Object-to-Object Color Transfer: Optimal Flows and SMSP Transformations

Daniel Freedman and Pavel Kisilev
Hewlett-Packard Laboratories
Haifa, Israel
{daniel.freedman, pavel.kisilev}@hp.com

Abstract

Given a source object and a target object, we consider the problem of transferring the "color scheme" of the source to the target, while at the same time maintaining the target's original look and feel. This is a challenging problem due to the fact that the source and target may each consist of multiple colors, each of which comes in multiple shades. We propose a two stage solution to this problem. (1) A discrete color flow is computed from the target histogram to the source histogram; this flow is computed as the solution to a convex optimization problem, whose global optimum may be found. (2) The discrete flow is turned into a continuous color transformation, which can be written as a convex sum of Stretch-Minimizing Structure-Preserving (SMSP) transformations. These SMSP transformations, which are computed based on the color flow, are affine transformations with desirable theoretical properties. The effectiveness of this two stage algorithm is validated in a series of experiments.

1. Introduction

Recent years have witnessed an interest in the development of automatic tools for image editing. Such tools, most commonly used by graphic artists, have found their way into popular software suites such as Adobe Photoshop. In this paper, we look at one such tool, which allows the transfer of colors from one object to another.

The problem we wish to solve is as follows. Given a source object and a target object, we would like to transfer the "color scheme" of the source to the target, while at the same time maintaining the target's original look and feel. From the algorithmic point of view, our goal is to compute an appropriate color transformation which achieves this transfer. This is a challenging problem due to the fact that the source and target may each consist of multiple colors, each of which comes in multiple shades.

We propose a two stage algorithm to solve this problem. In the first stage, a discrete color flow is computed from the target histogram to the source histogram. This flow is computed as the solution to a convex optimization problem, which is an extension of the classic Transportation Problem [3, 8]; the extension relaxes the conservation constraints, and imposes a smoothness term on the flow. Since the problem is convex, the globally optimal flow may be found.

In the second stage, the discrete color flow is turned into a continuous color transformation. The continuity of the transformation is necessary, as a discrete transformation – based only on the color flow – will not capture the multiple shades of each color, and leads to strong quantization artifacts. Hence, a discrete transformation has difficulty capturing the look and feel of the target; a continuous transformation, by contrast, is suited precisely to this purpose. We write our continuous transformation as a convex sum of Stretch-Minimizing Structure-Preserving (SMSP) transformations. These SMSP transformations, which are computed based on the color flow, are affine transformations with desirable theoretical properties.

1.1. Related Work

Many recoloring papers, such as [13], focus on the problem of transferring color to gray scale images and vice versa. In [14], the authors propose a recoloring method that is based on alpha matting and compositing, and a color transformation function which is single-valued and monotonically increasing in the destination pixel intensity domain. In [5], the authors propose a natural colorization approach which is based on the assumption that pixels having similar intensities should have similar color. In this paper, the input image (or sequence) is gray scale and the output image is a colored image. In [2], the authors present a method for recoloring a destination image according to the color scheme from the source image. The image is first segmented into groups of pixels with similar color; then, the color palette for an image is constructed by choosing most typical colors from the above segments. Color transfer is computed by matching the segment areas between the source and destination segments using a Euclidean metric. In [9], the authors apply a linear transformation that scales the mean and the variance of the target area according to the ones from the source area. In [12], the authors use the approach in [9] along with a simple segmentation technique to perform recoloring operations. In [7], the authors present generic machinery for transforming probability distributions, which can be applied to the problem to the problem of color transfer.

This paper greatly extends prior work by the authors [4]. The algorithm presented in [4] was entirely discrete, with the attendant quantization artifacts. Further, a simpler version of the Transportation Problem, without the smoothness extension, was used.

1.2. Outline

The remainder of the paper is organized as follows. Section 2 presents the computation of the discrete color flow. The problem is cast as an optimization, and properties of the optimal solution are proven. Section 3 defines and computes the Stretch-Minimizing Structure-Preserving (SMSP) transformation. These transformations are quite general, and of interest independent of the color transfer application. Section 4 shows how to use the discrete color flow to compute appropriate SMSP transformations; consequently, the overall color transform may be written as a convex sum of SMSPs. Section 5 presents results, showing the effectiveness of the method.

2. Color Flow: The Convex Formulation

In this section, we compute the *color flow*, which is a key step in the development of our color transformation. The color flow describes, in discrete terms, how target colors may be mapped to source colors. This discrete color flow will then be used in Section 4 to derive a continuous transformation from target colors to source colors.

2.1. Problem Setup

Let us assume that our target and source are characterized by probability distributions. For simplicity, we take these distributions to be discrete, that is histograms. We may thus represent the target and source distributions compactly as a list of histogram bins with non-zero probability, along with their probabilities. For example, the target distribution is written as $\{(t_i, p_i^t)\}_{i=1}^{n_t}$, where t_i is a target bin-center, p_i^t is the corresponding probability mass for that bin, and n_t is the number of such bins. Likewise, the source distribution will be written as $\{(s_j, p_j^s)\}_{j=1}^{n_s}$.

Where do these distributions come from? We assume that in the first stage of the algorithm, the user selects a small region of both the target and source from the image (or images) in question. Based on these selections, a semi-automatic segmentation scheme segments the target and source regions, from which these distributions can then be computed. Details of the segmentation scheme are given in Section 5.

2.2. The Transportation Problem

Now, given the target and source distributions, we would like to find a way to map from the target to the source. Our solution to this problem is to use the classic Transportation Problem to compute the transformation between the two distributions. The Transportation Problem, which is familiar to the vision community through the work on the Earth Mover's Distance (EMD) [10], is formulated as follows [3, 8]. Let the *flow* between the target and source distributions be given by f_{ij}, where the indices i and j range over the (non-empty) bins of the target and source distributions, respectively. That is, f_{ij} can be thought of as the part of target bin i which is mapped to source bin j. A key quantity is then the *color distance* between target and source colors[1], which we denote by $D(t, s)$. The color distance conveys how similar a target color is to a source color. While we may sometimes take the color distance to be the ordinary L_2 distance between RGB vectors, we may also prefer more interesting functions. For instance, if we wish to remove shadows, then we might factor out the brightness of the colors; this "Brightness-Invariant Distance" may be written

$$D(t, s) = \left\| \frac{t}{\|t\|_1} - \frac{s}{\|s\|_1} \right\|_2$$

where we have assumed an RGB representation of color (for both target and source).[2] More examples of color distances will be given in Section 5.

Taking the color distance as given for the moment, we would like to solve the following optimization:

$$\min_{\{f_{ij}\}} \sum_{i=1}^{n_t} \sum_{j=1}^{n_s} f_{ij} D(t_i, s_j)$$

$$\text{subject to} \quad \sum_{j=1}^{n_s} f_{ij} = p_i^t \quad i = 1, \ldots, n_t \qquad \text{(CT')}$$

$$\sum_{i=1}^{n_t} f_{ij} = p_j^s \quad j = 1, \ldots, n_s \qquad \text{(CS')}$$

$$f_{ij} \geq 0 \qquad \text{(NN)}$$

The goal of the objective function is to map the target colors t_i to corresponding source colors s_j in such a way that the color distance between them is as small as possible. However, we cannot reasonably expect that each bin of the target distribution maps neatly to exactly one bin of the source distribution. Thus, we allow target bins to be spread over several source bins, subject to the two equality constraints (CT') and (CS') which ensure "conservation of probability" for both the target and source distributions.

2.3. Relaxing the Conservation Constraints

A key aspect of the Transportation Problem formulation is the conservation of probability, as embodied in the equality constraints. In fact, in our case requiring conservation of probability is too extreme as it assumes that the source and target regions contain exactly the same amounts of "comparable colors." To understand this issue, consider the following example. Suppose that the source image is 50% light red and 50% dark red, and the target image is 40% light blue and 60% dark blue. Assume further that we would like to map target pixels with a given brightness to source pixels with a given brightness, independent of color; and that the light blue and light red pixels have the same brightness, as do the respective dark colors. (The corresponding color distance will then be the "Brightness Distance," i.e. $D(t, s) = |\|t\|_1 - \|s\|_1|$, where t and s are RGB representations of color.) In this case, inspection shows that by solving the Transportation Problem, the conservation of probability constraints will require coloring part (10%) of the *dark* blue section of the target image in *light* red, which is obviously not desirable.

[1] In the work on the Earth Mover's Distance, D is generally called the *ground distance*.

[2] Note that it is possible to choose D to be asymmetric, in that the target and source are treated differently, so that D is not a true distance. See Section 5 for an example of such a D.

As a result, we would like to relax the conservation of probability constraints. To do so, we replace the equality constraints (CT') and (CS') in the original Transportation Problem by the following inequality constraints:

$$p_i^t/\eta \leq \sum_{j=1}^{n_s} f_{ij} \leq \eta p_i^t \quad i = 1, \ldots, n_t \tag{CT}$$

$$p_j^s/\eta \leq \sum_{i=1}^{n_t} f_{ij} \leq \eta p_j^s \quad j = 1, \ldots, n_s \tag{CS}$$

and add in the extra equality constraint

$$\sum_{i,j} f_{ij} = 1 \tag{TP}$$

while retaining the constraint (NN). Here $\eta \geq 1$ is a parameter which describes the slackness of the conservation constraints. Note that the constraint (TP) which was enforced implicitly in the original Transportation Problem, is now made explicit.

The role of η is elucidated in the following theorem:

Theorem 1 *Consider the optimization problem with relaxed conservation constraints. Modify (CT) to read* $p_i^t/\eta + (1 - 1/\eta)\epsilon \leq \sum_{j=1}^{n_s} f_{ij} \leq \eta p_i^t$*, for some* $\epsilon > 0$ *(*ϵ *can be chosen arbitrarily small). Then:*

- *if* $\eta = 1$*, the problem reduces to the ordinary transportation problem; and*
- *if* $\eta \to \infty$*, then for a fixed* i*, the only* j *for which* $f_{ij} > 0$ *is* $j = \arg\min_{j'} D(t_i, s_{j'})$*.*

Proof: Omitted due to space constraints. ■

The theorem thus states that the optimization problem with relaxed conservation constraints interpolates between two cases: the Transportation Problem ($\eta = 1$), and the case in which each target color t_i flows *entirely* to the source color s_j which is closest to it, according to the color distance $D(t_i, s_j)$ ($\eta = \infty$). Mostly we will use an intermediate setting, such as $\eta = 2$.

2.4. The Smoothed Relaxed Transportation Problem

The modified optimization problem is still missing one critical property. Namely, we would like the flows to be similar for similar target colors. This is important, as the final (continuous) color transformation will be based on the flow, and we would like to ensure that this transformation is sufficiently smooth. Thus, we would like to encourage smoothness in the flow itself.

We therefore modify the objective function to be

$$\sum_{i=1}^{n_t} \sum_{j=1}^{n_s} f_{ij} D(t_i, s_j) + \gamma \sum_{i \neq i'} \omega_{i,i'} \left\| \frac{f_{i,\cdot}}{p_i^t} - \frac{f_{i',\cdot}}{p_{i'}^t} \right\|_p \tag{OBJ}$$

where $f_{i,\cdot}$ denotes the vector $[f_{i,1} \ldots f_{i,n_s}]^T$, $\| \cdot \|_p$ is the p-norm, $\omega_{i,i'}$ is a measure of the similarity of t_i and $t_{i'}$, and γ is a weighting factor between the transportation and smoothness terms. For pairs i and i' with similar target colors, i.e. for which $\omega_{i,i'}$ is large, we demand that the *normalized flows* $f_{i,\cdot}/p_i^t$ and $f_{i',\cdot}/p_{i'}^t$ are similar. The reason for the normalization will become clearer in Section 4, and can be explained briefly here as follows. We will consider the normalized flow as a kind of probability distribution over source colors j corresponding a particular target color i; it is these probability distributions that will be important for the final transformation, so it is these distributions that we wish to be smooth. Note that in practice, we use a similarity measure of the form $\omega_{i,i'} = (d(t_i, t_{i'}) + \epsilon)^{-\alpha}$, where d measures the distance between target colors, and ϵ and α are small positive parameters.

Finally, we have the **Smoothed Relaxed Tranportation Problem**:

$$\min_{f_{ij}} \text{(OBJ)} \quad \text{subject to} \quad \text{(CT), (CS), (NN), (TP)}$$

If the norm p in (OBJ) satisfies $p \geq 1$, this problem is a convex program, which may be solved efficiently. Cases of interest are $p = 1$, 2, and ∞. We take $p = 2$ in our implementation, which can be solved as a quadratic program.

3. Intermezzo: Stretch-Minimizing Structure-Preserving (SMSP) Transformations

In this section, we take a brief pause from the color transfer problem to develop the idea of the Stretch-Minimizing Structure-Preserving (SMSP) transformations. These affine transformations have desirable properties, and will aid in the development of the final, continuous color transform in Section 4. However, the SMSP transformations are quite general, and of interest independent of the color transfer application.

3.1. Intuition

Suppose that our goal is to compute a transformation ϕ which maps a vector random variable Z_1 with mean and covariance matrix μ_1, σ_1, so that the transformed variable $Z_2 = \phi(Z_1)$ has mean and covariance matrix μ_2, σ_2. We will focus on affine transformations, i.e. $Z_2 = \phi(Z_1) = B_{12}Z_1 + b_{12}$.

There are multiple such transformations, but not all are equally sensible. A simple example illustrates the problem. Suppose that the vectors Z represent RGB color vectors, and that the covariance matrices are

$$\sigma_1 = \begin{bmatrix} 1 & 0 & 0 \\ 0 & \epsilon & 0 \\ 0 & 0 & \epsilon \end{bmatrix} \quad \text{and} \quad \sigma_2 = \begin{bmatrix} \epsilon & 0 & 0 \\ 0 & \epsilon & 0 \\ 0 & 0 & 1 \end{bmatrix}$$

where ϵ is a very small number. Thus, the initial distribution has various shades of red, and the final distribution has various shades of blue. Ignoring the role of the means for a moment (or altenatively, assuming zero means), the most straightforward transformation which suggests itself is

$$B_{12} = \sigma_2 \sigma_1^{-1} = \begin{bmatrix} \epsilon & 0 & 0 \\ 0 & 1 & 0 \\ 0 & 0 & \epsilon^{-1} \end{bmatrix}$$

Such a transformation will indeed ensure that the transformed variable Z_2 has covariance matrix σ_2. However, it is also clear that in some sense, this is quite an extreme transformation, which greatly squashes the red axis (by a factor of ϵ) and greatly expands the blue axis (by a factor of ϵ^{-1}). In fact, if $\epsilon \to 0$, this is a very problematic transformation.

Instead, one might consider the transformation

$$B_{12} = \begin{bmatrix} 0 & 0 & 1 \\ 0 & 1 & 0 \\ 1 & 0 & 0 \end{bmatrix}$$

This transformation does not lead to any stretching or squashing, and is in fact valid for $\epsilon = 0$. We can therefore plausibly claim that this transformation is in some sense better.

3.2. Desiderata

Let us now capture this intuition in a formal optimization framework. We would like to compute an affine transformation $Z_2 = \phi(Z_1) = B_{12}Z_1 + b_{12}$, with the following properties.

1. Structure Preservation (SP). We would like the transformation ϕ to preserve the structure of the random variable Z_1. There are many possible definitions of structure; here we focus on two natural properties.

(a) The first two moments are preserved. That is, the mean matches the desired mean:

$$\mu_2 = E[Z_2] = B_{12}E[Z_1] + b_{12} = B_{12}\mu_1 + b_{12}$$

Thus, $b_{12} = \mu_2 - B_{12}\mu_1$, and $Z_2 = \phi(Z_1) = B_{12}(Z_1 - \mu_1) + \mu_2$.

Also, the convariance matrix must match the desired covariance matrix:

$$\begin{aligned} \sigma_2 &= E[(Z_2 - \mu_2)(Z_2 - \mu_2)^T] \\ &= E[B_{12}(Z_1 - \mu_1)(Z_1 - \mu_1)^T B_{12}^T] \\ &= B_{12}\sigma_1 B_{12}^T \end{aligned}$$

(b) Orthogonality of the principal directions is preserved. Here, we focus on the principal directions of the covariance matrix σ_1. Specifically, we insist that the orthogonality of the principal directions is preserved under the transformation ϕ.

Let $\sigma_1 = U_1 D_1 U_1^T$ be the eigendecomposition of σ_1, so that the principal directions are the columns of the matrix U_1. Then structure preservation says that $B_{12}U_1$ is an orthgonal, though *not* orthonormal, matrix. We may rewrite this as $(B_{12}U_1)^T(B_{12}U_1)$ is diagonal.

2. Stretch Minimization (SM). We measure the *stretch* of the transformation ϕ in terms of how much the transformation squashes or stretches the axes (see the example in Section 3.1). A natural way to measure this is in terms of the singular values of the matrix B_{12}. In particular, let the singular values of B_{12} be denoted $s_i(B_{12})$. Then we can define the stretch of the transformation ϕ as

$$\Theta(\phi) = \Theta(B_{12}) = \sum_{i=1}^{3} \theta(s_i(B_{12}))$$

where $\theta(s)$ punishes values that are far away from 1, i.e. that constitute a stretch of a squash. Possibilities for θ include

$$\theta(s) = |\log(s)| \quad \text{and} \quad \theta(s) = \max\{s, s^{-1}\} = e^{|\log(s)|}$$

Note that in cases where one of the covariance matrices is singular or nearly singular, we may wish to use a robust version of the stretch penalty, i.e. $\theta_r(s) = \min\{\theta(s), \theta_0\}$. For example, if $\theta(s) = \max\{s, s^{-1}\}$, we might take $\theta_0 = 10$.

3.3. Definition and Computation of the SMSP

To summarize: we would like to compute an affine transformation $Z_2 = \phi(Z) = B_{12}Z_1 + b_{12}$ which preserves structure and minimizes stretch. We have already shown that preservation of means implies that $b_{12} = \mu_2 - B_{12}\mu_1$. It remains to compute B_{12}. Using the other properties described above, this leads to the following optimization:

$$\begin{aligned} \min_{B_{12}} \Theta(B_{12}) \text{ subject to } & B_{12}\sigma_1 B_{12}^T = \sigma_2 \\ & (B_{12}U_1)^T(B_{12}U_1) \text{ is diagonal} \end{aligned} \tag{1}$$

The solution to this optimization is contained in the following theorem.

Theorem 2 *Let $\sigma_k = U_k D_k U_k^T$ and let $d_{k,i} = (D_k)_{ii}$ for $k = 1, 2$ and $i = 1, 2, 3$. Then the solution to (1) is $B_{12}^* = U_2 S_{12} P_{12} U_1^T$, where*

- *P_{12} is a permuation matrix, corresponding to the permutation π which solves*

  $$\min_{\pi} \sum_{i=1}^{3} \theta((d_{2,i}/d_{1,\pi(i)})^{1/2})$$

  This is an instance of the assignment problem, which can be solved e.g. by the Hungarian Method [6].

- *S_{12} is a diagonal matrix whose ii^{th} entry is $(d_{2,i}/d_{1,\pi(i)})^{1/2}$ (where π is as above).*

Proof: Omitted due to space constraints. ■

The SMSP transformation is thus easy to compute, requiring only diagonalization and singular value decomposition of small (3×3) matrices, as well as the solution of a small Assignment Problem (again 3×3).

We will now turn to the issue of incorporating these SMSPs within the color transfer algorithm.

4. The Color Transform as a Convex Sum of SMSP Transformations

Our goal is to compute a color transformation, which maps target colors to source colors. We have computed a discrete color flow, f_{ij}, from the target distribution to the source distribution, by solving the Smoothed Relaxed Transportation Problem. In this section we show how to turn such a discrete flow into a continuous transformation, using the machinery of SMSP transformations.

4.1. From Flows to Pairs of Random Variables

We would like to map from the discrete color flow to a continuous color transformation. The key to this procedure is to think of the *normalized flow* for each target bin i and each source bin j as probability distributions, which can tell us which target colors ought to be mapped to which source colors. These normalized flows will lead to the definition of pairs of random variables for each target bin i; we can then map from the first such variable to the second using an SMSP transformation.

Let us begin with the normalized flows. Formally, the normalized target and source flows are, respectively,

$$p_{ij}^{ts} = \frac{f_{ij}}{\sum_{j'} f_{ij'}} \quad \text{and} \quad p_{ji}^{st} = \frac{f_{ij}}{\sum_{i'} f_{i'j}}$$

The normalized target flow (denoted with the superscript ts to indicate the mapping from target to source) gives the distribution of source colors corresponding to the i^{th} target color. Similarly, the normalized source flow (denoted with superscript st) gives the distributions of target colors corresponding to the j^{th} source color. Our goal is to use these normalized flows to compute statistics, i.e. means and covariance matrices, of a pair of vector random variables associated with the i^{th} target color.

Let us begin with the *source random variable* associated with target bin i. The distribution of source bins associated with target color i is simply the normalized source flow, p_{ij}^{ts}. We wish to compute the statistics of the source random variable associated with bin i, that is the mean $\tilde{\mu}_i^s$ and covariance matrix $\tilde{\sigma}_i^s$. Suppose that we have the mean and second moment associated with each source bin, i.e.

$$\mu_j^s = E[s|s \in \text{bin}_j] \quad \text{and} \quad v_j^s = E[ss^T|s \in \text{bin}_j]$$

(In practice, μ_j^s and v_j^s are computed from the samples within bin j in the standard way.) Then the statistics of the source colors associated with bin i are

$$\tilde{\mu}_i^s = \sum_j p_{ij}^{ts} \mu_j^s \quad \text{and} \quad \tilde{\sigma}_i^s = \sum_j p_{ij}^{ts} v_j^s - \tilde{\mu}_i^s (\tilde{\mu}_i^s)^T \quad (2)$$

Now, let us turn to the *target random variable* associated with target bin i, by examining the distribution of target bins associated with target color i. One might think that this is naturally the degenerate distribution with all of its weight on bin i. However, using our knowledge of the flow f, we can be more sophisticated in our analysis. In particular, the normalized target flow p_{ij}^{ts} indicates the source bins j associated with the target bin i; but for each such source bin j, the normalized source flow p_{ji}^{st} indicates the target bins i associated with that bin. Thus, we can combine the normalized target and source flows to understand which target bins are related, where by related we mean that they map to the same source bins j. Specifically, we define

$$p_{ii'}^{tt} = \sum_j p_{ij}^{ts} p_{ji'}^{st}$$

or in matrix notation, $p^{tt} = p^{ts} p^{st}$. This distribution relates target bin i to target bin i'; $p_{ii'}^{tt}$ will be large if target bins i and i' map to similar source bins, and small if this is not the case. It is easy to verify that p^{tt} is a stochastic matrix, i.e. $\sum_{i'} p_{ii'}^{tt} = 1$.

Given $p_{ii'}^{tt}$, we wish to compute the statistics of the target random variable associated with bin i, that is the mean $\tilde{\mu}_i^t$ and covariance matrix $\tilde{\sigma}_i^t$. We proceed precisely as in the case of the source statistics. Suppose that we have the mean and second moment associated with each target bin, i.e.

$$\mu_i^t = E[t|t \in \text{bin}_i] \quad \text{and} \quad v_i^t = E[ss^T|y \in \text{bin}_i]$$

(In practice, μ_i^t and v_i^t are computed from the samples within bin i in the standard way.) Then the statistics of the target colors associated with bin i are

$$\tilde{\mu}_i^t = \sum_{i'} p_{ii'}^{tt} \mu_{i'}^t \quad \text{and} \quad \tilde{\sigma}_i^t = \sum_{i'} p_{ii'}^{tt} v_{i'}^t - \tilde{\mu}_i^t (\tilde{\mu}_i^t)^T \quad (3)$$

4.2. The Color Transform

Now, we would like to compute the optimal transformation for bin i. In particular, we would like to transform a distribution with mean and covariance matrix $\tilde{\mu}_i^t, \tilde{\sigma}_i^t$ to one with mean and covariance matrix $\tilde{\mu}_i^s, \tilde{\sigma}_i^s$. To do so, we use the SMSP machinery, and denote the resulting transformation as $\Psi_i(t)$. We know from Section 3 that the resulting transformation will have as small a stretch as possible, while preserving the structure of the target random variable.

Finally, we write the overall transformation as a convex sum of such SMSP transformations, that is

$$\Psi(t) = \sum_{i=1}^{n_t} w_i(t) \Psi_i(t) \quad (4)$$

where the non-negative weights $w_i(t)$ sum to 1, i.e., $\sum_i w_i(t) = 1$, with $w_i(t_i) = 1$. In other words, the overall transformation $\Psi(t)$ interpolates the bin-by-bin transformations $\Psi_i(t)$. In practice, we take

$$w_i(t) = \frac{d(t, t_i)^{-\alpha}}{\sum_{i'} d(t, t_{i'})^{-\alpha}}$$

where d is the distance between two target colors, and $\alpha > 0$.

An interesting property of this algorithm occurs when the smoothness weight, γ, goes to ∞ in the computation of the color flow in the Smoothed Relaxed Transportation Problem. Let $\mu_{tot}^s, \sigma_{tot}^s$ be the mean and covariance matrix of the *entire* collection of source vectors (i.e. ignoring the histogram bins, and simply pooling all source vectors together), and likewise for $\mu_{tot}^t, \sigma_{tot}^t$. Then:

Theorem 3 *Let the color flow f_{ij} be computed as the solution to the Smoothed Relaxed Transportation Problem, with $\eta = 1$ and $\gamma \to \infty$. Then the color transformation $\Psi(t)$ is the SMSP transformation from the random variable with mean and covariance matrix $\mu_{tot}^t, \sigma_{tot}^t$ to the random variable with mean and covariance matrix $\mu_{tot}^s, \sigma_{tot}^s$.*

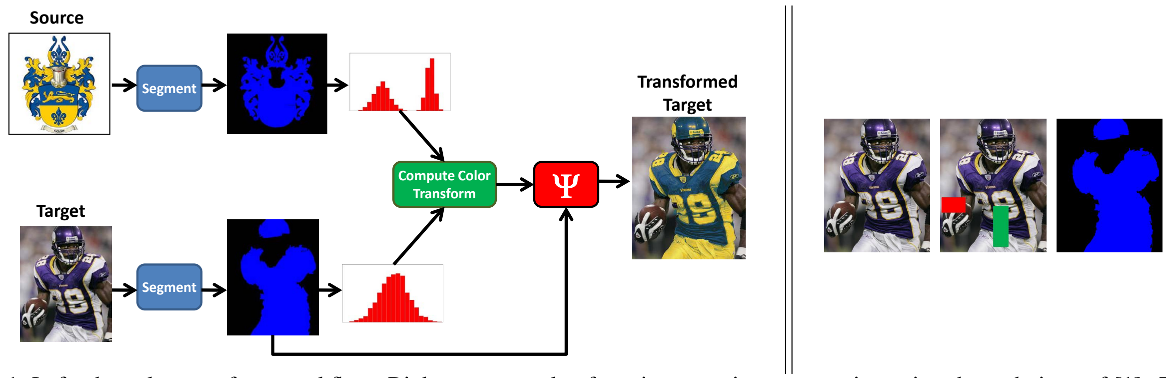

Figure 1. Left: the color transform workflow. Right: an example of semi-automatic segmentation using the technique of [1]. The green rectangle is the object seed, and the red rectangle is the background seed.

Proof: Omitted due to space constraints. ■

This property is quite useful. It says that as the smoothness term becomes the dominant term, the transformation approaches the transformation one would get by considering the source and target distributions to be unimodal. In that sense, in the limit as $\gamma \to \infty$, the solution should approach something like that of Reinhard *et al.* [9]. However, the solution will still differ, as the SMSP transformation is more robust than the straighforward affine transformation of [9].

4.3. Putting it All Together

Let us summarize the construction of our color transformation.

1. Compute the optimal color flow f_{ij} by solving the Smoothed Relaxed Transportation Problem.

2. Compute the statistics of each target bin i using the color flow: the target mean and covariance matrix $(\tilde{\mu}_i^t, \tilde{\sigma}_i^t)$ and source mean and covariance matrix $(\tilde{\mu}_i^t, \tilde{\sigma}_i^t)$, using Equations (3) and (2) respectively.

3. For each target bin i, compute the SMSP transformation $\Psi_i(t)$ which maps the target mean and covariance matrix $(\tilde{\mu}_i^t, \tilde{\sigma}_i^t)$ to the source mean and covariance matrix $(\tilde{\mu}_i^t, \tilde{\sigma}_i^t)$ using Theorem 2.

4. Transform any target color t as a weighted sum of the optimal SMSPs, according to Equation (4).

5. Results

We begin by describing the overall color transform workflow, which is detailed on the left side of Figure 1. Both source and target images are segmented to yield the source and target regions, respectively. The segmentation technique used is an interactive, or semi-automatic technique, based on the random walker with priors technique of Grady [1]. The user is required to select a rectangle[3] for both the object and the background; based on these "seeds," the algorithm computes the most likely segmentation. The procedure is illustrated on the right side of Figure 1.

[3]It is also possible to select more than one rectangle for either the object or background; however, this was only necessary for the hockey player image, in which the object and background have similar color distributions.

$D(t,s)$	Description	Examples
$\lVert t-s \rVert_2$	"Ordinary L_2"	1, 2, 4
$\lvert \lVert t \rVert_1 - (1 - \lVert s \rVert_1) \rvert$	"Inverted Brightness Distance"	3
$\left\lVert \frac{t}{\lVert t \rVert_1} - \frac{s}{\lVert s \rVert_1} \right\rVert_2$	"Brightness-Invariant Distance"	5

Table 1. The color distances used in the experiments. The example number refers to the row number in Figure 2.

Returning to the color transform workflow, the source and target segmentations yield histograms, from which the color transformation Ψ may be computed, as described in Sections 2, 3 and 4. This transform is then applied to all pixels within the target region, finally yielding the transformed output image.

We chose the following parameter values for the experiments. η, the slackness in the conservation constraints, is set to 2. p, the norm used in the smoothness term, is also set to 2, so that quadratic programming may be used. γ, the weight of the smoothness term, is set to 1. The stretch function $\theta(s) = \min\{\max\{s, s^{-1}\}, 10\}$. Finally, α, which is used to compute the interpolation weights $w_i(t)$, is set to 4. The color distance $D(t,s)$ varies from example to example, and is listed in Table 1.

5.1. Experiments

In each experiment, the proposed method is compared with the popular method of Reinhard *et al.* [9]. The latter technique transforms the source and target regions to Lab color space, and then transforms the mean and variance of each of the three channels separately in the natural way. We of course use the segmentations as preprocessing for both the proposed method and the method of Reinhard *et al.*

The results are shown in Figure 2, with the color distance for each example given in Table 1; in all cases, the colors used are RGB. Note that in the case of the Inverted Brightness Distance, the formula for $D(t,s)$ assumes that s and t have been normalized so that $\lVert s \rVert_1, \lVert t \rVert_1 \in [0,1]$ for all

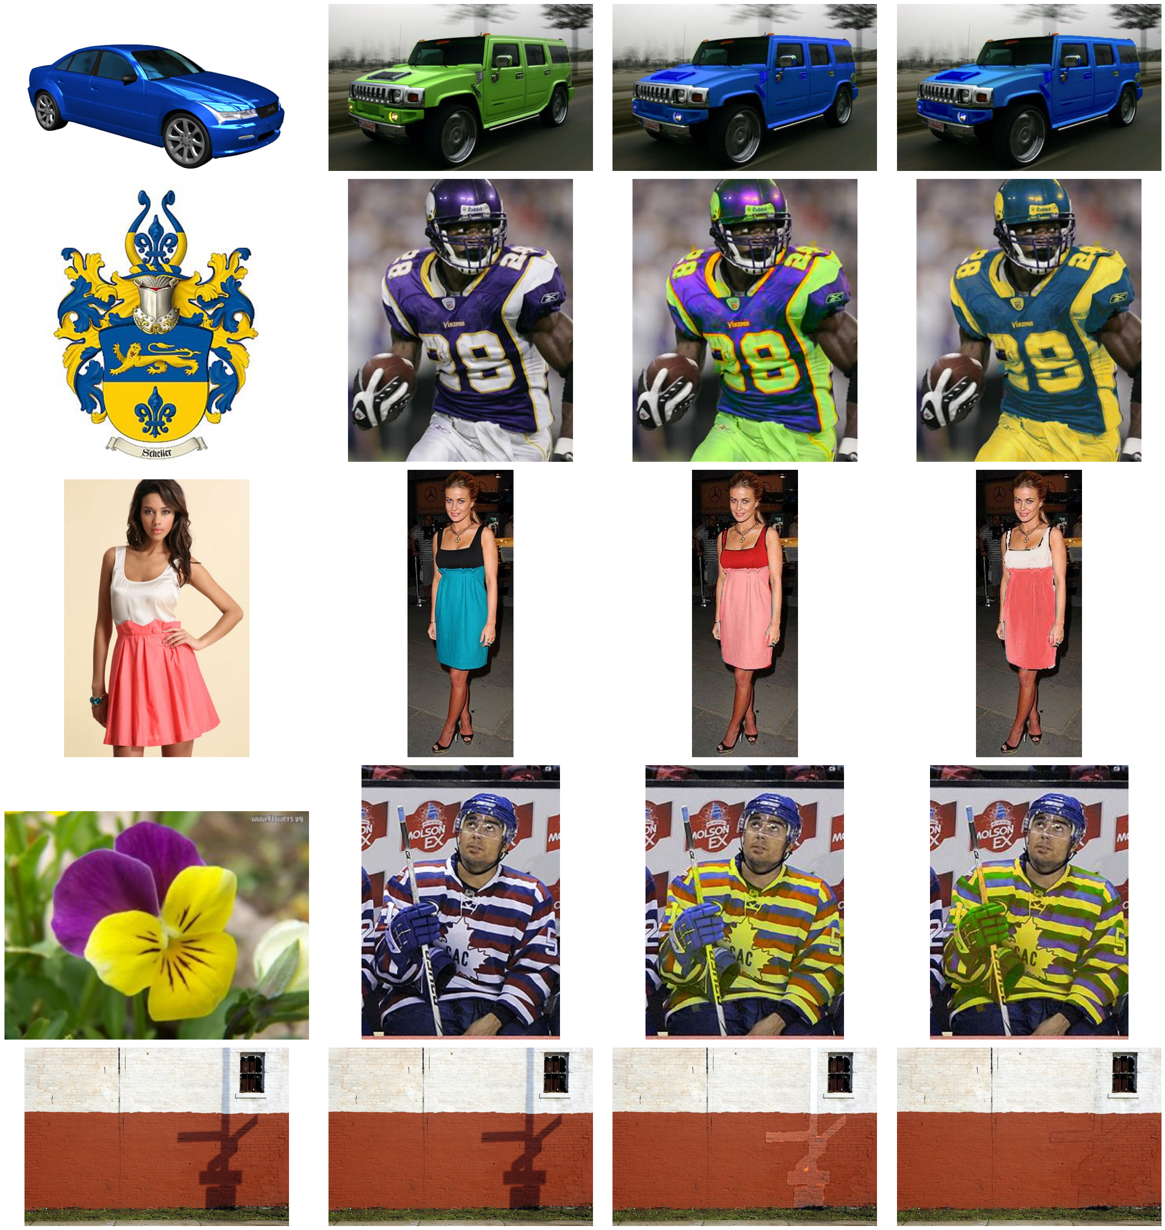

(a) (b) (c) (d)

Figure 2. (a) Source image. (b) Target image. (c) Method of Reinhard *et al.* [9]. (d) Proposed Method.

source and target vectors. The segmentation masks for both source and target are omitted due to space constraints.

The first row of Figure 2 shows a green Hummer being recolored blue. Both the proposed method and the Reinhard method succeed here, which is not surprising given the unimodal nature of both source and target distributions. Note that the proposed method uses the Ordinary L_2 distance between RGB vectors.

The situation becomes more interesting in the second row, in which the source is a crest and the target is a football player's uniform. The color distance used here is again the Ordinary L_2 distance between RGB vectors. In this case, both the source and the target have two modes. Not surprisingly, the method of Reinhard fails here; the purple on the uniform is turned to a slightly brighter purple, while the white is turned to a pale green. Neither of these colors is found in the source image. By contrast, the proposed method performs in the natural way, turning the purple on the uniform to the blue of the crest, and the white on the uniform to the yellow of the crest. Note that the look and feel of the uniform is retained – for example, creases in the jersey are maintained, and darker purple regions are mapped to darker blue regions. This can be credited to the locally SMSP behavior of the color transform.

The third row presents another example in which both source and target are bimodal. For a different take, we use a different color distance, the Inverted Brightness Distance. This distance matches dark source colors to bright target colors and vice versa. In this case, note the success of the proposed method in correctly mapping the white source region to the black target region, and the pink source region to the blue target region. Due to the use of this "inverted" color distance, there are some artifacts produced: the shadows on the dress (darker blue regions) are mapped to lighter pink regions. Nonetheless, the look and feel of the target is largely retained, but with the source color scheme. The Reinhard algorithm produces a random coloration again, with the jersey colored in red and the skirt in a very light pink, neither of which appear in the source image. (Note that both algorithms suffer some artifacts at the top of the shirt, which is colored in black; this is due to a mistake in the segmentation, and is unrelated to the color transforms.)

The fourth row shows a trimodal example. Here, the source consists of the purple and yellow flowers, along with the green grass and leaves; while the target consists of the hockey player's jersey and glove, which are colored in white, red, and blue. Here we used the Ordinary L_2 distance between RGB vectors as the color distance. The proposed method plausibly recolors the shirt, mapping white to yellow, red to purple, and blue to green. The Reinhard method also maps white to yellow (albeit a paler yellow), but retains the red and blue of the original target, simply brightening each.

The fifth row shows a shadow removal example. In this case, the source and target images are the same, but the source region is the region which is lit, and the target region is in shadows. The color distance is now naturally the Brightness-Invariant Distance. The proposed method does a good job; while leaving some artifacts around the edge of the shadow (perhaps penumbra), the shadow is largely removed. The Reinhard method suffers the usual problems, inventing colors which are not part of the source. Note that there are many well developed techniques for shadow removal, such as [11], which may perform very well on such an image. Our goal is not to claim dominance in the area of shadow removal, but simply to suggest another possible use for the proposed framework.

Our final result is shown in Figure 3. Here we show the results of the proposed method with two different color distances. The result are shown in the bottom row: the left image uses the Brightness Distance $D(t,s) = |\|t\|_1 - \|s\|_1|$ while the right image uses the Inverted Brightness Distance (see Table 1). Thus, the former image maps the stripes on the target fish to white stripes, whereas the latter image maps them to black stripes. Due to the segmentation, which includes the fish's eye, the latter example also whitens the eye of the fish.

Figure 3. Results with different color distances. Top left: source image. Top right: target image. Bottom left: transformed image using Brightness Distance. Bottom right: transformed image using Inverted Brightness Distance.

References

[1] L. Grady. Multilabel random walker image segmentation using prior models. In *CVPR*, volume 1, page 763, 2005.

[2] G. Greenfield and D. House. Image recoloring induced by palette color associations. *J. WSCG*, 11(1):189–196, 2003.

[3] F. Hitchcock. The distribution of a product from several sources to numerous localities. *J. Math. Phys. Mass. Inst. Tech*, 20:224–230, 1941.

[4] P. Kisilev and D. Freedman. Color transforms for creative image editing. In *Color Imaging Conference*, 2009.

[5] A. Levin, D. Lischinski, and Y. Weiss. Colorization using optimization. *ACM TOG*, 23(3):689–694, 2004.

[6] J. Munkres. Algorithms for the assignment and transportation problems. *J. SIAM*, pages 32–38, 1957.

[7] F. Pitie, A. Kokaram, and R. Dahyot. N-Dimensional Probability Density Function Transfer and its Application to Colour Transfer. In *ICCV*, pages 1434–1439, 2005.

[8] S. Rachev. The Monge–Kantorovich mass transference problem and its stochastic applications. *Theory of Probability and its Applications*, 29:647, 1985.

[9] E. Reinhard, M. Ashikhmin, B. Gooch, and P. Shirley. Color transfer between images. *IEEE Computer Graphics and Applications*, pages 34–41, 2001.

[10] Y. Rubner, C. Tomasi, and L. Guibas. The earth mover's distance as a metric for image retrieval. *International Journal of Computer Vision*, 40(2):99–121, 2000.

[11] Y. Shor and D. Lischinski. The Shadow Meets the Mask: Pyramid-Based Shadow Removal. In *Computer Graphics Forum*, volume 27, pages 577–586, 2008.

[12] Y. Tai, J. Jia, and C. Tang. Local color transfer via probabilistic segmentation by expectation-maximization. In *CVPR*, volume 1, 2005.

[13] T. Welsh, M. Ashikhmin, and K. Mueller. Transferring color to greyscale images. In *Proceedings of the 29th annual conference on Computer graphics and interactive techniques*, pages 277–280. ACM New York, NY, USA, 2002.

[14] X. Xiao, L. Ma, and M. Kunze. Object-Based Image Recoloring Using Alpha Matte and Color Histogram Specification. *Lecture Notes in Computer Science*, 4270:194, 2006.