optimized 1D LUTs, it can be seen that there are more output values near the neutrals (a* and b* near zero) than the number of output values in the high chroma region (large positive or negative values of a* and b*). Also, the shape of the distribution of the output values from the optimized LUTs is very similar to the shape of the corresponding curves seen in Figure 4. These optimized LUTs can therefore be used to readdress the grid points of a succeeding CLUT ensuring that the distances between grid points are perceptually more uniform.

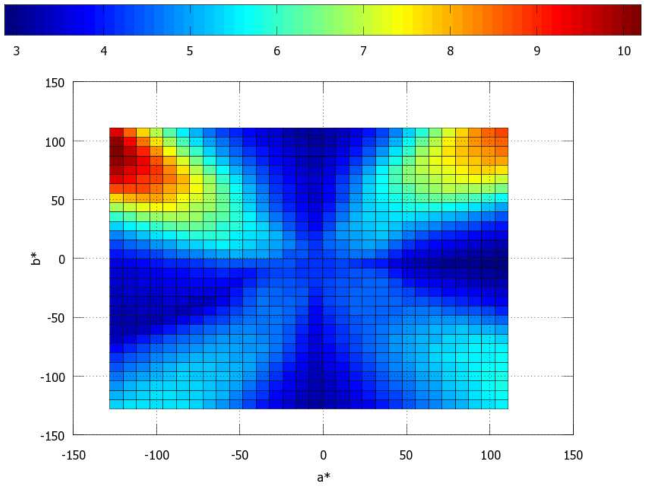

***Figure 7**. Maximum CIEDE2000 color difference distribution across the resampled a*-b* plane with a granularity of 33x33 at constant L* of 50*

A surface plot similar to the one seen in Figure 1 is plotted in Figure 7 where the grid points in Figure 7 are readdressed using the optimized LUTs shown in Figure 5. Each CIELAB address for grid point in a 33x33 a*-b* plane was first transformed through the optimized 1D LUTs, and then for each transformed grid point address, the CIEDE2000 color differences against its eight neighbors were calculated and the maximum of these was stored at that point. These maximum differences were then plotted as a surface plot shown in Figure 7. The color blue indicates the lowest CIEDE2000 difference of about 2.8 while the color brown indicates the highest CIEDE2000 difference of about 10.0.

As can be clearly seen in Figure 7, the maximum color differences between grid point addresses are now found in the high chroma regions and the lowest color differences are near the neutrals. Typical device gamuts will be centered around the neutrals and these maximum perceptual differences will therefore be typically less of an issue. Moreover the higher CIEDE2000 differences are in the very high chroma regions which are rarely able to be reproduced by an output device. Additionally, the color differences along both the a* and b* axes are more or less equal to the neutrals and are the lowest.

However, this approach is only an approximation. In the CIEDE2000 equation, the a* and b* dimensions are not independent of each other whereas the optimized a* LUT that was built assumes a b* value of zero and vice-versa for the b* LUT. Hence, the optimized a* and b* 1D LUTs have the dependency only along the axes giving the low color differences. The color differences will be higher as we diagonally move away from both the axes to the corners. This is clearly seen in Figure 7. These higher perceptual differences in the high chroma region should be acceptable due to the fact that the differences near the neutrals are now much less. Also, these higher differences are actually lower than the high differences seen in the neutral region (from Figure 1) using the uniform sampling method.

To summarize, the perceptual distances of grid point addresses as measured by CIEDE2000 color differences across most of the plane (except the corners of the high chroma green and orange regions) are mostly uniform. And it can be concluded that the grid points of a uniformly spaced CLUT applied afterwards are more perceptually uniformly spaced as compared to the uniform sampling method, which is the aim of the technique. One important point to note is that this improved perceptual uniformity is achieved without increasing the number of grid points in the CLUT.

Since, the optimized 1D LUTs can be embedded inside the inverse tag of an ICC profile along with the CLUT; the performance of such ICC profiles can be compared against the ones without these optimized 1D LUTs. Two sets of CMYK printer ICC profiles were built. The first set of profiles, referred to as "Profile Linear", were built using linear 1D LUTs and granularities of 27, 33 and 41 for the CLUT. The second set of profiles, referred to as "Profile Optimized", were built using a linear 1D LUT for L* while the optimized 1D LUTs for a* and b* and the granularities of 27, 33 and 41 for the CLUT. The printer device used to build these profiles was a forward CLUT, representing a CMYK to CIELAB transform of an actual printer. Using a CLUT as a printer device rather than an actual printer has several benefits in evaluating the method proposed. First, no actual colors would have to be printed and measured making sure that any printer and measurement device variability will not contribute to the evaluation results. Second, we could evaluate a large number of color samples which could have been prohibitively time consuming to print and measure. An even sampling of CIELAB colors inside the gamut of the printer device formed the test dataset for a total of 12439 colors. These colors were then converted to CMYK through each profile's inverse tag using the absolute rendering intent. The CMYK values were then converted to CIELAB using the printer device CLUT. The CIEDE2000 color differences were then calculated between the original CIELAB values and the converted ones. Table 1 lists some vital statistics of these color differences.

The statistics in Table 1 clearly show that there is an improvement in the overall CIEDE2000 errors when using the optimized profile. For a given granularity, there is a decrease in the mean and maximum CIEDE2000 errors when using the optimized profile. The total percent of colors with a CIEDE2000 difference less than 1 also increased (improved) for the optimized profile. Moreover, a "Profile Linear" even with a granularity of 41 grid points is not able to match the accuracy of a "Profile Optimized" with a reduced granularity of 27. Thus, the proposed method could actually be used to build a CLUT with reduced number of grid points while maintaining or exceeding the accuracy of a CLUT with more grid points using a uniform sampling. This would reduce the time required to build the CLUT while also reducing storage cost of such CLUTs. These results also clearly indicate that the higher perceptual sampling distances shift to the high chroma region when using the optimized 1D LUTs. Nevertheless, this does not cause an overall increase in the CIEDE2000 color difference errors, but in fact helped reduce errors over the entire color space.

Since the optimized profile has more grid points near the neutrals a closer look at the performance of these profiles near the neutral colors was investigated. Table 2 lists the same statistics as in Table 1 but only for a subset of colors near the neutrals. The colors with original chroma values less than 10 were considered as near neutrals and there were 1256 such colors in the test dataset.

As can be seen from Table 2, the improvement in CIEDE2000 errors near neutrals is similar to the overall improvements. The mean color difference for each optimized profile is almost half of the mean color difference of each corresponding profile using uniform sampling method with corresponding granularities. It can also be seen from Table 2 that the mean difference for an optimized profile using a granularity of 27 is significantly less than the linear profile built using a granularity of 41. Thus again emphasizing that the proposed method could be used to build profiles with lower number of grid points but with increased accuracy. The number of colors with a CIEDE2000 color difference less than 1 also improves significantly when using the optimized profile indicating a reduction in interpolation errors for these colors. Since a CIEDE2000 value of greater than 1 is considered as noticeable difference to a human observer, a grayscale image printed through the optimized profile should look a lot more accurate than one printed through a linear profile.

It should be noted that these are only example numbers and they can differ considerably depending on the behavior of the device. If the device behavior is highly linear near the neutrals then increasing the number of grid points in this region might not be as beneficial. On the other hand, if the device behavior is significantly non-linear for near neutral colors then using the optimized profile can improve the performance of the profile for near neutral colors.

Conclusions

The perceptual non-uniformity of the uniform sampling method when constructing a CLUT was demonstrated. Urban et al [6] previously proposed a method to convert the CIELAB space into a perceptually uniform Euclidean space which was discussed as a means of achieving perceptually uniform grid point addressing. A proposal for encoding Urban's method [6] using multiProcessElementsType tags was discussed. However, this method cannot generally be used to directly encode a CLUT into more typical ICC profiles. We then presented a simplification of the method using one dimensional adjustments to approximate perceptual uniformity with a computational technique to construct 1D LUTs for the a* and b* dimensions. These 1D LUTs can more readily be used to readdress the grid points of CLUTs in typical ICC output profiles. The CIEDE2000 equation was used to compute these optimized 1D LUTs. Computationally, it was seen that using these optimized 1D LUTs improved the perceptual uniformity of the CLUT (especially for near neutral colors). Finally it was demonstrated that using this technique in actual device output profiles improved the interpolation accuracy of a test dataset including a significant improvement in the accuracy of neutral colors. This was achieved without increasing the number of grid points in the CLUT. In fact, it was seen that to achieve the same level of accuracy for near neutral colors as that of optimized profile, the uniform sampling method would need a very large number of grid points. Thus, the method presented provides a means to maintain or improve the accuracy of an inverse CLUT while minimizing the cost of populating and storing such a CLUT.

Table 1. CIEDE2000 statistics for in gamut colors

CLUT Grid Points	27			33			41		
CIEDE2000	Mean	Max	< 1	Mean	Max	< 1	Mean	Max	< 1
Profile Linear	0.76	3.64	85.0%	0.71	3.45	89.0%	0.68	3.19	91.1%
Profile Optimized	0.48	3.04	94.9%	0.45	2.99	96.6%	0.43	3.08	97.2%

Table 2. CIEDE2000 statistics for near neutral in gamut colors

CLUT Grid Points	27			33			41		
CIEDE2000	Mean	Max	< 1	Mean	Max	< 1	Mean	Max	< 1
Profile Linear	1.00	3.29	69.3%	0.96	3.16	75.8%	0.94	2.95	78.7%
Profile Optimized	0.53	2.75	90.8%	0.50	2.82	91.9%	0.48	2.71	92.5%

References

[1] International Color Consortium, Specification ICC.1:2004-10 (Profile version 4.2.0.0) Image technology color management – Architecture, profile format and data structure (2004). http://color.org/ICC1v42_2006-05.pdf

[2] G. Sharma, Digital Color Imaging (CRC Press, 2003).

[3] H. R. Kang, Computational Color Technology (SPIE – The International Society for Optical Engineering, 2006).

[4] R. S. Berns, Billmeyer and Saltzman's Principles of Color Technology, 3rd Edition (John Wiley and Sons, 2000).

[5] CIE Publication No. 142, Improvement to Industrial Colour Difference Evaluation, CIE Central Bureau, Vienna (2001).

[6] P. Urban, R. S. Berns, M. R. Rosen, Constructing Euclidean Color Spaces based on Color Difference Formulas, Proc. 15th Color Imaging Conference, Albuquerque, New Mexico, pg. 77-81 (2007).

[7] International Color Consortium, Revision to the Specification ICC.1:2004-10, Floating Point Device Encoding Range (2006). http://color.org/ICCSpecRevision_02_11_06_Float.pdf

Author Biography

Rohit A. Patil received his B. Tech degree in Chemical Technology from the Mumbai University Institute of Chemical Technology, India, in 2002 and his Master's degree in Color Science from the Munsell Color Science Laboratory at Rochester Institute of Technology, in 2004. Currently he is working as a senior Color Science Engineer at ONYX Graphics, Inc. His interests lie in the area of color management and digital color imaging, especially researching improvements to device modeling, tone mapping and gamut mapping algorithms.