Accuracy-Preserving Smoothing Of Color Transformation LUTs

Ján Morovič, Angel Albarran, Jordi Arnabat, Yvan Richard and Martí Maria
Hewlett–Packard Company, Barcelona, Spain

Abstract

Color transformations are frequently embodied in look-up tables (LUTs) whose content is the result of computation performed on the basis of color measurements. Noise in color measurements and device output gives LUTs that result in lack of smoothness in color transitions. Smoothing of color LUTs is therefore typically performed, often with a significant loss of accuracy. This paper proposes a method to both smooth LUTs and preserve a significantly greater degree of colorimetric accuracy than conventional methods. The key to such a result is to smooth along 1D paths where lightness change is greatest. E.g., along constant a*b* lines where CIELAB is the indexing space or along paths parallel to the line connecting the black and white points where the indexing space is device–dependent.

Introduction

Color transformation look-up tables (LUTs) are the most wide-spread way of taking digital color data from one domain and transforming it into another. ICC profiles[1] use LUTs in predicting colorimetry from device color data and conversely to compute device color data that matches a given, desired color. Imaging devices use internal color LUTs to transform from standard color spaces (e.g. sRGB, Adobe RGB) to their device color spaces and from those to inputs to their colorant channels. The computation of the content of such LUTs in turn is the result of a series of components, including device characterization, color appearance modeling, color enhancement and gamut mapping, and is typically done on the basis of measuring the colors obtained using a device for a sampling of its inputs.

In terms of their purpose, color LUTs need to deliver all the desired color attributes of the result, which often include smoothness of transitions, naturalness, contrast, shadow and highlight detail and either colorimetric accuracy or the colorimetrically accurate rendering of the result of some form of color enhancement. Given this list, it is not surprising to see the degree of complexity that goes into computing color transformation LUTs.

The fly in the ointment is the fact that there is noise in the color measurements of device outputs that form the basis of computing color transformation LUTs. The noise comes both from variability in the device's output and the measuring instrument's operation. Consequently LUTs obtained from acting directly on measured data tend to result in smoothness artifacts when used for color transformation. Take for example the following case (Fig. 1), where a single-colorant device, whose output is smooth with respect to its digital inputs, has a set of inputs sent to it and measured. As a consequence of noise the result is not smooth. Inverting the relationship expressed by the measurements when attempting to generate a smooth transition using the device then results in non-smooth device inputs, which in turn give artifacts when output on the device. The solution to this problem, which is very well known to imaging engineers, is to smooth either the measured data directly or to smooth a LUT built on its basis. The drawback though is that this tends to reduce the overall accuracy with which the device is controlled.

The aim of this paper is to review three approaches to smoothing color transformation LUTs, including a novel one that exploits the human visual system's sensitivity to luminance changes being greater than to chrominance changes.[2] Results of evaluating the new technique and comparing it with a typical, current alternative will also be presented. Before proceeding, please, note that this paper is about smoothing color transformation LUTs rather than images that may have been processed using such LUTs. Smoothing LUTs does not result in blurred images, but only in the degree with which smooth transitions are rendered as a result of their use.
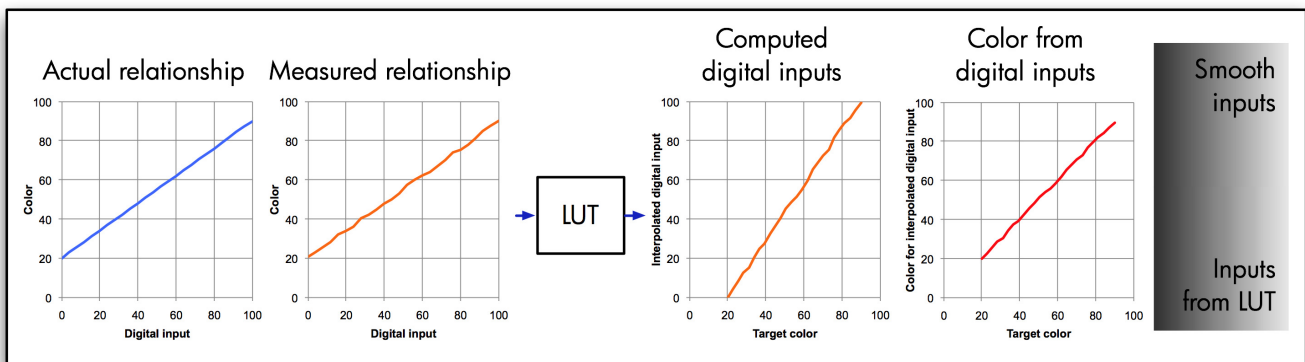

Fig. 1. Effect of noise on color transformations.

Alternative smoothing strategies

Three smoothing approaches will be presented here, with a focus on application to LUTs that are indexed in a colorimetric space and where the output is in a device color space (e.g., such as used in ICC profile BToA1 tags).

3D direct device data smoothing

The most typical approach used today is to take each LUT entry and replace its content with the weighted sum of that entry's $(2*n+1)^3$ neighborhood, centered on the entry being smoothed. The weights sum to one and can either be the same for all neighborhood entries or decrease as distance increases from the entry being smoothed.

This approach certainly compensates for noise in the data that a LUT is derived from and can give rise to smooth transitions. The problem though is that it can dramatically reduce the accuracy of the resulting transformation (e.g., introducing hue shifts). E.g., in the case of the test setup described later, the maximum error of an ICC profile jumps from 4.6 ΔE2000[3] when no smoothing is used to 6.4 ΔE2000 when this approach is used even with the smallest neighborhood (n=1). The challenge therefore is to find a new approach that delivers at least the degree of smoothing obtained from the 3D direct method without the accuracy penalty incurred here (Tab. 1).

Tab. 1. What should a new smoothing algorithm deliver?

No smoothing	Artifacts	Good accuracy
3D direct smoothing	No artifacts	Bad accuracy
New smoothing	No artifacts	Good accuracy

Modeled device data smoothing

The first idea here is to say that minimal changes due to smoothing can be achieved by estimating how smooth a LUT is to begin with and only making changes where non-smoothness is detected. Furthermore, the knowledge that smoothness is predominantly judged by the lightness changes in a transition (as opposed to chroma or hue changes) leads to considering lines of constant a*b*.

The analysis can be done by taking such a line from the BToA1 tag of an unsmoothed profile, looking at the device RGBs (dRGBs) in its nodes and using the profile's AToB1 tag to predict the lightnesses of the nodes. Note that the same principle applies to device CMYKs even though the remainder of this paper will focus on the RGB case. The predicted lightnesses are then judged in terms of their smoothness and changed to new, smoother values if needed. Changed lightnesses are then used as targets for a search in the AToB1 tag that looks for the dRGBs that most closely match the target lightness at the given a*b* values. Initially the hypothesis was formed that a smooth transition is one where the slopes of the piecewise linear function defined by the LUT vary monotonically (i.e. change keeps getting more or less steep across the entire trajectory).

Taking the black–to–white line in a pair of corresponding non–smoothed and 3D–direct–smoothed BToA1 LUTs, predicting the L*s of their content using the profiles' AToB1 tag and expressing them as differences from the straight line connecting the darkest and lightest in-gamut nodes' L*s (Fig. 2) destroys the above reasoning though. Transforming the non–smoothed data to CIELAB results in a path that is smoother than the predictions made for the 3D–direct–smoothed case. However, prints made using these profiles exhibit the opposite – i.e. 3D–direct–smoothing does indeed improve smoothness.
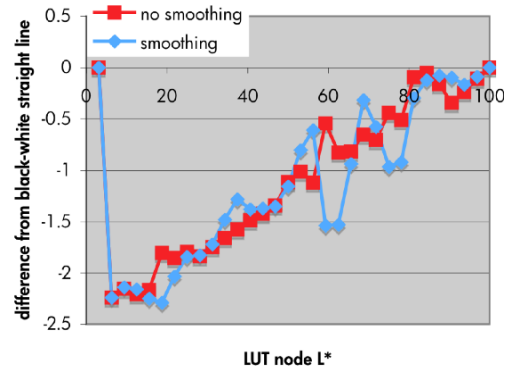

*Fig. 2. ΔL*s of non-smoothed (square) versus 3D-direct-smoothed (diamond) dRGBs of black–white, constant a*b* line.*

The reason for this is that the starting point is a profile that has good round-trip accuracy – i.e., in-gamut nodes contain dRGBs that match the nodes' colorimetry. E.g., if a node represents the L*a*b*=[50,0,0] location then it will contain dRGBs that map to that L*a*b* when passed through the AToB1 tag – and similarly for all the nodes of the lightness axis that are in gamut. In other words, for in-gamut relative colorimetric nodes in a profile with low round-trip errors, the profile will always think that its BToA1 output is smooth. This is so by definition. If a print made with such an unsmoothed profile is not smooth the reason for it is a mismatch between the profile's AToB1 tag and the printing system's behavior (e.g., due to noise in the data from which it was derived) and this has as a consequence that the AToB1 tag cannot be used for interpreting the dRGBs that the BToA1 tag outputs. Therefore the solution to smoothing the effects of noise cannot be along these AToB1 modeled lines and an alternative must be sought.

1D lightness-path smoothing

Since interpreting dRGBs using the AToB1 tag is a dead end, a new solution needs to come from focusing solely on dRGBs that are contained in a BToA1 tag and from changing them less than the 3D-direct method does while providing visual smoothing. Such a move, however, brings with it an inability to directly attempt a minimization of color change (since no reliable method of predicting it is available – the AToB1 tag having been disqualified as a model). Instead, dRGBs need to be altered in a way that can *a priori* be expected to result in lesser change and therefore a lesser reduction in accuracy.

Since it is lightness changes that are predominantly responsible for the perception of contours (i.e., edges where there should be none)

or other non-smoothness artifacts, the proposal here is to apply smoothing along such lines too. Instead of considering a 3D neighborhood of a LUT entry, here it will only be a $2*n+1$ 1D neighborhood from which LUT entries will be used for replacing a given LUT node's content. In the case of LUTs indexed in a colorimetric space this is straightforward and it will be neighbors along lines of constant a*b* that will be used. When indexing is done in a device color space (e.g., dRGB), the equivalent is to consider neighborhoods that involve the greatest lightness change. These in fact are such paths through the device color space that are parallel to the line connecting the black and white points.

The lightness–path (L–path) smoothing algorithm therefore replaces each LUT entry by a weighted sum of itself and the $2*n$ LUT nodes that are its neighbors along a lightness path. In the simplest case where n=1, nodes are affected only by two instead of 26 neighbors, which can be expected to result in less change than for 3D-direct smoothing (Fig. 3).
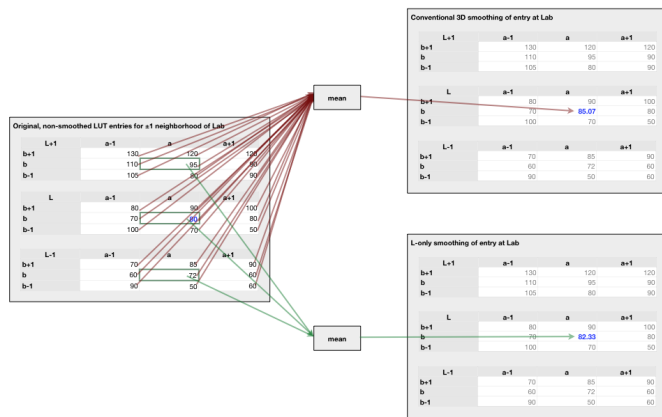

Fig. 3. Example of a LUT entry being smoothed based on its 3D neighborhood (top) versus a 1D, L–path neighborhood (bottom).

Note that this is not equivalent to just processing lightness or luminance channels, which is a well known technique[2] where, e.g., only the lightness values of a neighborhood would be taken into account. Instead we present here an adaptation of that idea by using lightness as the criterion for choosing an appropriate, small neighborhood and then processing either all colorimetric channels or all device color channels.

Evaluating L-path smoothing

To quantify the performance of L-path smoothing, both colorimetric and psychovisual tests were performed on a set of three ICC profiles using the following smoothing options: none, 3D and L-path. Prints for the test were made using a calibrated HP Designjet Z3100 printer on Hahnemühle Smooth Fine Art paper.

Psychovisual evaluation

Since smoothness is an attribute that cannot reliably be predicted computationally, a psychovisual category–judgment experiment[4] was conducted to quantify it. A composite test image (Fig. 4) was printed using the three smoothing approaches and shown to 11 observers with normal color vision. Each observer was asked to judge the smoothness of prints on a scale from one to seven, where one represented the worst smoothness they could imagine, seven represented the best one and categories are equally spaced.

The results of the experiment are shown in Fig. 5, where raw observer judgment data was analyzed using the mean category method.[4] L–path smoothing is judged to deliver a clear gain in smoothness while the traditional 3D smoothing only results in a moderate (and statistically insignificant) gain.
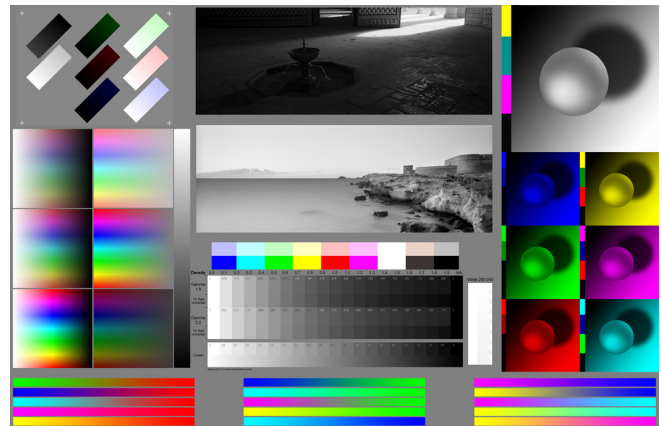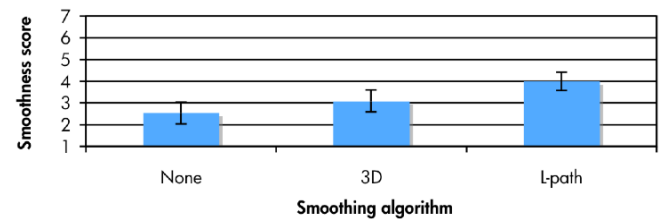

Fig. 4. Composite test image.

Fig. 5. Smoothness scores obtained in psychovisual experiment. Error bars indicate the 95% confidence interval.

Round-trip, accuracy and gamut

In addition to the primary effect of a smoothing algorithm on the resulting smoothness of color transitions, its effect on the colorimetric properties of the LUT is also important. Three attributes will therefore be computed for each of the three smoothing approaches: round-trip differences, accuracy and gamut coverage.
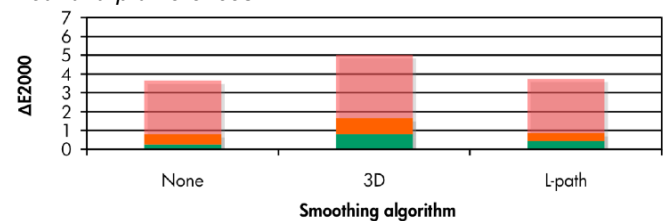
Round–trip differences

Fig. 6. Round–trip differences in ΔE2000 in terms of median (green), 95th percentile (orange) and maximum (transparent red).

Given an ICC profile's AToB1 and BToA1 tags, it is possible to compute the error introduced by taking in-gamut LABs, transforming them via the BToA1 tag to device color values and then using AToB1 to go back to LAB. The color differences between the initial and final LABs expresses the degree of symmetry between the two LUTs. The results for this test are shown in Fig. 6 and it can be seen that the loss of symmetry due to L–path smoothing is minimal compared to full 3D smoothing.

Colorimetric accuracy
Given in-gamut LABs, computing device color values for them using the BToA1 tag, then printing and measuring the result and finally computing $\Delta E2000$ color differences between the two LAB sets expresses the degree of accuracy to which desired LABs can be matched. It can be seen in Fig. 7 that L–path smoothing introduces a much lesser reduction in accuracy than 3D smoothing does.
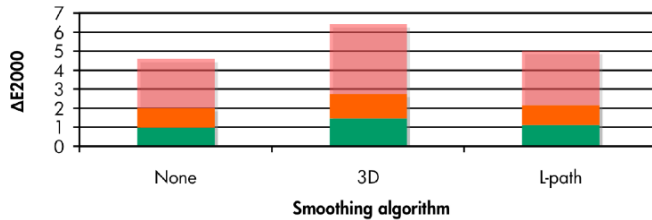

Fig. 7. Colorimetric accuracy in $\Delta E2000$ in terms of median (green), 95[th] percentile (orange) and maximum (transparent red).

Gamut coverage
How much of the device color space's (hyper)cube is used by an ICC profile's BToA1 tag determines how much of the device's color gamut is accessible when the profile is used. Since smoothing can affect this attribute of a profile, Tab. 2 shows the distance between 8–bit RGB values output for gamut boundary colors and the closest RGB cube coordinates. The greater these values are the further inside the RGB cube is the LUT output that ought to be on the RGB cube's surface and the less of the device's gamut is accessible. The results of the three smoothing alternatives are also compared to the performance of Bhachech *et al.*'s extrapolation scheme[5] that addresses this issue directly.

Tab. 2. 8–bit RGB Euclidean distance from cube surface.

	None	3D	L–path	Bhachech extrapolation
Median	0	0.47	0	0
95[th] percentile	0.14	6.07	2.18	0
Maximum	6.99	19.06	10.58	2.34

Conclusions

The L–path smoothing technique presented here has been shown to outperform typical 3D smoothing of color LUTs both in terms of the smoothness that it delivers and in terms of the colorimetric cost at which it does so. To reiterate this point, Fig. 8 shows the col-

orimetric accuracy results plotted against the smoothness scores for the three alternatives compared here and it can be seen that L–path smoothing represents a combination of attributes that outperforms both 3D smoothing and the use of no smoothing in color LUTs.
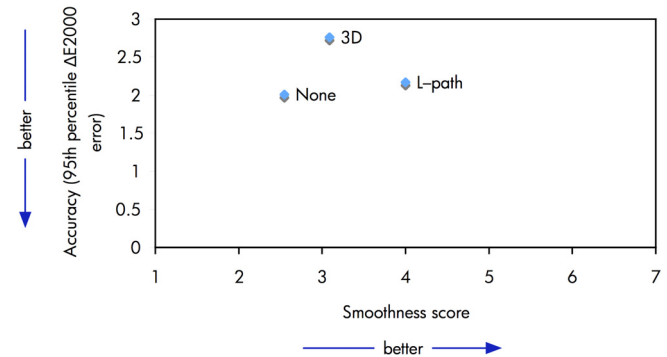

Fig. 8. Accuracy versus smoothness of alternative smoothing algorithms.

In essence L–path smoothing is another way of exploiting the greater sensitivity of the human visual system to lightness (luminance) changes and the greater frequency of spatial detail being predominantly present in the same domain.[2] These facts have long been known and used to engineering advantage and the present work simply represents the identification of a new area of application. What is arguably most novel here is that the concept of treating luminance differently from chrominance also applies when it is not directly these domains that are transformed. This is the case when smoothing BToA1 tags, where it is device color data that is being smoothed – albeit along paths where lightness alone changes in the indexing space.

Acknowledgments

References

1. ICC (2004) *Specification ICC.1:2004–10 (Profile version 4.2.0.0) Image technology colour management — Architecture, profile format, and data structure*, International Color Consortium, http://www.color.org/ICC1V42.pdf.
2. Hunt, R. W. G. (1995) Why is Black-and-White so Important in Color?, *4th IS&T/SID Color Imaging Conference*, Scottsdale, AZ, 54-57.
3. CIE (2001) *CIE 142:2001: Improvement to industrial colour difference evaluation*, CIE Central Bureau, Vienna.
4. Bartleson, C. J. (1984) *Optical Radiation Measurement. Vol. 5 – Visual Measurements*, Bartleson C. J. and Grum F. (eds.), Academic Press Inc.
5. Bhachech, M., Shaw, M. and DiCarlo, J. (2006) Improved Color Table Inversion Near The Gamut Boundary, *14th IS&T/SID Color Imaging Conference*, Scottsdale, AZ, 44–49.