# Color-matching Shots from Different Cameras Having Unknown Gamma or Logarithmic Encoding Curves 

Raquel Gil Rodríguez<br>Universitat Pompeu Fabra, Barcelona, Spain, raquel.gil@upf.edu<br>Javier Vazquez-Corral<br>Universitat Pompeu Fabra, Barcelona, Spain, javier.vazquez@upf.edu<br>Marcelo Bertalmío<br>Universitat Pompeu Fabra, Barcelona, Spain, marcelo.bertalmio@upf.edu

Written for presentation at the
SMPTE 2017 Annual Technical Conference \& Exhibition


#### Abstract

In cinema and TV it is quite usual to have to work with footage coming from several cameras, which show noticeable color differences among them even if they are all the same model. In TV broadcasts, technicians work in camera control units so as to ensure color consistency when cutting from one camera to another. In cinema post-production, colorists need to manually colormatch images coming from different sources. Aiming to help perform this task automatically, the Academy Color Encoding System (ACES) introduced a color management framework to work within the same color space and be able to use different cameras and displays; however, the ACES pipeline requires to have the cameras characterized previously, and therefore does not allow to work 'in the wild', a situation which is very common. We present a color stabilization method that, given two images of the same scene taken by two cameras with unknown settings and unknown internal parameter values, and encoded with unknown non-linear curves (logarithmic or gamma), is able to


[^0]correct the colors of one of the images making it look as if it was captured with the other camera. Our method is based on treating the in-camera color processing pipeline as a combination of a $3 \times 3$ matrix followed by a non-linearity, which allows us to model a color stabilization transformation among two shots as a linear-nonlinear function with several parameters. We find corresponding points between the two images, compute the error (color difference) over them, and determine the transformation parameters that minimize this error, all automatically without any user input. The method is fast and the results have no spurious colors or spatio-temporal artifacts of any kind. It outperforms the state of the art both visually and according to several metrics, and can handle very challenging real-life examples.

Keywords. Color stabilization, color matching, non-linearity estimation, logarithmic encoded images, gamma corrected images.

The authors are solely responsible for the content of this technical presentation. The technical presentation does not necessarily reflect the official position of the Society of Motion Picture and Television Engineers (SMPTE), and its printing and distribution does not constitute an endorsement of views which may be expressed. This technical presentation is subject to a formal peer-review process by the SMPTE Board of Editors, upon completion of the conference. Citation of this work should state that it is a SMPTE meeting paper. EXAMPLE: Author's Last Name, Initials. 2011. Title of Presentation, Meeting name and location.: SMPTE. For information about securing permission to reprint or reproduce a technical presentation, please contact SMPTE at jwelch@smpte.org or 914-761-1100 (445 Hamilton Ave., White Plains, NY 10601).

## Introduction

This paper addresses the color stabilization problem given image pairs codified with any combination of non-linearities. In short, color stabilization can be defined as the situation where two pictures taken from the same scene differ in terms of color and where we need to match the colors of one of the images (that we call source image) to those of the other image (that we call reference image). These color differences are due to the use of different camera models (each camera sensor captures colors in different ways), or even when shooting with the same camera but using different settings (white balance, exposure time, aperture, etc.).

Let us consider the generic color processing pipeline in regular digital cameras as proposed in [1],

$$
\begin{equation*}
[R, G, B]_{\text {out }}^{t}=\left(A \cdot[R, G, B]_{\text {in }}^{t}\right)^{1 / \gamma} \tag{1}
\end{equation*}
$$

where $[R, G, B]_{\text {in }}^{t}$ corresponds to a raw triplet read by the camera sensor, $A$ is a $3 \times 3$ matrix containing information about white balance and color encoding, the power exponent $1 / \gamma$ is the non-linearity (gamma correction) applied to the linear colored image, and $[R, G, B]_{\text {out }}^{\mathrm{t}}$ is the resulting pixel value (this is not the actual value produced as an output by the camera since there are other processes we are omitting, like quantization, contrast enhancement or compression, but it is a good enough approximation: for details see [2]). Cinema cameras allow the possibility to replace the power exponent with a logarithmic function, which in a general form (common to the most popular log-encoding methods) can be expressed as:

$$
\begin{equation*}
[R, G, B]_{\text {out }}^{t}=c \cdot \log _{10}\left(a \cdot\left(A \cdot[R, G, B]_{\text {in }}^{t}\right)+b\right)+d \tag{2}
\end{equation*}
$$

where the parameters $a, b, c$ and $d$ are constant real values (varying for different camera manufacturers and camera settings), and $[R, G, B]_{\text {in }}^{t}$, A and $[R, G, B]_{\text {in }}^{t}$ are defined as before. The idea of logarithmic encoding was firstly envisioned by Kodak in a format called Cineon Log in 1992 with the goal of better storing scanned film negatives data by using more bits for dark regions, similar to the fact that humans are more sensitive to differences in dim regions. In Figure 1 we present the graphs considering linear, logarithmic and gamma responses from normalized values. Notice the difference between the output ranges for dark values, the logarithmic curve allocating a wider range for them than with gamma correction; for instance, in the example in Figure 1, gamma encoding devotes roughly $35 \%$ of the range to the darkest 10\% of linear intensity values, versus more than $70 \%$ of the range in the case of log-encoding.

Summarizing, digital photo and video cameras using gamma correction or log-encoding represent the data in a manner that can be roughly approximated by Equation (1) or (2), respectively. Different camera sensors and camera settings produce different values for $\gamma$, matrix $A$, and parameters $\{\mathrm{a}, \mathrm{b}, \mathrm{c}, \mathrm{d}\}$.
In the industry, there exist several solutions for bringing consistency across shots. They normally involve very skilled manual work, done by colorists during color grading in movie postproduction and by technicians using camera control units (CCU) [3] in live TV broadcasts; require a proper characterization of the cameras used and their settings like with the ACES framework [4], or the presence of color-charts in the shots.

In academia, to the best of authors' knowledge, except from a preliminary work of two of the authors of this paper [5], current color-matching solutions do not consider the very common cinema industry situation of logarithmic encoding, i.e. that of Equation (2).


Figure 1. Linear response vs logarithmic and gamma response. The logarithmic curve (red) was defined as $0.26 \log _{10}(72.88 \cdot \mathrm{x}+0.2)+0.39$. The gamma curve (blue) was defined as $x^{\frac{1}{2.2}}$.

Our main contribution in this work is to show how we can transform a log-encoded image (Eq. 2) into the form of a gamma-corrected image (Eq. 1). In this way, we can apply a successful colormatching method limited to work with only gamma-corrected images. Our method outperforms state-of-the-art color-matching methods, both quantitative and qualitatively.

## Related Work

In image processing and computer vision research, it is a challenge to color match a pair of pictures (reference and source) which may or may not share some content. The aim of color transfer methods is to give the appearance/look of the reference image to the source image. A seminal work in color transfer was proposed by Reinhard et al. [6], where the pair of RGB images are first converted to a decorrelated color space and then some statistics from the reference (mean and variance) are transferred to the source. Pitié et al. [7] also focused on the statistics of the images, which guide the estimation of a color mapping transform through an iterative process. The method presented by Kotera [8] computes k-means for clustering the colors of the image in order to match the principal components of the clusters of the source to the reference by a matrix multiplication (rotation and scaling). Xiao and Ma [9] also worked with color statistics (mean and covariance), but unlike Reinhard's method, the color space treated is not important. Nguyen et al. [10] presented a color transfer method that first applies color constancy to the input images, then it performs luminance matching, and finally the color gamuts are aligned by a linear transformation. Notice that [10] algorithm works for RAW input images in order to color match the linear RGB color spaces of different cameras (source and reference). In addition, it needs several images taken under different illuminants from both source and reference cameras.

Color stabilization can be understood as a constrained color transfer problem, where some regions or objects appear in both the reference and the source images. An example is the work of HaCohen et al. [11], who presented a method to compute dense correspondences between the images, combined with a global color mapping model. Vazquez-Corral and Bertalmío [2] proposed a color stabilization algorithm that consists of estimating a power law ( $\gamma$ value) for each of the images, and a single $3 \times 3$ matrix $H$, to color match the source image to the reference. It is built on the assumption that in digital cameras the color encoding can be expressed as a matrix multiplication followed by a power law (gamma correction), as expressed in Equation (1). We give some details on the method as it is the basis of the approach we introduce in this paper in the next section, but for more information we refer the reader to [2].

To the best of the authors' knowledge, there is only one color stabilization work that deals with logarithmic encoding images, that of Vazquez-Corral and Bertalmío [5]. The method relies on finding a sufficiently large number of achromatic matches among source and reference, which in some situations may be a challenging limitation.

## Background: Color stabilizing Rec. 709

Let us introduce in more detail Vaquez-Corral and Bertalmío [2] method, since it is the building part of our new methodology. Given a reference image $I_{\text {ref }}$ and a source image $I_{s r c}$, let us consider a point $p$ in the scene appearing in both images. This point produces the same irradiance triplet $[R G B]_{p}$ (linear notation) in the sensor for both images, but in general appears in different locations ( $p_{r e f}$ and $p_{s r c}$ ) and with different pixel values, $\left[R^{\prime} G^{\prime} B^{\prime}\right]_{p_{r e f}}$ and [ $\left.R^{\prime} G^{\prime} B^{\prime}\right]_{p_{\text {src }}}$ (non-linear notation), in the reference and the source. These different values can be written following Equation (1) as

$$
\left[\begin{array}{c}
R^{\prime}  \tag{3}\\
G^{\prime} \\
B^{\prime}
\end{array}\right]_{p_{r e f}}=\left(A_{r e f}\left[\begin{array}{l}
R \\
G \\
B
\end{array}\right]_{p}\right)^{\gamma_{r e f}} ;\left[\begin{array}{c}
R^{\prime} \\
G^{\prime} \\
B^{\prime}
\end{array}\right]_{p_{s r c}}=\left(A_{s r c}\left[\begin{array}{c}
R \\
G \\
B
\end{array}\right]_{p}\right)^{\gamma_{s r c}}
$$

Let us suppose we are able to estimate and undo the non-linear correction, therefore obtaining the linearized values for both images

$$
\left[\begin{array}{l}
R  \tag{4}\\
G \\
B
\end{array}\right]_{p_{r e f}}=\left(\left[\begin{array}{l}
R^{\prime} \\
G^{\prime} \\
B^{\prime}
\end{array}\right]_{p_{r e f}}\right)^{1 / \gamma_{r e f}}=A_{\text {ref }}\left[\begin{array}{l}
R \\
G \\
B
\end{array}\right]_{p} ;\left[\begin{array}{c}
R \\
G \\
B
\end{array}\right]_{p_{s r c}}=\left(\left[\begin{array}{c}
R^{\prime} \\
G^{\prime} \\
B^{\prime}
\end{array}\right]_{p_{s r c}}\right)^{1 / \gamma_{s r c}}=A_{s r c}\left[\begin{array}{c}
R \\
G \\
B
\end{array}\right]_{p}
$$

Then, by linear algebra operations, we can isolate $[R G B]_{p}$ in both equations, obtaining

$$
A_{r e f}^{-1}\left[\begin{array}{l}
R  \tag{5}\\
G \\
B
\end{array}\right]_{p_{\text {ref }}}=\left[\begin{array}{l}
R \\
G \\
B
\end{array}\right]_{p}=A_{s r c}^{-1}\left[\begin{array}{c}
R \\
G \\
B
\end{array}\right]_{p_{s r c}}
$$

Combining Equations (4) with Equation (5) we can match the colors from the source to the reference image as,

$$
\left(\left[\begin{array}{l}
R^{\prime}  \tag{6}\\
G^{\prime} \\
B^{\prime}
\end{array}\right]_{p_{r e f}}\right)^{1 / \gamma_{r e f}}=A_{r e f} \cdot A_{s r c}^{-1}\left(\left[\begin{array}{c}
R^{\prime} \\
G^{\prime} \\
B^{\prime}
\end{array}\right]_{p_{s r c}}\right)^{1 / \gamma_{s r c}}
$$

From now on we define matrix $H$ as $H=A_{\text {ref }} \cdot A_{\text {src }}^{-1}$. Therefore, Equation (6) presents 11 unknowns (the 9 positions of the matrix $H$ plus the two non-linearities). Fortunately, it is common to obtain a large enough number of pixel correspondences between images that share the same content, e.g. using SIFT [12]. Each correspondence will provide us with a different instance of Equation (6), and therefore we will end with an overdetermined system of equations, whose solution can be obtained via optimization. In particular the authors of [2] propose a 2-step minimization approach, obtaining in a first step the non-linearities and in a second step the matrix $H$.

After finding $H$ and the gamma-nonlinearities, the source image can be corrected with the following formula so that its colors match the reference:

$$
\begin{equation*}
\left(H \cdot I_{s r c}^{1 / /_{s r c}}\right)^{\gamma_{r e f}} \sim I_{r e f} \tag{7}
\end{equation*}
$$

## Methodology

We propose a general color matching method that is able to deal with source and reference images regardless of their encoding non-linearity, be it gamma or log.
Essentially, our method does the following:

- given two input images, we apply a power 10 transform to the log-encoded image(s) so that both inputs adopt the form of regular gamma-corrected images;
- on the gamma-corrected pair we apply the color stabilization method proposed in [2];
- finally, the power 10 transform is undone if necessary (i.e. if the original source image is logencoded).


## Power 10 function: from log-encoded image to gamma corrected image

Given a logarithmic encoded image LogI, we apply to it a power 10 function obtaining

$$
\begin{aligned}
\log \mathrm{I} & =\operatorname{cog}_{10}(\mathrm{a} \cdot \mathrm{I}+\mathrm{b})+\mathrm{d} \\
10^{\mathrm{Log} \mathrm{I}} & =10^{\log _{10}(\mathrm{a} \cdot \mathrm{I}+\mathrm{b})^{c}} \cdot 10^{\mathrm{d}}(8)
\end{aligned}
$$

Let us note that the parameter $b$ usually gets small values, and that setting $b=0$ does not change significantly the logarithmic curve (see Figure 2). Therefore, we neglect $b$, and Equation (8) becomes

$$
10^{\mathrm{LogI}}=(\mathrm{a} \cdot \mathrm{I})^{c} \cdot 10^{\mathrm{d}}=(\mathrm{k} \cdot \mathrm{I})^{c}(9)
$$

where $=a \cdot 10^{d / c}$. Please note the similarities between Equation (1) and Equation (9). By applying a power 10 function to the logarithmic encoded image, we have been able to obtain an image that behaves as a gamma corrected one.


Figure 2. Graph of 3 logarithmic encoded curves in three different colors: green, magenta and blue (continuous lines). In addition, the same logarithmic curves by setting $\mathrm{b}=0$ in their definitions (dashed lines). Note that the distance between the dashed and continuous lines from the same color is small.

## Color-matching: three cases

There are three cases depending on the non-linearities of the input pair (two gamma corrected inputs, one gamma corrected and the other log-encoded, or two log-encoded inputs) and for all these cases we compute the set of correspondences between the images using SIFT [15] and consider the set of matching pairs $\left(p_{r e f}, p_{s r c}\right)$ and their corresponding pixel values ( $\left[R^{\prime} G^{\prime} B^{\prime}\right]_{p_{r e f}}$ and $\left[R^{\prime} G^{\prime} B^{\prime}\right]_{p_{s r c}}$ ) in both images.

In the case that both inputs are log-encoded, we apply the power 10 function to both images and consider only the set of correspondences $10^{\left.\left[R^{\prime} G^{\prime} B^{\prime}\right]\right]_{r e f}}, 10^{\left[R^{\prime} G^{\prime} B^{\prime}\right]_{\text {src }}}$ for applying the method [5]. Then, the transformed source can be color-matched to the transformed reference:

$$
\begin{equation*}
\left[H \cdot\left(10^{I_{s r c}}\right)^{1 / \gamma_{s r c}}\right]^{\gamma_{r e f}} \sim 10^{I_{r e f}} \tag{10}
\end{equation*}
$$

The last step is to undo the transformation by applying the logarithm of base 10 to the left side of Equation (10). See Figure 3 for a diagram.
If just one input is log-encoded, we apply the power 10 transform only to this image, perform the color-matching using the method in [2], and then undo the power 10 transform if the logencoded input is the source.
If both inputs are gamma encoded then this case can directly be handled by [2].


Figure 3. Flowchart for the proposed color stabilization given a pair log-encoded images. Given two log-encoded images, $I_{r e f}$ and $I_{s r c}$, 1) we apply the power 10 function to the images, obtaining $10^{I_{r e f}}$ and $10^{I_{s r c}}, 2$ ) we estimate the parameters $\left\{\gamma_{r e f}, \gamma_{s r c}, H\right\}, 3$ ) and undo the power 10 function to the already color corrected image $\left[H \cdot\left(10^{I_{s r c}}\right)^{1 / \gamma_{s r c}}\right]^{\gamma_{r e f}}$ by applying a logarithm base 10. As sketched in the upper part of the diagram, the estimation of the non-linearities and the matrix $\left\{\gamma_{r e f}, \gamma_{s r c}, H\right\}$ is done in the following way: 1) find correspondences using SIFT [12], 2) apply power 10 to the corresponding triplets intensity values from each image, and 3) estimate the parameters using method [2].

## Results and Comparison

In this section we compare our approach with five popular color matching methods, several of which are state of the art: Reinhard et al. [6], Kotera [8], Xiao and Ma [9], Pitié et al. [7], HaCohen et al. [11]. We want to emphasize that for the last mentioned color transfer method (Pitié et al. [7]), the last stage for reducing the grain is not applied. We present results from three different experiments, by considering distinct pairs of non-linear data: 1) two log-encoded images, 2) a log-encoded and a gamma corrected one, and 3) two gamma corrected inputs. We compute quantitative results as well as qualitative for the three experiments.

## Data and Ground Truth

We have created a dataset and its corresponding ground truth (GT) to evaluate our approach in the following way. First, we captured in RAW mode using a Nikon D3100 camera model 20 different pairs of images, where each pair represents the same scene under two different viewpoints. We then apply to each of the RAW images a random $3 \times 3$ color correction matrix,
followed by a random non-linearity (either using common log-encoding parameters or random $\gamma$ values, depending on the performed experiment). In this way, we obtain a set of 20 reference and source images $\left\{I_{r e f}, I_{s r c}\right\}$ from which we know exactly the non-linearities (parameters $\left\{N L_{r e f}, N L_{s r c}\right\}$ defining the curves) and the linear transformations ( $3 \times 3$ color correction matrices $\left\{A_{r e f}, A_{s r c}\right\}$ ) applied to them. Then, for each particular pair, the ground-truth image is obtained by applying to the RAW of the source image the linear transformation and the non-linearity of the reference image, as explained in Figure 4 (by RAW we refer to the demosaiced and white balanced image in the sensor RGB color space).


Figure 4. Methods' evaluation pipeline. Given two linear input images $R A W_{1}, R A W_{2}$, we apply to them two non-linearities $\left\{N L_{r e f}, N L_{s r c}\right\}$ and two linear transformations $\left\{A_{r e f}, A_{s r c}\right\}$. In this way we obtain the reference and the source image inputs for the methods, $\left\{I_{r e f}, I_{s r c}\right\}$. A colormatching method is applied, and its outputs $I_{s r c}$ is compared to the ground truth GT (i.e. $R A W_{2}$ after applying the reference non-linearities $N L_{r e f}$ and $3 \times 3$ matrix $A_{r e f}$. All color metrics are performed in sRGB space.

## Results and Evaluation

In Figures 5, 6 and 7 we show the results from all the different methods for the three experiments we consider. For each scene (column) we present, from top to bottom: reference image, source image, ground truth image, our result, HaCohen et al. [11], Xiao and Ma [9], Reinhard et al. [6], Pitié et al. [7], and Kotera [8]. For quantitative evaluation, we select the following color metrics $\Delta E_{a b}^{*}$ [13] (the smaller the better), and color PSNR defined as CPSNR (the larger the better). The CPSNR is calculated considering the mean of the MSE of the each RGB channels. For each metric we show the maximum, mean and median values. Figure 4 shows the pipeline for the evaluation. The three experiments follow the same steps, only varying the non-linearities applied to the linear inputs. All images presented in this paper are converted into Rec. 709 images, in order to facilitate their visualization (this is possible as we have the actual values of all the non-linearities).

## Experiment 1: Log-encoded inputs

In Table 1 we present the quantitative results for the first experiment. Our algorithm outperforms the rest of algorithms in both metrics CPSNR, and $\Delta E_{a b}^{*}$, expcept for the maximum value
obtained in CPSNR by HaCohen et al. method. Xiao and Ma, and Pitié et al. obtain very similar results in $\Delta E_{a b}^{*}$, which are comparable. However, Xiao and Ma outperforms Pitié et al. in terms
Figure 5 we present the output of the algorithms for four scenarios. In the first, almost all the algorithms obtain close results to the GT, but HaCohen et al. and Pitié et al. show artifacts on the sky. In the second scene, Reinhard et al. and Xiao and Ma show problems in the bright areas, since those regions appear darker in their outputs. In the third scenario, Kotera brings the bright blue color to almost black color. In the last scene, Kotera presents the same problem as before, inverting the colors (red and green in this case).

## Experiment 2: One gamma corrected input, the other log-encoded

In the second experiment, we evaluate the performance of the methods by considering as reference a Rec. 709 image, and as source a log-encoded one. In Table 2 we present the results for both metrics CPSNR, and $\Delta E_{a b}^{*}$. Notice that the quantitative evaluation is done on all the image pairs where HaCohen et al. could compute a solution. In this experiment our algorithm outperforms the other methods, except for the maximum value in $\Delta E_{a b}^{*}$ by HaCohen et al. method. Their results are comparable to ours for both metrics.
Table 1. We show the maximum, mean and median of each metric ( $\Delta E_{a b}^{*}$ and CPSNR) considering the 20 pair of images.

|  | $\Delta \boldsymbol{E}_{\boldsymbol{a b}}^{*}$ |  |  | CPSNR |  |  |
| :--- | :---: | :---: | :---: | :---: | :---: | :---: |
|  | Max | Mean | Median | Max | Mean | Median |
| HaCohen et al. | 14.80 | 5.87 | 5.12 | $\mathbf{4 1 . 6 9}$ | 29.97 | 30.45 |
| Kotera | 26.42 | 13.41 | 11.92 | 34.67 | 24.67 | 22.97 |
| Pitié et al. | 12.84 | 6.23 | 5.26 | 36.18 | 25.77 | 24.11 |
| Reinhard et al. | 18.59 | 5.82 | 5.22 | 37.62 | 29.56 | 30.57 |
| Xiao and Ma | 21.60 | 6.69 | 5.40 | 36.62 | 29.54 | 30.13 |
| Ours | $\mathbf{1 4 . 3 6}$ | $\mathbf{3 . 9 3}$ | $\mathbf{2 . 6 9}$ | 39.91 | $\mathbf{3 2 . 1 2}$ | $\mathbf{3 3 . 0 1}$ |

Table 2. We show the maximum, mean and median of each metric ( $\Delta E_{a b}^{*}$ and CPSNR).

|  | $\Delta \boldsymbol{E}_{\boldsymbol{a b}}^{*}$ |  |  | CPSNR |  |  |
| :--- | :---: | :---: | :---: | :---: | :---: | :---: |
|  | Max | Mean | Median | Max | Mean | Median |
| HaCohen et al. | $\mathbf{5 . 7 3}$ | 3.10 | 2.61 | 42.91 | 34.57 | 36.53 |
| Kotera | 23.87 | 13.94 | 14.08 | 32.02 | 23.74 | 22.37 |
| Pitié et al. | 8.90 | 4.70 | 4.05 | 35.60 | 29.07 | 28.47 |
| Reinhard et al. | 7.65 | 4.52 | 4.32 | 35.39 | 31.09 | 31.93 |
| Xiao and Ma | 9.67 | 5.04 | 4.40 | 33.91 | 30.14 | 31.38 |
| Ours | 8.62 | $\mathbf{3 . 0 3}$ | $\mathbf{2 . 1}$ | $\mathbf{5 0 . 6 1}$ | $\mathbf{3 5 . 4 1}$ | $\mathbf{3 7 . 3 9}$ |

In Figure 6 we present the results for each algorithm, together with the reference, source and ground truth images. Although the images are presented as Rec. 709, notice that the source is a log-encoded image, and therefore most of the intensity values are shifted towards the highest value. For that reason, this experiment becomes more challenging. HaCohen et al. encounter difficulties to compute correspondences between the inputs, therefore for the third and fourth scenarios there are no outputs from this method since the first stage of the algorithm could not be performed. In the first scene, Reinhard et al. and Xiao and Ma outputs brighten the pink color
of the wall. On the second scenario, Kotera inverts the colors, bringing the blue to the red and viceversa. On the next scene, if we focus on the border of the green painting (peach pink), we see that our method is the closest one to the GT. In the last scene, Xiao and Ma presents a more intense yellow cast, see the concrete blocks. Pitié et al. presents some problems in the shadows on the building and on the pavement.

## Experiment 3: Gamma corrected inputs

In Table 3 we show the results for both metrics CPSNR, and $\Delta E_{a b}^{*}$. In this case only one image pair could not be computed by HaCohen et al. algorithm. We can see that in terms of $\Delta E_{a b}^{*}$ our algorithm outperforms the rest of approaches except for the maximum value by the Xiao and Ma method. For CPSNR metric, Reinhard et al. shows better result on median value.

In Figure 7 we show the results of the different methods given two gamma corrected images, Rec. 709, in four different scenarios. In the first scene, Kotera brings the orange color of the wall to blue. Reinhard et al., Xiao and Ma and Pitié et al. outputs brighten the orange part with respect to the ground truth. In the second one, our method preserves more the contrast of the whole image, and HaCohen et al. cannot find corresponding patches between the images. The third scene is a challenging one, since it is very colorful. If we focus on the top yellow/orange wheel, we can see how our method appears to be the closest to the ground truth. In the last scene, [9] fails by inverting colors in the image, the rest perform good, although if we focus on the top part of the image, the ground truth appears to be brighter than in the outputs of Reinhard et al., Xiao and Ma, and Pitié et al.

Our method relies on the correspondences found using SIFT. This may be an issue when no correspondences between regions of the same color are computed, since the computed $3 \times 3$ transformation matrix will not carry the color information of that particular color. Notice that other approaches for finding correspondences between pair of images could be used and integrated easily in our algorithm.
Let us mention that although our algorithm outperforms the rest of the methods both quantitatively (CPSNR and $\Delta E_{a b}^{*}$ ) and visually, the results might not be sufficient for an expert colorist. As it can be seen in Figure 5: in the second scene (second column), the color of the wall contains more orange color than the one presented as the GT; if we focus on the last scene (fourth row) green color appears darker in the GT than in our result.
Table 3. We show the maximum, mean and median of each metric ( $\Delta E_{a b}^{*}$ and CPSNR).

|  | $\Delta \boldsymbol{E}_{\boldsymbol{a b}}^{*}$ |  |  | CPSNR |  |  |
| :--- | :---: | :---: | :---: | :---: | :---: | :---: |
|  | Max | Mean | Median | Max | Mean | Median |
| HaCohen et al. | 13.45 | 4.12 | 3.75 | 44.61 | 33.22 | 31.40 |
| Kotera | 25.16 | 12.53 | 12.76 | 48.48 | 24.97 | 23.80 |
| Pitié et al. | 11.80 | 4.79 | 4.28 | 45.43 | 26.66 | 24.94 |
| Reinhard et al. | 10.03 | 3.90 | 3.73 | 44.53 | 32.51 | $\mathbf{3 3 . 4 6}$ |
| Xiao and Ma | $\mathbf{9 . 6 3}$ | 4.24 | 3.59 | 48.49 | 31.44 | 32.51 |
| Ours | 14.15 | $\mathbf{3 . 4 4}$ | $\mathbf{2 . 1 5}$ | $\mathbf{5 5 . 2 1}$ | $\mathbf{3 3 . 9 2}$ | 32.39 |



Figure 5. Results from experiment 1 (two log-encoded inputs). For each scene (columns), we present the reference, source, and GT images, and the outputs from our method, HaCohen et al. [11], Xiao and Ma [9], Reinhard et al. [6], Pitié et al. [7] and Kotera [8] (rows).


Figure 6. Results from experiment 2 (gamma corrected, and log-encoded). For each scene (columns), we present the reference, source, and GT images, and the outputs from our method, HaCohen et al. [11], Xiao and Ma [9], Reinhard et al. [6], Pitié et al. [7] and Kotera [8] (rows). The method [11] is not able to produce outputs in two of the cases (No Solution).


Figure 7. Results from experiment 3 (two gamma corrected inputs). For each scene (columns), we present the reference, source, and GT images, and the outputs from our method, HaCohen et al. [11], Xiao and Ma [9], Reinhard et al. [6], Pitié et al. [7] and Kotera [8] (rows). The method [11] is not able to produce outputs in one of the cases (No Solution).

## Conclusion

In this paper we have presented a method for color match any combination of non-linear encoded images: gamma corrected to gamma corrected, logarithmic-encoded to gamma corrected or logarithmic-encoded to logarithmic-encoded. The method is based on the modification of logarithmic-encoded images so as they behave as gamma corrected ones. In this way, we are able to apply the method defined in [2]. Our results show that our method outperforms state-of-the-art algorithms quantitatively and qualitatively, considering the three different cases. Future work will focus on generalizing the method to the case where there are no matches among the two views (e.g., the shots come from different scenes).

## Acknowledgements

This work was supported by the European Research Council, Starting Grant ref. 306337, by the Spanish government and FEDER Fund, grant ref. TIN2015-71537-P(MINECO/FEDER,UE), and by the Icrea Academia Award. The work of J. Vazquez-Corral was supported by the Spanish government under Grant IJCI-2014-19516.

## References

[1] S. Bianco, A. Bruna, F. Naccari and R. Schettini, "Color space transformations for digital photography exploiting information about the illuminant estimation process," JOSA A, vol. 29, no. 3, pp. 374-384, 2012.
[2] J. Vazquez-Corral and M. Bertalmío, "Color Stabilization Along Time and Across Shots of the Same Scene, for One or Several Cameras of Unknown Specifications," IEEE Transactions on Image Processing, vol. 23, no. 10, pp. 4564-4575, Oct 2014.
[3] W. C. Operations, " http://www.mediacollege.com/vido/production/camera-control/.," 2012. [Online].
[4] P. Postma and B. Chorley, "Colour Grading with Colour Management," in SMPTE15: Persistence of Vision - Defining the Future, 2015.
[5] J. Vazquez-Corral and M. Bertalmío, "Log-encoding Estimation for Color Stabilization of Cinematic Footage," in IEEE International Conference on Image Processing, 2016.
[6] E. Reinhard, M. Ashikhmin, B. Gooch and P. Shirley, "Color Transfer Between Images," IEEE Comput. Graph. Appl., vol. 21, no. 5, pp. 34-41, 2001.
[7] F. Pitié, A. C. Kokaram and R. Dahyot, "Automated colour grading using colour distribution transfer," Computer Vision and Image Understanding, vol. 107, no. 1, pp. 123-137, 2007.
[8] H. Kotera, "A scene-referred color transfer for pleasant imaging on display," in IEEE International Conference on Image Processing 2005, 2005.
[9] X. Xiao and L. Ma, "Color Transfer in Correlated Color Space," in Proceedings of the 2006 ACM International Conference on Virtual Reality Continuum and Its Applications, New York, NY, USA, 2006.
[10] R. Nguyen, D. K. Prasad and M. S. Brown, "Raw-to-Raw: Mapping between Image Sensor Color Responses," in Conference on Computer Vision and Pattern Recognition (CVPR), 2014.
[11] Y. HaCohen, E. Shechtman, D. B. Goldman and D. Lischinski, "Non-rigid Dense Correspondence with Applications for Image Enhancement," ACM Transactions on Graphics, vol. 30, no. 4, pp. 70:1--70:10, 2011.
[12] D. G. Lowe, "Object recognition from local scale-invariant features," in Proceedings of the Seventh IEEE International Conference on Computer Vision, 1999.
[13] CIE. Colorimetry, CIE Publication No. 15.2, Central Bureau of CIE, Vienna, 1986.


[^0]:    The authors are solely responsible for the content of this technical presentation. The technical presentation does not necessarily reflect the official position of the Society of Motion Picture and Television Engineers (SMPTE), and its printing and distribution does not constitute an endorsement of views which may be expressed. This technical presentation is subject to a formal peer-review process by the SMPTE Board of Editors, upon completion of the conference. Citation of this work should state that it is a SMPTE meeting paper. EXAMPLE: Author's Last Name, Initials. 2011. Title of Presentation, Meeting name and location.: SMPTE. For information about securing permission to reprint or reproduce a technical presentation, please contact SMPTE at jwelch@smpte.org or 914-761-1100 (445 Hamilton Ave., White Plains, NY 10601).

