Transformation of the Beta Distribution for Color Transfer

Hristina Hristova, Olivier Le Meur, Rémi Cozot and Kadi Bouatouch

University of Rennes 1, 263 Avenue Général Leclerc, 35000 Rennes, France

Abstract: In this paper, we propose a novel transformation between two Beta distributions. Our transformation progressively and accurately reshapes an input Beta distribution into a target Beta distribution using four intermediate statistical transformations. The key idea of this paper is to adopt the Beta distribution to model the discrete distributions of color and light in images. We design a new Beta transformation which we apply in the context of color transfer between images. Experiments have shown that our method obtains more natural and less saturated results than results of recent state-of-the-art color transfer methods. Moreover, our results portray better both the target color palette and the target contrast.

1 INTRODUCTION

The Gaussian distribution is a well-known and well-studied continuous unbounded distribution with many applications to image processing. The Gaussian distribution is commonly adopted to fit the distributions of various image features, such as color and light (Reinhard et al., 2001). The analytically tractable function and relative simplicity of the Gaussian distribution reveal its significance to problems like transportation optimization (Olkin and Pukelsheim, 1982), color correction for image mosaicking (Oliveira et al., 2011), example-based color transfer (Faridul et al., 2014), etc.

Color transfer between images has raised a lot of interest during the past decade. Color transfer transforms the colors of an input image so that they match the color palette of a target image. Color transfer applications include image enhancement (Hristova et al., 2015), time-lapse image hallucination (Shih et al., 2013), example-based video editing (Bonneel et al., 2013; Hwang et al., 2014), etc. Color transfer is often approached as a problem of a transfer of distributions, where the Gaussian distribution plays a significant role. Early research works on color transfer assume that the color and light distributions of images follow a Gaussian distribution. This assumption has proved beneficial for computing several global Gaussian-based transformations (Reinhard et al., 2001; Pitié and Kokaram, 2006). However, those global color transformations may produce implausible results in cases when the Gaussian model is not accurate enough. To tackle this limitation, image clustering has been incorporated into the framework of color transfer methods. A number of local color transfer methods (Tai et al., 2005; Bonneel et al., 2013; Hristova et al., 2015) adopt more precise models, such as Gaussian mixture models (GMMs), and cluster the input and target images into Gaussian clusters. This approach significantly improves the results of the color transfer.

So far, color transfer have been limited to Gaussian-based transformations. Despite the fact that color and light in images are bounded in a finite interval, such as [0, 1], they are still modelled using the unbounded Gaussian distribution. Indeed, the Gaussian distribution is commonly preferred over other types of distributions thanks to its beneficial analytical properties and its simplicity. Unfortunately, performing a Gaussian-based transformation between bounded distributions may result in out-of-range values. Such values are simply cut off and eliminated, causing over-/under-saturation, out-of-gamut values, etc., as shown in results *(a)* and *(b)* in figure 1. Furthermore, as a symmetrical distribution, the Gaussian distribution cannot model asymmetric distributions. In practice, the majority of the light and color distributions of images are left- or right-skewed, *i.e.* asymmetric. This reveals an important limitation of the Gaussian model when applied to image processing tasks and, in particular, to color transfer. To tackle these limitations of the Gaussian-based transformations, in this paper we adopt bounded distributions, and more specifically, the Beta distribution. Figure 2 illustrates the benefit of using a bounded Beta distribution to model color

Figure 1: Results *(a)* and *(b)* are obtained using two global color transfer methods (Reinhard et al., 2001; Pitié and Kokaram, 2006). The input image consists of two lightness clusters (corresponding to the two main peaks in its lightness histogram), whereas the target image is composed of three color clusters (as shown in its lightness-hue plot). Due to this fact, the global color transformations fail to produce plausible results. Result *(a)* does not match the target colors, whereas result *(b)* is significantly over-saturated, which compromises its photo-realism. Result *(c)* is obtained using a local color transfer (Hristova et al., 2015), which increases significantly the quality of the color transfer and the result.

and light.

The Beta distribution is a bounded two-parameter-dependent distribution, which can admit different shapes and thus, fit various data, bounded in a discrete interval. Adopting the Beta distribution to model color and light distributions of images is our key idea and motivation.

In this paper, we propose a novel transformation between two Beta distributions. Our transformation consists of four intermediate statistical transformations which progressively and accurately reshape an input Beta distribution into a target Beta distribution. We apply our Beta transformation both globally and locally in the context of a color transfer between images. The results, obtained using the proposed Beta transformation, appear more natural and less saturated than results of recent state-of-the-art methods. Additionally, our results represent accurately the target color palette and truthfully portray the target contrast.

The rest of the paper is organized as follows. Section 2 presents state-of-the-art color transformations. Section 3 introduces our Beta transformation. Results from applying the proposed Beta transformation in the context of color transfer are shown in section 4. The final section concludes the paper.
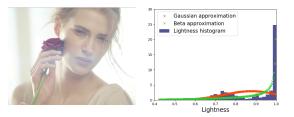

Figure 2: The right-hand plot illustrates the lightness histogram of the left-hand image as well as two approximations, namely Gaussian and Beta. The Gaussian distribution provides a poor approximation of the asymmetric lightness distribution. In contrast, the Beta distribution models the image lightness more accurately by accounting for its right-skewness.

2 RELATED WORK

This paper presents a novel Beta transformation applied to color transfer between images. Therefore, hereafter we discuss existing transformations related to color transfer. Color transfer methods are generally classified into statistical methods and geometric methods (Faridul et al., 2014). In this paper, we focus on statistical color transformations as they are more general (*i.e.* can easily be applied to a variety of image pairs), content-independent and do not require user assistance.

Statistical color transformations modify specific image features, such as color histograms and distributions. The very first statistical color transfer method is the non-parametric histogram matching. The latter transfers the target color palette to the input image by matching the input and target cumulative density functions. The full histogram transfer often results in visual artifacts. To resolve this problem, Pouli et al. (Pouli and Reinhard, 2010; Pouli and Reinhard, 2011) propose two methods for partial histogram matching at different scales.

Furthermore, a number of parametric color transformations have been proposed. They are computed using the Gaussian distribution. In this sense, Reinhard et al. (Reinhard et al., 2001) are the first to assume that each color channel of the input and target images can be well-described by a univariate Gaussian distribution. They introduce a simple color mapping between two Gaussian distributions, which is applied between each pair of input/target color channels. This method proves to be effective for natural scenes but it may introduce false colors in the result.

Reinhard et al.'s transformation is a one-dimensional transformation which does not take the correlation between the channels of the color space into account. To overcome this drawback, Pitié et al. (Pitié and Kokaram, 2006) propose a three-dimensional (3D) color transformation. Similarly to Reinhard et al., Pitié et al. describe the 3D color

distributions of the input and target images using the multivariate Gaussian distribution (MGD). This helps for their color mapping to be computed as a closed-form solution to the well-known Monge-Kantorovich's optimization problem (Evans, 1997). Pitié et al.'s color transformation is more robust than Reinhard et al.'s, as the former minimizes the displacement cost of the color transfer, preventing color swaps and preserving the intended geometry of the result.

Despite the efficiency of Pitié et al.'s method, the global Gaussian assumption may become too strong to ensure a good color transfer, as shown in figure 1. To improve the quality of the color transfer, local color transfer methods have been introduced (Tai et al., 2005; Bonneel et al., 2013; Hristova et al., 2015). Such local methods first partition the input and target images into several clusters which can be fitted by an MGD. The clustering is commonly performed using GMMs (Tai et al., 2005; Hristova et al., 2015). After the image clustering, either Reinhard et al.'s or Pitié et al.'s color transformations are carried out between the corresponding input/target clusters. The results of the local color transfer methods strongly depend on how the target clusters are mapped to the input ones. For instance, Bonneel et al. (Bonneel et al., 2013) propose a mapping function based only on the lightness of the input and target images. Hristova et al. (Hristova et al., 2015) go further and introduce four new mapping policies, which are functions of both the lightness and the chroma of the input and target images.

Local color transfer methods adopt more precise distribution models, *i.e.* GMMs, to fit the input and target color distributions, which improves significantly the results of the color transfer, as shown in figure 1. Evidently, this indicates that the more accurate the adopted distribution model, the better the results of the color transfer. We believe that we can further improve the precision of the model by using the Beta distribution. The Beta distribution can account for the skewness of the color and light distributions. Moreover, the Beta distribution, which is bounded, is a more appropriate model for the discrete color and light distributions of images than the continuous Gaussian distribution. The following section presents our novel transformation between two Beta distributions, which is later applied to color transfer between images.

3 BETA TRANSFORMATION

The present section consists of two parts: a first part presenting both the Beta distribution and well-known statistical transformations of the Beta distribution, and a second part introducing our Beta transformation.

3.1 Beta Distribution

First, we present the density function of the Beta distribution as well as several well-known statistical transformations which play an important role in the derivation of our Beta transformation.

3.1.1 Density Function

The density function $f(\cdot)$ of a Beta distributed random variable $\mathbf{x}$ with shape parameters $\alpha, \beta > 0$ (denoted $\mathbf{x} \sim Beta(\alpha, \beta)$) is given as follows:

$$f(\mathbf{x}) = \frac{1}{\mathbf{B}(\alpha, \beta)} \mathbf{x}^{\alpha-1}(1-\mathbf{x})^{\beta-1}, \qquad (1)$$

where $\mathbf{x} \in [0,1]$, and $\mathbf{B}(\alpha, \beta)$ denotes the Beta function. The Beta distribution is a univariate distribution. The variable $\mathbf{x}$, distributed according to the Beta law, can be directly transformed into a Fisher variable using a simple statistical transformation, as presented hereafter.

3.1.2 Beta-Fisher Relationship

Let $\mathbf{x} \sim Beta(\alpha, \beta)$. Then, a variable $\mathbf{y}$, obtained as $\mathbf{y} = f_{BF}(\mathbf{x}, \alpha, \beta)$, where function $f_{BF}(\cdot)$ is defined as follows:

$$f_{BF}(\mathbf{x}, \alpha, \beta) = \frac{\beta \mathbf{x}}{\alpha(1 - \mathbf{x})}, \qquad (2)$$

is a Fisher variable with shape parameters 2α and 2β (denoted $\mathbf{y} \sim \mathcal{F}(2\alpha, 2\beta)$). Equation (2) maps the bounded interval [0, 1] into the semi-bounded interval $[0, \infty)$ with $\lim_{\mathbf{x} \to 1} \mathbf{y} = \infty$.

Reversely, a variable $\mathbf{z} = f_{FB}(\mathbf{y}, \alpha, \beta)$, where $f_{FB}(\cdot)$ is obtained by the following formula:

$$f_{FB}(\mathbf{y}, \alpha, \beta) = \frac{\alpha \mathbf{y}}{\beta + \alpha \mathbf{y}}, \qquad (3)$$

and $\mathbf{y} = f_{BF}(\mathbf{x}, \alpha, \beta)$, is a Beta variable, *e.g.* $\mathbf{z} \sim Beta(\alpha, \beta)$. Equation (3) maps the semi-founded interval $[0, \infty)$ into the bounded interval [0, 1] with $\lim_{\mathbf{y} \to \infty} \mathbf{z} = 1$.
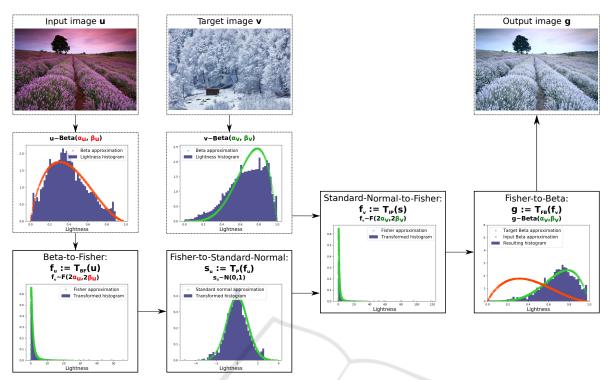
Figure 3: Our Beta transformation consists of four intermediate transformations (steps). The first two steps progressively transform the input distribution into a standard normal distribution. Then, using the target distribution, the last two transformations reshape the standard normal distribution into a Beta distribution with shape parameters close to the *target* shape parameters. For each step of our method, the flowchart shows the transformed distribution and its corresponding approximation. For illustration's sake, the input and target distributions are extracted from the lightness channels of two images. The same sequence of steps can be applied independently to each of the two chroma channels of the input and target images. The flowchart illustrates one iteration of our method.

3.2 Transformation of the Beta Distribution

Hereafter, we introduce our one-dimensional (1D) Beta transformation (it is a transformation between two univariate Beta distributions). Let **u** and **v** be 1D input and target random variables, following a Beta distribution, *i.e.* $\mathbf{u} \sim Beta(\alpha_\mathbf{u}, \beta_\mathbf{u})$ and $\mathbf{v} \sim Beta(\alpha_\mathbf{v}, \beta_\mathbf{v})$. We aim to transform variable **u** so that its distribution becomes similar to the target distribution (*i.e.* the distribution of **v**).

Transforming one Beta distribution into another one in a single pass could be challenging. To the best of our knowledge, no such transformation has yet been proposed. In this paper, we progressively transform the input Beta distribution by benefiting from the potential of well-known statistical transformations and approximations. We propose a transformation of the Beta distribution, consisting of four intermediate transformations, as follows:

1. Transformation of the input variable **u** into a Fisher variable $\mathbf{f_u} \sim \mathcal{F}(2\alpha_\mathbf{u}, 2\beta_\mathbf{u})$;

2. Transformation of the Fisher variable $\mathbf{f_u}$ into a standard normal variable $\mathbf{s_u}$;

3. Transformation of the standard normal variable $\mathbf{s_u}$ into a Fisher variable $\mathbf{f_v} \sim \mathcal{F}(2\alpha_\mathbf{v}, 2\beta_\mathbf{v})$;

4. Transformation of the Fisher variable $\mathbf{f_v}$ into a Beta variable $\mathbf{g} \sim Beta(\alpha_\mathbf{v}, \beta_\mathbf{v})$.

The key idea of our transformation consists in transforming the input Beta distribution into a standard normal distribution. The standard normal distribution can then be easily reshaped using the shape parameters of the *target* Beta distribution. However, the Beta distribution cannot be directly transformed into a standard normal distribution. To this end, we first perform an intermediate transformation to compute a Fisher distribution from the input Beta distribution. Then, we use the computed Fisher distribution to approximate the standard normal distribution. The output of the proposed transformation is the Beta variable **g** which has a distribution similar to the target Beta distribution. The steps of our transformation are illustrated in figure 3 and described hereafter.

3.2.1 Beta-to-Fisher

During the first step of our transformation, we transform the input Beta variable $\mathbf{u}$ into a Fisher variable $\mathbf{f_u} = \mathcal{T}_{BF}(\mathbf{u})$ with shape parameters $2\alpha_{\mathbf{u}}$ and $2\beta_{\mathbf{u}}$ as follows:

$$\mathcal{T}_{BF} : \mathbf{u} \rightarrow f_{BF}(\mathbf{u}, \alpha_{\mathbf{u}}, \beta_{\mathbf{u}}), \qquad (4)$$

where function $f_{BF}(\cdot)$ is defined in (2). Once we have obtained the Fisher variable $\mathbf{f_u}$, we transform it into a standard normal variable $\mathbf{s_u}$ in the second step of our transformation.

3.2.2 Fisher-to-Standard-Normal

The transformation of the Fisher variable $\mathbf{f_u}$, obtained during the first step of our method, is based on Paulson's equation (Paulson, 1942; Ashby, 1968).

In general, Paulson's equation transforms a given Fisher variable $\mathbf{f} \sim \mathcal{F}(\alpha, \beta)$ into a standard normal variable $\mathbf{s} \sim \mathbf{N}(0,1)$ as follows (see appendix A for a detailed derivation of Paulson's equation):

$$\mathbf{s} = f_P(\mathbf{f}, \mu_{\mathbf{x}}, \mu_{\mathbf{y}}, \sigma_{\mathbf{x}}, \sigma_{\mathbf{y}}, p) = \frac{\mathbf{f}^{\frac{1}{p}}\mu_{\mathbf{y}} - \mu_{\mathbf{x}}}{\sqrt{\mathbf{f}^{\frac{2}{p}}\sigma_{\mathbf{y}}^2 + \sigma_{\mathbf{x}}^2}}, \qquad (5)$$

with $\sigma_x^2 = f_\sigma(\alpha, p)$, $\sigma_y^2 = f_\sigma(\beta, p)$, $\mu_x = f_\mu(\sigma_x)$ and $\mu_y = f_\mu(\sigma_y)$, where $p \in \mathbb{N}$, and the functions $f_\sigma(\cdot)$ and $f_\mu(\cdot)$ are defined as follows:

$$f_\sigma(\gamma, p) = \frac{2}{\gamma p^2} \text{ and } f_\mu(\sigma) = 1 - \sigma^2. \qquad (6)$$

We adopt Paulson's equation (5) to transform the Fisher variable $\mathbf{f_u} \sim \mathcal{F}(2\alpha_{\mathbf{u}}, 2\beta_{\mathbf{u}})$, computed with transformation (4), into a standard normal variable $\mathbf{s_u} = \mathcal{T}_P(\mathbf{f_u})$. The transformation $\mathcal{T}_P$ is defined as follows:

$$\mathcal{T}_P : \mathbf{f_u} \rightarrow f_P(\mathbf{f_u}, \mu_\alpha^{\mathbf{u}}, \mu_\beta^{\mathbf{u}}, \sigma_\alpha^{\mathbf{u}}, \sigma_\beta^{\mathbf{u}}, p), \qquad (7)$$

where $p \in \mathbb{N}$ and $(\sigma_\alpha^{\mathbf{u}})^2 = f_\sigma(2\alpha_{\mathbf{u}}, p)$, $(\sigma_\beta^{\mathbf{u}})^2 = f_\sigma(2\beta_{\mathbf{u}}, p)$, $\mu_\alpha^{\mathbf{u}} = f_\mu(\sigma_\alpha^{\mathbf{u}})$ and $\mu_\beta^{\mathbf{u}} = f_\mu(\sigma_\beta^{\mathbf{u}})$. Function f_P is defined in (5).

3.2.3 Standard-Normal-to-Fisher

Once we have computed the standard normal variable $\mathbf{s_u}$, we inverse Paulson's equation (5) to transform $\mathbf{s_u}$ into a Fisher variable $\mathbf{f_v}$. We carry out the inversed Paulson's equation using the *target* shape parameters $\alpha_{\mathbf{v}}$ and $\beta_{\mathbf{v}}$ instead of the input shape parameters $\alpha_{\mathbf{u}}$ and $\beta_{\mathbf{u}}$.

Let α and β be any two shape parameters. We first present the inversed Paulson's equation for transforming any standard normal variable $\mathbf{s}$ into a Fisher variable $\mathbf{f} \sim \mathcal{F}(2\alpha, 2\beta)$:

$$(\mathbf{s}^2\sigma_{\mathbf{y}}^2 - \mu_{\mathbf{y}}^2)\mathbf{f}^{\frac{2}{p}} + 2\mu_{\mathbf{x}}\mu_{\mathbf{y}}\mathbf{f}^{\frac{1}{p}} + \mathbf{s}^2\sigma_{\mathbf{x}}^2 - \mu_{\mathbf{x}}^2 = 0, \qquad (8)$$

where $\sigma_x^2 = f_\sigma(2\alpha, p)$, $\sigma_y^2 = f_\sigma(2\beta, p)$, $\mu_x = f_\mu(\sigma_x)$ and $\mu_y = f_\mu(\sigma_y)$. The solutions of the inversed Paulson's equation are derived in appendix B.

Now, we solve the inversed Paulson's equation (8) using the target shape parameters $\alpha_{\mathbf{v}}$ and $\beta_{\mathbf{v}}$. In (8), we replace $\mathbf{s}$ by $\mathbf{s_u}$ (computed with transformation (7)) and the functions σ_x^2, σ_y^2, μ_x and μ_y by the following functions respectively: $(\sigma_\alpha^{\mathbf{v}})^2 = f_\sigma(2\alpha_{\mathbf{v}}, p)$, $(\sigma_\beta^{\mathbf{v}})^2 = f_\sigma(2\beta_{\mathbf{v}}, p)$, $\mu_\alpha^{\mathbf{v}} = f_\mu(\sigma_\alpha^{\mathbf{v}})$ and $\mu_\beta^{\mathbf{v}} = f_\mu(\sigma_\beta^{\mathbf{v}})$.

After solving equation (8) using the aforementioned parameters, we obtain a Fisher variable $\mathbf{f_v} \sim \mathcal{F}(2\alpha_{\mathbf{v}}, 2\beta_{\mathbf{v}})$. Each sample $\mathbf{f_v}^i$ of $\mathbf{f_v}$ is computed using a transformation $\mathcal{T}_{IP}$, i.e. $\mathbf{f_v}^i = \mathcal{T}_{IP}(\mathbf{s_u}^i)\ \forall i \in \{1, \ldots, n\}$:

$$\mathcal{T}_{IP} : \begin{cases} \mathbf{s_u}^i \rightarrow \left(f_{IP}(\mathbf{s}_i, \mu_\alpha^{\mathbf{v}}, \mu_\beta^{\mathbf{v}}, \sigma_\alpha^{\mathbf{v}}, \sigma_\beta^{\mathbf{v}}, 1)\right)^p, \text{if } \mathbf{s_u}^i < 0, \\ \mathbf{s_u}^i \rightarrow \left(f_{IP}(\mathbf{s}_i, \mu_\alpha^{\mathbf{v}}, \mu_\beta^{\mathbf{v}}, \sigma_\alpha^{\mathbf{v}}, \sigma_\beta^{\mathbf{v}}, -1)\right)^p, \text{if } \mathbf{s_u}^i \geq 0, \end{cases} \qquad (9)$$

where $p \in \mathbb{N}$ and function $f_{IP}(\cdot)$ is defined in equation (20) (appendix B).

3.2.4 Fisher-to-Beta

In the final step of our transformation, we transform the Fisher variable $\mathbf{f_v}$ into a Beta variable $\mathbf{g}$ using the following transformation $\mathcal{T}_{FB}$:

$$\mathcal{T}_{FB} : \mathbf{f_v} \rightarrow f_{FB}(\mathbf{f_v}, \alpha_{\mathbf{v}}, \beta_{\mathbf{v}}), \qquad (10)$$

where function $f_{FB}(\cdot)$ is defined in (3). The variable $\mathbf{g} = \mathcal{T}_{FB}(\mathbf{f_v})$ is approximately distributed according to a Beta distribution with shape parameters $\alpha_{\mathbf{v}}$ and $\beta_{\mathbf{v}}$, *i.e.* its distribution is similar to the target distribution.

3.2.5 Choice of p

The inverse Paulson's equation (8) has two solutions (as shown in appendix B), which can be both negative and positive, depending on the parameters $\mu_\alpha^{\mathbf{v}}$, $\mu_\beta^{\mathbf{v}}$, $\sigma_\alpha^{\mathbf{v}}$ and $\sigma_\beta^{\mathbf{v}}$. In contrast, the values of the variables, distributed according to the Fisher law, are non-negative. To ensure that each component $\mathbf{f_v}^i$ of the Fisher variable $\mathbf{f_v}$ is non-negative (see (9)), we make the parameter p equal to 4, following Hawkins et al.'s proposition (Hawkins and Wixley, 1986). By choosing $p = 4$, we also make sure that Paulson's equation (5) holds for small values of the shape parameters $\alpha_{\mathbf{u}}$ and $\beta_{\mathbf{u}}$, as discussed in (Hawkins and Wixley, 1986).

3.2.6 Iterations

For the sake of robustness, our Beta transformation can be performed iteratively using a *dynamic input variable*. In this case, at the beginning of each iteration (after the first one) the input variable is updated

using the resulting variable from the previous iteration, *i.e.*:

$$\mathbf{g}^{(t)} = \begin{cases} \mathcal{T}_{Beta}(\mathbf{u}, \mathbf{v}), \text{if } t = 1; \\ \mathcal{T}_{Beta}(\mathbf{g}^{(t-1)}, \mathbf{v}), \text{if } t \in [2, N], \end{cases} \quad (11)$$

where $\mathbf{g}^{(t)}$ denotes the result after iteration t, $\mathcal{T}_{Beta}$ denotes our Beta transformation (which is a combination of the four steps, presented earlier in this section) and N denotes the number of iterations. The benefit of iterating the proposed transformation and the choice of an appropriate value for N are discussed in the following section.

3.3 Evaluation of Beta Transformation

To evaluate the performance of our method, we carry out our Beta transformation on 111 different pairs of input and target data samples, which we model using the Beta distribution. The data samples are extracted from a collection of images containing both indoor and outdoor images.

We compute the percentage error between the shape parameters of each resulting (from the transformation) distribution and the target shape parameters. The Box-and-Whisker plots in figure 4 show the distributions of the percentage errors (for 1 and 5 iterations of our transformation). After 1 iteration, the mean percentage error is less than 0.05 and it converges towards 0 after 5 iterations.

The plots in figure 4 illustrate the benefit of performing our Beta transformation iteratively. After a single iteration, the mean percentage error is already significantly small. The more we iterate, the smaller the mean percentage error. For the application, shown in this paper (*i.e.* color transfer), our experiments have indicated that the change in the distribution of the result $\mathbf{g}^{(t)}$ becomes negligible after the fifth iteration, *i.e.* for $t > 5$, as illustrated in figure 5. Therefore, the results, shown in this paper, are obtained using five iterations of our Beta transformation, *i.e.* $N = 5$.
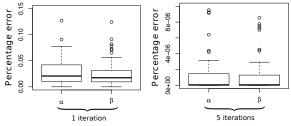

Figure 4: Box-and-Whisker plots of the percentage errors for the two Beta shape parameters α and β (computed after 1 and 5 iterations of our transformation).

4 RESULTS

We apply our Beta transformation in the context of color transfer between input and target images. We carry out our 1D Beta transformation independently on each pair of input/target color channels. To this end, we use CIE Lab, as the color space provides efficient channel decorrelation. Each of our results in this paper is obtained using 5 iterations of our transformation. We perform the Beta transformation globally (*i.e.* on the distributions of entire images) as well as locally (*i.e.* on the distributions of image clusters).

4.1 Global Color Transfer

To perform a global color transfer, we first model the distribution of each input/target channel by a Beta distribution and then, we carry out our Beta transformation between the channel distributions. In figure 6, our global Beta-based color transfer is compared to the Gaussian-based 3D linear color transformation by Pitié et al. (Pitié and Kokaram, 2006) (also carried out in CIE Lab). Target image *(a)* is a highly bright image, characterized by a low contrast (*i.e.* there is an absence of strongly defined highlights). Similarly, our first result in figure 6 does not contain highly contrasting regions, *i.e.* the illumination of our result appears uniform (much like the illumination of target image *(a)*). Moreover, our first result has a lower contrast than the input image but it looks sharper than its corresponding result by Pitié et al. We observe that our result *(a)* preserves more details from the input image (note the sharpness of the girl's face and hair in our first result). In contrast, Pitié et al.'s first result contains regions of over-saturated pixels, appearing on the girl's face, shoulder and flower (snippets *(1)* and *(2)*), causing a loss of details. Such highlights are uncharacteristic for target image *(a)* and their presence increases the contrast of Pitié et al.'s result.

Furthermore, target image *(b)* in figure 6 is characterized by a presence of strong highlights and deep shadows. Likewise, highly contrasting regions appear in our result *(b)*, whereas the absence of strong highlights in Pitié et al.'s result *(b)* influences the contrast decrease and makes the image appear flat.

4.2 Local Color Transfer

Instead of modeling the entire distributions of color and light by Beta distributions, we can build an even more accurate model using Beta mixture models (BMMs) (Ma and Leijon, 2009). We use BMMs to cluster the input and target images according to one of two components, *i.e.* lightness or hue. We adopt

Figure 5: Results of our Beta transformation (applied iteratively in the context of color transfer between images) for different number of iterations. Bottom-right corner: images of the observed difference between each result and the result, obtained after the fifth iteration. As observed, the more we iterate, the smaller the difference. The difference between the result, obtained after the fifth iteration, and the result, obtained after the tenth iteration, is negligible and therefore, we can assume that our method converges after the fifth iteration (this conclusion has been drawn using various results of color transfer). The input and target images for the color transfer are presented in figure 3.
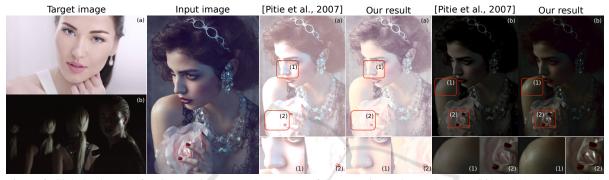

Figure 6: Global color transfer. Our results portray the contrast of the target image better than the Pitié et al.'s results and they appear sharper than Pitié et al.'s results. Snippets *(1)* and *(2)* illustrate differences between our results and Pitié et al.'s results (best viewed on screen).

the classification method, proposed in (Hristova et al., 2015), to determine the best clustering component for each image. Additionally, we let clusters overlap and we use the BMMs soft segmentation to compute the overlapping pixels. Then, we apply the mapping policies, proposed in (Hristova et al., 2015), to map the target to the input clusters and we carry out our Beta transformation between each pair of corresponding clusters (see (Hristova et al., 2015)). The overlapping pixels are influenced by more than one Beta transformation. To determine the final value of an overlapping pixel, we compute the average of all transformations, containing the pixel, using exponential decay weights like in (Hristova et al., 2015).

Figure 7 illustrates the benefit of applying a local color transfer between images. In figure 7, our global Beta transformation fails to transfer properly the target color palette to the input image. This is due to the fact that the color distributions of the input and target images in figure 7 cannot be modelled well enough by a single Beta distribution. To improve the result of the global transfer, we apply our Beta transformation locally, using 2 clusters. That way, BMMs, as more precise distribution models for the input/target distributions, help to transfer the target colors more accurately. We compare our local method with Bonneel et

al.'s (Bonneel et al., 2013) and Hristova et al.'s (Hristova et al., 2015) local methods, both of which perform clustering and carry out Gaussian-based transformations. As shown in figure 7, Bonneel et al.'s method fails to match the target floral color (making the flower in the result appear less white and slightly saturated), whereas Hristova et al.'s method causes overexposure of the foreground pixels in the result. In contrast, our local result portrays the natural cream white color of the target rose without overexposing it.

Furthermore, our *local* Beta transformation can also influence the contrast of the result. For instance, figure 8 shows that when we apply our Beta transformation *globally*, we accurately transfer the target color palette. However, that way we also decrease the contrast of the result, making it appear unnaturally under-saturated. Due to the specific contents of the input/target images in figure 8, users may expect a more saturated result with a higher contrast. To preserve the input contrast, we perform our Beta transformation *locally* using two lightness clusters, *i.e.* highlights and shadows. That way, we can transfer the target colors without compromising the input contrast.

As the input and target images in figure 8 consist of a single dominant color, the average hue provides a decent statistic for measuring the similarity of each

Figure 7: Local color transfer. The local Beta transformation is more efficient than the global Beta transformation in cases when the input/target color and light distributions cannot be well-modelled by a single Beta distribution. We use 2 clusters to obtain a naturally-looking result, matching accurately the target color palette and contrast.

Figure 8: Our local color transfer may significantly influence the contrast of the resulting image. Our global method accurately transfers the target color palette, but it also decreases the contrast, resulting in an unnaturally flat image. On the other hand, our local Beta transformation preserves the input contrast.
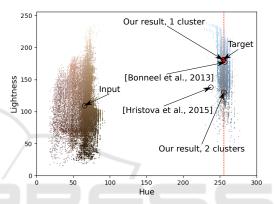

Figure 9: Lightness-hue distribution plot of the input and target images from figure 8. The left cluster corresponds to the color distribution of the input image, whereas, the right cluster corresponds to the color distribution of the target image. The red circle illustrates the mean target lightness and hue, whereas the dashed red line visualizes the mean target hue.

result to the target color palette. Figure 9 presents a plot of the lightness-hue distributions of the input and target images appearing in figure 8. In the plot of figure 9, the mean lightness and the mean hue of each result in figure 8 are displayed in circles. We observe that our local result, obtained using two clusters, has the same mean hue as the target image. In contrast, Hristova et al.'s result has an out-of-gamut mean hue. Indeed, a closer visual comparison between our local result and Hristova et al.'s result (figure 8) indicates that our local method transfers the target color palette more accurately (note the slightly different (from the target image) shade of blue in Hristova et al.'s result). The hue statistic supports this observation: our local result matches the target colors better than Hristova et al.'s result. Furthermore, the difference between our local result and the target image in terms of mean lightness is expected and can be explained by the intentional preservation of the input contrast. Finally,

more results, obtained with our local Beta transformation, are presented in figure 10.

5 CONCLUSION AND FUTURE WORK

In this paper, we have presented a transformation between two Beta distributions. Our main idea involved modelling color and light image distributions using the Beta distribution (or BMMs) as an alternative to the Gaussian distribution (or GMMs). We have applied our Beta transformation in the context of color transfer between images, though the transformation can be applied to any bounded data, following the Beta distribution law. Our results have shown the robustness of our method for transferring given target color palette and contrast.

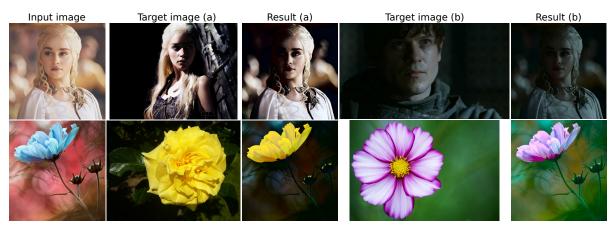The proposed 1D Beta transformation has been

Figure 10: More results, obtained with our Beta transformation. We apply our Beta transformation locally using two clusters (either lightness or hue clusters). The mapping between the input/target clusters is carried out using the mapping policies, proposed in (Hristova et al., 2015).

applied on each input channel separately without accounting for the channels inter-dependency. The latter highlights the main limitation of our method. To improve the results of our method, we can extend the dimensionality of our Beta transformation by considering the Dirichlet distribution. The Dirichlet distribution (Kotz et al., 2004) is the multivariate generalization (for more than two shape parameters) of the Beta distribution. The Dirichlet distribution with four shape parameters could be used to model any three-dimensional data, such as the joint 3D distribution of color and light in images. Given the promising color transfer results, obtained using our Beta transformation, we believe that the Dirichlet distribution would also prove beneficial to image editing and that is why, we consider it as a future avenue for improvement.

REFERENCES

Ashby, T. (1968). A modification to paulson's approximation to the variance ratio distribution. *The Computer Journal*, 11(2):209–210.

Bonneel, N., Sunkavalli, K., Paris, S., and Pfister, H. (2013). Example-based video color grading. *ACM Transactions on Graphics (Proceedings of SIGGRAPH 2013)*, 32(4):2.

Evans, L. C. (1997). Partial differential equations and monge-kantorovich mass transfer. *Current developments in mathematics*, 1999:65–126.

Faridul, H. S., Pouli, T., Chamaret, C., Stauder, J., Trémeau, A., Reinhard, E., et al. (2014). A survey of color mapping and its applications. In *Eurographics 2014-State of the Art Reports*, pages 43–67. The Eurographics Association.

Fieller, E. (1932). The distribution of the index in a normal bivariate population. *Biometrika*, pages 428–440.

Hawkins, D. M. and Wixley, R. (1986). A note on the transformation of chi-squared variables to normality. *The American Statistician*, 40(4):296–298.

Hristova, H., Le Meur, O., Cozot, R., and Bouatouch, K. (2015). Style-aware robust color transfer. *EXPRESSIVE International Symposium on Computational Aesthetics in Graphics, Visualization, and Imaging*.

Hwang, Y., Lee, J.-Y., Kweon, I. S., and Kim, S. J. (2014). Color transfer using probabilistic moving least squares. In *Computer Vision and Pattern Recognition (CVPR), 2014 IEEE Conference on*, pages 3342–3349. IEEE.

Kotz, S., Balakrishnan, N., and Johnson, N. L. (2004). *Continuous multivariate distributions, models and applications*. John Wiley & Sons.

Ma, Z. and Leijon, A. (2009). Beta mixture models and the application to image classification. In *Image Processing (ICIP), 2009 16th IEEE International Conference on*, pages 2045–2048. IEEE.

Oliveira, M., Sappa, A. D., and Santos, V. (2011). Unsupervised local color correction for coarsely registered images. In *Computer Vision and Pattern Recognition (CVPR), 2011 IEEE Conference on*, pages 201–208. IEEE.

Olkin, I. and Pukelsheim, F. (1982). The distance between two random vectors with given dispersion matrices. *Linear Algebra and its Applications*, 48:257–263.

Paulson, E. (1942). An approximate normalization of the analysis of variance distribution. *The Annals of Mathematical Statistics*, 13(2):233–235.

Pitié, F. and Kokaram, A. (2006). The linear monge-kantorovitch linear colour mapping for example-based colour transfer. In *CVMP'06*. IET.

Pouli, T. and Reinhard, E. (2010). Progressive histogram reshaping for creative color transfer and tone reproduction. In *Proceedings of the 8th International Symposium on Non-Photorealistic Animation and Rendering*, pages 81–90. ACM.

Pouli, T. and Reinhard, E. (2011). Progressive color transfer

for images of arbitrary dynamic range. *Computers & Graphics*, 35(1):67–80.

Reinhard, E., Adhikhmin, M., Gooch, B., and Shirley, P. (2001). Color transfer between images. *Computer Graphics and Applications, IEEE*, 21(5):34–41.

Shih, Y., Paris, S., Durand, F., and Freeman, W. T. (2013). Data-driven hallucination of different times of day from a single outdoor photo. *ACM Transactions on Graphics (TOG)*, 32(6):200.

Tai, Y.-W., Jia, J., and Tang, C.-K. (2005). Local color transfer via probabilistic segmentation by expectation-maximization. In *Computer Vision and Pattern Recognition, 2005. CVPR 2005. IEEE Computer Society Conference on*, volume 1, pages 747–754. IEEE.

APPENDICES

A: Paulson's Equation

Hereafter, we derive Paulson's equation (5). First, we present a well-known transformation of the Fisher distribution, which plays a key role for the derivation of Paulson's equation.

Fisher-Chi-square Relationship. A Fisher variable can be expressed as a ratio of two Chi-square variables. Let $\mathbf{x}$ and $\mathbf{y}$ be two Chi-square random variables with α and β degrees of freedom respectively ($\mathbf{x} \sim \chi^2_\alpha$, $\mathbf{y} \sim \chi^2_\beta$). Then:

$$\frac{\mathbf{x}/\alpha}{\mathbf{y}/\beta} \sim \mathcal{F}(\alpha,\beta). \qquad (12)$$

Normal-Chi-square Relationship. Let $\mathbf{z} \sim \mathbf{N}(\mu,\sigma^2)$ be a normal variable and $\mathbf{w} \sim \chi^2_\gamma$ be a Chi-square variable with γ degrees of freedom. Then, $\mathbf{z}$ can be expressed as follows (Hawkins and Wixley, 1986):

$$\mathbf{z} = \left(\frac{\mathbf{w}}{\gamma}\right)^{\frac{1}{p}}, \qquad (13)$$

where $p \in \mathbb{N}$ and μ and σ are functions of γ and p:

$$\sigma^2 = f_\sigma(\gamma, p) \quad \text{and} \quad \mu = f_\mu(\sigma). \qquad (14)$$

Functions $f_\sigma(\cdot)$ and $f_\mu(\cdot)$ are defined in (6).

Derivation of Paulson's Equation. We derive Paulson's equation using Fieller's approximation (Fieller, 1932). Fieller's approximation transforms the ratio of two normally distributed variables into a standard normal variable. Let $\mathbf{x} \sim \mathbf{N}(\mu_\mathbf{x},\sigma^2_\mathbf{x})$ and $\mathbf{y} \sim \mathbf{N}(\mu_\mathbf{y},\sigma^2_\mathbf{y})$ be two normally distributed random variables. Then, Fieller approximation (Fieller, 1932) admits the following form:

$$\mathbf{s} = f_{FL}(\mathbf{x},\mathbf{y},\mu_\mathbf{x},\mu_\mathbf{y},\sigma_\mathbf{x},\sigma_\mathbf{y}) = \frac{\frac{\mathbf{x}}{\mathbf{y}}\mu_\mathbf{y} - \mu_\mathbf{x}}{\sqrt{\left(\frac{\mathbf{x}}{\mathbf{y}}\right)^2\sigma^2_\mathbf{y} + \sigma^2_\mathbf{x}}} \qquad (15)$$

Variable $\mathbf{s}$, where $\mathbf{s} = f_{FL}(\mathbf{x},\mathbf{y},\mu_\mathbf{x},\mu_\mathbf{y},\sigma_\mathbf{x},\sigma_\mathbf{y})$, is a standard normal variable.

Paulson's equation can be derived from (15) by computing the normal variables $\mathbf{x}$ and $\mathbf{y}$ using the Chi-square distribution. In this sense, we express the normal variables $\mathbf{x}$ and $\mathbf{y}$ in (15) using equation (13) as follows:

$$\mathbf{x} = (\mathbf{w}_x/\alpha)^{\frac{1}{p}} \quad \text{and} \quad \mathbf{y} = (\mathbf{w}_y/\beta)^{\frac{1}{p}}, \qquad (16)$$

where $\mathbf{w}_x \sim \chi^2_\alpha$ and $\mathbf{w}_y \sim \chi^2_\beta$, and $p \in \mathbb{N}$. Then, we compute a new variable $\mathbf{f}$ as follows:

$$\mathbf{f} = \frac{\mathbf{x}^p}{\mathbf{y}^p} = \frac{\mathbf{w}_x/\alpha}{\mathbf{w}_y/\beta}. \qquad (17)$$

From (12), it becomes clear that $\mathbf{f}$ is a Fisher variable with shape parameters α and β, *i.e.* $\mathbf{f} \sim \mathcal{F}(\alpha,\beta)$. Once we have replaced the ratio $\mathbf{x}/\mathbf{y}$ in Fieller's approximation (15) by its equivalence from equation (17), we obtain Paulson's equation (5). Paulson's equation (5) transforms a Fisher variable into a standard normal variable.

B: Inversed Paulson's Equation

For each pair of shape parameters (α, β) and each standard normal variable $\mathbf{s}$, the inversed Paulson's equation (8) can be solved as a quadratic equation for $\mathbf{t} = \mathbf{f}^{\frac{1}{p}}$. Let $\mathbf{t} = (\mathbf{t}_1,\ldots,\mathbf{t}_n)$ and $\mathbf{s} = (\mathbf{s}_1,\ldots,\mathbf{s}_n)$, where $\mathbf{t}_i$ and $\mathbf{s}_i$ are the samples of $\mathbf{t}$ and $\mathbf{s}$ respectively. Then, the two solutions $\mathbf{t}_1$ and $\mathbf{t}_2$ of (8) are computed $\forall i \in \{1,\ldots,n\}$ using a function $f_{IP}(\mathbf{s}_i,\mu_\mathbf{x},\mu_\mathbf{y},\sigma_\mathbf{x},\sigma_\mathbf{y},C)$ (Ashby (Ashby, 1968) presents the solutions for $p = 3$):

$$\mathbf{f}_1^{\frac{1}{p}} = \mathbf{t}_1 = f_{IP}(\mathbf{s}_i,\mu_\mathbf{x},\mu_\mathbf{y},\sigma_\mathbf{x},\sigma_\mathbf{y},1), \qquad (18)$$

$$\mathbf{f}_2^{\frac{1}{p}} = \mathbf{t}_2 = f_{IP}(\mathbf{s}_i,\mu_\mathbf{x},\mu_\mathbf{y},\sigma_\mathbf{x},\sigma_\mathbf{y},-1), \qquad (19)$$

where the function $\mathbf{f}_{IP}(\cdot)$ is defined as follows:

$$f_{IP}(\mathbf{s}_i,\mu_\mathbf{x},\mu_\mathbf{y},\sigma_\mathbf{x},\sigma_\mathbf{y},C) = \frac{-\mu_\mathbf{x}\mu_\mathbf{y} + C\sqrt{D}}{\mathbf{s}_i^2\sigma^2_\mathbf{y} - \mu^2_\mathbf{y}} \qquad (20)$$

for determinant D: $D = \mathbf{s}_i^2(\sigma^2_\mathbf{x} + \sigma^2_\mathbf{y} + \sigma^2_\mathbf{x}\sigma^2_\mathbf{y}(\sigma^2_\mathbf{x} + \sigma^2_\mathbf{y} - \mathbf{s}_i^2 - 4))$, and coefficient $C = \pm 1$.