

# Optimal Transportation for Example-Guided Color Transfer

Oriel Frigo, Neus Sabater, Vincent Demoulin, Pierre Hellier

Technicolor Research and Innovation  
Avenue des Champs Blancs, 35570, Cesson-Sévigné, France

**Abstract.** In this work, a novel and generic method for example-based color transfer is presented. The color transfer is formulated in two steps: first, an example-based Chromatic Adaptation Transform (CAT) has been designed to obtain an illuminant matching between input and example images. Second, the dominant colors of the input and example images are optimally mapped. The main strength of the method comes from using optimal transportation to map a pair of meaningful color palettes, and regularizing this mapping through thin plate splines. In addition, we show that additional visual or semantic constraints can be seamlessly incorporated to obtain a consistent color mapping. Experiments have shown that the proposed method outperforms state-of-the-art techniques for challenging images. In particular, color mapping artifacts have been objectively assessed by the Structural Similarity (SSIM) measure [26], showing that the proposed approach preserves structures while transferring color. Finally, results on video color transfer show the effectiveness of the method.

## 1 Introduction

Color transfer is the process to modify the color of an *input* image according to the color palette of an *example* image. The color changes could be more or less realistic (or artistic) depending on the choice of the *example* image and the final sought appearance. Image and video color transfer is useful for a variety of computer vision applications. For example, coloring gray scale images, picture contrast change acquired in bad climatic conditions or image denoising using flash/no flash pairs.

In the movie industry, to ensure that a movie has the appropriate color mood, color grading (changes in color, contrast, white balance and hues) is performed manually by the colorist, a highly qualified art professional, who could clearly benefit from automatic and consistent color transfer methods. Similarly, editing tools for large personal photo collections could be improved with example-based algorithms [8, 11], specially considering the last trend of editing sets of multi-contributor images of a given event (party, wedding, *etc*).

A common drawback of color transfer methods is the strong tendency to create undesired visual artifacts. For instance, existing noise or compression “block effects”, that are initially barely noticeable, can become prominent. Hence, in

order to achieve automatic color transfer, the considered method must achieve a visually-plausible and appropriate color transformation. Artifacts and unnatural colors should be avoided, while computational space and time complexity should be kept as low as possible. Considering these requirements, in this paper we propose an example-based method for automatic color transfer where the illuminant matching and the transfer of dominant colors of the images are treated separately. Moreover, our method carries out a final regularization of the color transform avoiding new parasite structures in the image. Thanks to this regularization, the color transformation can easily be applied to videos without any color flickering. Experimental video results are included on the supplementary material.

## 2 Related Work

Reinhard *et al.* [21] were pioneers in establishing the concept of color transfer, with an approach to modify the color distribution of one given original image based on the global color statistics (mean and variance) of an example image in the decorrelated color space  $l\alpha\beta$ . Other works have proposed global color transfer in terms of non-linear histogram matching [16, 17, 19] or N-Dimensional Probability Distribution Function (N-PDF) transfer [18]. Removing the inherent artifacts due to color modifications is the main goal of some other works. For example, in [20] a nonlocal filter is studied to remove spatial artifacts.

Related work also include [24] in which the color transfer is defined on color segments given by an Expectation-Maximization adapted algorithm. However their color mapping is essentially different from our color mapping which is based on the Monge-Kantorovitch optimal transportation problem. In [6, 5] the color transfer problem is also presented in terms of optimal transportation. Other works adapted the flow-based color transfer representing the colors in the image by compact signatures given by Gaussian Mixture Models [15] or super-pixel segmentation [27].

Another class of methods such as [7, 23, 10] assumes that there are spatial correspondences to be found between the input and example image, these correspondences being used to derive a color transformation. The assumption of geometrical relationship drastically reduces the scope and genericity of the method.

Few works have introduced semantic constraints in the context of color transfer or color enhancement. In [3], a semantic annotation of input and example images is obtained by training a Support Vector Machines and classifying regions in the image according to a trained model. The method proposed by [13] performs semantic image enhancement based on correspondences between image characteristics and semantic concepts.

Despite the significant progress made since the Reinhard *et al.* seminal paper [21], color transfer remains a challenging problem. Indeed, we think that the current approaches are still limited in some situations (strongly dependent on the selected images) and are prone to create image artifacts. On the one hand, linear color transfer methods are robust and usually do not introduce noise. However,

they do not perform well when there are several objects with different colors, since the linear transformation cannot account for the magnitude of color change that is expected. On the other hand, highly non-linear transformations seem to be more robust but at the cost of amplifying noise and introducing undesired structures when a local histogram stretching arises. Besides, all techniques may transform the image with unnatural results. For example, an object receiving non-plausible colors as a green face.

### 3 Our Approach

In this section, we present two contributions to color transfer: an example-based CAT for illuminant matching (Sec. 3.1) and a color transfer based on automatic color palette associations (Sec. 3.2). These two methods are independent and complementary, and achieve convincing results for challenging images when used together. Moreover, we show how our color transfer can optionally be constrained with semantic attributes like saliency or faces (Sec. 3.3).

In order to limit the aforementioned artifacts, we propose to process separately the luminance and chroma channels of the image. Basically, the luminance channel will be addressed using a novel example-based CAT, accounting for the illuminant change, while the chroma channels will be transformed using optimal transportation. In fact, we have observed a substantial improvement in our results with this approach compared to other color transfer techniques that treat jointly luminance and chroma.

#### 3.1 Example-based CAT

In [21], it is mentioned that one interesting application for color transfer is to remove undesirable color cast from an image, such as the yellowish colors in photos taken under incandescent illumination. Although this description reminds the color constancy problem, as far as we know, color constancy and chromatic adaptation has not been approached in the color transfer literature. In digital photography, adjusting the lighting is known as white-balance or color-balance and is modified with respect to a standard illuminant. In this work we propose to modify the illuminant of the input image with respect of the example image illuminant.

In most cases, no information is available about the spectral power distribution of the illuminant in the scene. Hence, the solution is to estimate the color of the illuminant (the white point in the scene) based on the digital image. A simple approach to address this problem is to consider a variant of the “grey world assumption” [9], which assumes that the mean value of a natural image tends to be a greyish color corresponding to the color of the illuminant. Formally, given an input image  $I$  and an example image  $E$ , the goal is to modify  $I$  so as to adapt its illuminant to the estimated illuminant of  $E$ . For that, we propose the following example-based CAT algorithm:

1. Estimate the white point (illuminant) of image  $E$ . For a given value  $t$  (we set  $t = 0.3$ , more discussion on [9]), the white point of an image is defined as the mean color value of all pixels such that

$$\frac{|a^*| + |b^*|}{L^*} < t, \quad (1)$$


where  $a^*$ ,  $b^*$  and  $L^*$  denote the pixel coordinates in the CIELAB color space.

2. Estimate similarly the white point of image  $I$ .
3. Perform the chromatic adaptation transform (CAT) on  $I$  to adapt its white point to the white point of  $E$ . This transformation is described below.
4. Repeat Steps 2 and 3 until (a) the maximum number of iterations has been reached or; (b) the  $I$  white point has not changed from the previous iteration.
5. Return image  $I'$  which has the same geometry of  $I$  but with colors adapted to the illuminant of  $E$ .

Now, let us describe the CAT transform (Step 3). Let  $(L_I, M_I, S_I)$  and  $(L_E, M_E, S_E)$  denote respectively the estimated white points in  $LMS$  color space. Then, the CAT linear transform is defined as:

$$M_A^{-1} \cdot \text{diag}(L_E/L_I, M_E/M_I, S_E/S_I) \cdot M_A, \quad (2)$$

where  $M_A$  is a CAT matrix<sup>1</sup> that transforms from XYZ to LMS cone space.



**Fig. 1.** Illustration of the example-guided chromatic adaptation transform (CAT), where the illuminant of an example image is estimated and transferred to another image. Left column: input image and cropped image around the face. Middle column: result of the example-guided CAT with a warm light estimated from the example drawing on the top. Right column: result of the example-guided CAT with a cold light estimated from the example image on the top. In both cases, the light cast has been estimated and transferred to the input image, specially visible on the face. Images are best viewed on the electronic version.



This transformation rescales the color values of  $I$  based on the ratio of input and example white points so that the colors of  $I$  appear to have the same illuminant of  $E$ .

The algorithm is performed iteratively, hence the user can control the desired degree of adaptation to the example image according to the maximum number of iterations parameter. Experimentally, it was assessed that no more than 30 iterations are needed for an acceptable illuminant adaptation, limiting the risk of over-adaptation.

Figure 1 shows a result of the example-based CAT, where an input image (left) is corrected so as to match the warm cast of an example image (middle column) or the cold cast of another example (right column).

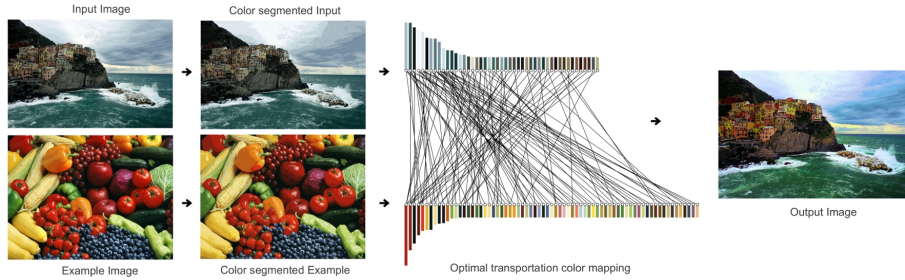
Notice that the example-based CAT depicted here can be used either as a pre-processing step before applying chroma color transfer, or as a standalone technique to perform a smooth color transfer accounting only for the illuminant of the example image.

### 3.2 Color Chroma Transfer

The intuition of our color chroma transfer method is to use optimal transportation as an efficient tool to map two color distributions approximated by their palettes, regardless of the number of colors in each set. In order to define the color chroma transfer, we propose to use a compact and robust description of the image by its set of meaningful color modes. In particular, we rely on a non-parametric color segmentation known as ACoPa (Automatic Color Palette) [4]. This is a non-supervised technique, so the user does not need to specify the number of color modes in advance, as they are automatically extracted from the color histogram based on meaningful (*a contrario*) peaks. After extracting the set of modes from the input and example images, the color transfer based on the optimal mapping between these two sets of modes is performed (see Fig. 2).

More precisely, given an input image  $I$  and an example image  $E$  with its set of meaningful modes  $P$  (and  $Q$  respectively), the mode mapping function  $f: P \rightarrow Q$  matches each input image mode with one or more example image modes. In practice, we propose a soft assignment method to compute many-to-many matches that minimizes the transportation cost between the set of modes  $P$  and  $Q$ . An effective solution for this problem comes from the Monge-Kantorovich theory of optimal transportation. Optimal transportation is a well-grounded and solid mathematical field that proves and characterizes the existence of optimal solutions for the transportation problem, minimizing the total displacement cost between two probability distributions. This displacement cost is also known as Wasserstein distance or Earth Mover’s Distance (EMD) [22]. Let  $P = \{\mathbf{p}_i\}_{i \in [1,m]}$  and  $Q = \{\mathbf{q}_j\}_{j \in [1,n]}$  be the input and example signatures with  $m$  and  $n$  modes

<sup>1</sup> Many CAT matrices exist in literature, such as CAT02, Bradford, CMCCAT2000, Sharp, *etc.* The state-of-the-art CAT02 transformation matrix [14] is used in our work.



**Fig. 2.** Overview of the proposed example-guided color transfer methodology. After extraction of the meaningful color palettes [4], a color mapping is estimated as the solution of the optimal transportation problem. Finally, a smooth color transform, computed as a 3D thin plate spline interpolation [1], generates an artifact-free image.

respectively, where  $\mathbf{p}_i$  and  $\mathbf{q}_j$  are the mode representatives. Each mode representative is a six-dimensional vector  $\mathbf{p}_i = (\mu_i^l, \mu_i^a, \mu_i^b, \sigma_i^l, \sigma_i^a, \sigma_i^b)$  composed of its mean and standard deviation (both defined as three-dimensional points in the CIELAB color space).

Let  $\mathbf{D} = [d_{ij}]$  be the distance matrix where  $d_{ij}$  denotes the distance between modes  $i$  and  $j$ :  $d_{ij} = \|\mu_i - \mu_j\|_2 + \|\sigma_i - \sigma_j\|_2$ , where  $\mu_i$  and  $\mu_j$  (resp.  $\sigma_i$  and  $\sigma_j$ ) are the mean (resp. standard deviation) color values of  $I$  and  $E$  over the modes  $i$  and  $j$ . Thus we aim to find  $\mathbf{F} = [f_{ij}]$ , with  $f_{ij}$  being the flow of the assignment between  $i$  and  $j$  minimizing the cost  $\sum_{i=1}^m \sum_{j=1}^n f_{ij} d_{ij}$ , subject to the four following constraints:

$$\begin{aligned}
 (1) \quad \forall j \in [1, n], \quad \sum_{i=1}^m f_{ij} &\leq \frac{1}{n}; & (2) \quad \forall i \in [1, m], \quad \sum_{j=1}^n f_{ij} &\leq \frac{1}{m}; \\
 (3) \quad \forall i \in [1, m], \forall j \in [1, n], \quad f_{ij} &\geq 0; & (4) \quad \sum_{i=1}^m \sum_{j=1}^n f_{ij} &= 1.
 \end{aligned}$$

In practice, the soft assignment matrix  $\mathbf{F} = [f_{ij}]$  is obtained using the Simplex linear programming algorithm [22]. After finding  $\mathbf{F}$ , for each input mode  $i = 1, \dots, m$ , an averaged and weighted mean is computed

$$\hat{\mu}_i^k = \frac{\sum_{j=1}^n f_{ij} \mu_j^k}{\sum_{j=1}^n f_{ij}}, \quad (3)$$

where  $k = \{a, b\}$  stands for chroma channels in the CIELAB color space. Based on the fact that the human visual system is more sensitive to differences in luminance than in chroma, we only use chroma channels in the color transformation, avoiding artifacts that would occur if the luminance channel was also used. Then, the color transfer between  $E$  and  $I$  is encoded giving the set of color

correspondences  $\mathcal{Y} = \{(L_i, \mu_i^a, \mu_i^b), (L_i, \hat{\mu}_i^a, \hat{\mu}_i^b)\}_i$ . Now, we have seen in practice that using  $\mathcal{Y}$  to apply a piecewise linear color transform to image  $I$  would create an output image with new undesired color edges at the color segment borders. Instead, we apply a *smooth* color transform to image  $I$ , computed as a 3D thin plate splines interpolation [1] using the set of color correspondences  $\mathcal{Y}$  in the RGB color space. Hence, it is guaranteed that the color transform applied to  $I$  is the best color transform in terms of optimal transportation and smoothness. Thin plate splines interpolation is only used as a final regularization to reduce edge artifacts. Note that the described method could create new colors i.e. colors that are not present in the input nor in the example image. Indeed, color statistics values  $\hat{\mu}_i^k$  are computed based on the weighted average of associated color statistics from the example image (Eq. 3). So, these averaged values can be seen as the result of an additive color mixture which is likely to create new colors. The risk of such color mixture model is to modify the input image with false or non realistic colors. However, this risk is limited thanks to the matching between the  $I$  illuminant and the  $E$  illuminant (cf. Sec. 3.1). Furthermore, the mapping can be constrained with visual or semantic priors as described in the next section.

### 3.3 Semantic Constraints on Color Transfer

The color chroma transfer described above does not need any prior knowledge of the input and example images. Nonetheless the color mode mapping can be easily adapted to take into account some semantic information. The main idea is to constrain the color transfer in such a way that modes corresponding to the same semantic components of the input and example images are matched together. Given the two images  $I$  and  $E$ , let us assume that the color modes can be separated into two classes  $P = \{\hat{P} \cup \tilde{P}\}$  (resp.  $Q = \{\hat{Q} \cup \tilde{Q}\}$ ) based on a spatial knowledge as visual attention or object segmentation. Then, two different color mappings  $g$  and  $h$  are computed as solutions to the bipartite graph matching in terms of Earth Mover’s distance such that:

$$g : \hat{P} \rightarrow \hat{Q}, \text{ and } h : \tilde{P} \rightarrow \tilde{Q}. \quad (4)$$

In order to satisfy these constraints, the optimization is split into two different transportation problems. In the following, we describe how semantic constraints as visual saliency and faces can be easily adapted to this framework.

**Saliency** The saliency map is a two-dimensional representation of conspicuity (gaze points) for every pixel in an image. In this work we use the saliency map  $S$  deduced from a coherent psychovisual space proposed in [12]. Given the saliency map  $S$  of an image, each color mode  $i$  is considered as salient if

$$\frac{1}{\#R_i} \sum_{x,y \in R_i} S(x,y) > \rho,$$

where  $R_i$  is the list of pixel coordinates  $(x, y)$  belonging to the color mode  $i$  and  $\#R_i$  is its cardinal. The parameter  $\rho$  is typically set to  $\rho = 0.7$ , meaning that at least 70% of the pixels belonging to a color mode are salient. Finally, salient modes are mapped to salient modes (and non salient modes to non salient modes), as described in Eq. (4).

**Faces** Face detection can be also easily incorporated in the color transfer method. The main objective is to ensure fidelity to skin tones and avoid unrealistic colors being assigned to faces and skin. Here, the popular face detector methodology<sup>2</sup> of Viola *et al.* [25] is used. The face detection is performed on both images  $I$  and  $E$ . Two cases of interest are considered:

- Faces are found in  $I$  and  $E$ . In this situation, we impose that the modes extracted from faces in  $I$  are mapped with the modes extracted from faces in  $E$  to ensure skin tones transfer.
- Faces are found in  $I$ , but not in  $E$ . In this case, colors corresponding to face and skin are not modified to ensure skin tones fidelity.


## 4 Results

In this section we present experimental results that illustrate the efficiency of our color transfer method. We strongly recommend the reader to look the figures on the digital version of the paper to appreciate the results. For all the experiments we use the example-based CAT followed by the color chroma transfer. We present four result sections. In Sec. 4.1, a comparison with five state-of-the-art color transfer methods is performed, and an objective assessment of the transformation consistency is proposed. Then, a comparison to a state-of-the-art local color transfer method is presented in Sec. 4.2. We also show the benefit of adding additional semantic constraints in specific situations in Sec. 4.4. Finally, results on video color transfer are presented in Sec. 4.3.

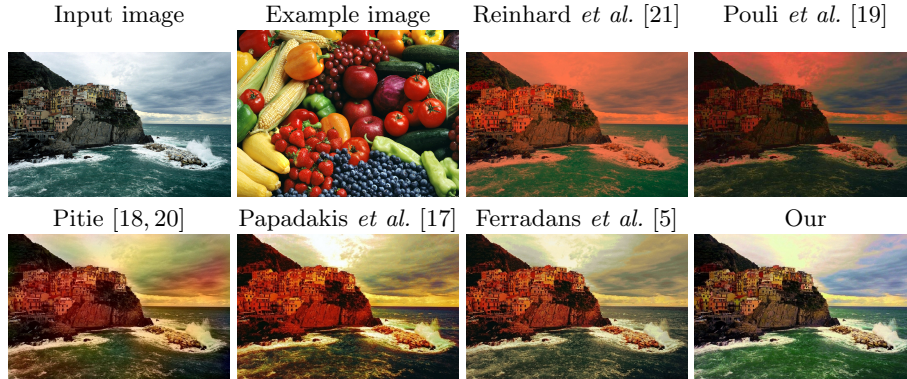
### 4.1 Evaluation of Image Color Transfer

First of all, we compare our results with five state-of-the-art global color transfer techniques from which authors have made their code available (see figures 3, 4 and 5): the seminal method of Reinhard [21], the N-PDF method from Pitié *et al.* followed by the regularization proposed by Rabin *et al.* [18, 20]; the variational method from Papadakis *et al.* [17], the histogram reshaping method described in [19]; and the regularized transportation method from Ferradans *et al.* [5]. Note that for the experiment on Fig. 3 (Scotland landscape, firstly appeared on [18]), methods [19, 17] produces a result with noise amplification artifacts, while the result of [5] has undesired edges on regions that were originally homogeneous on the input image. On the other hand, our method produces a result without


<sup>2</sup> implementation available in the OpenCV library. <http://opencv.org/>



**Fig. 3.** Results obtained by state-of-the-art color transfer techniques compared to our method for the Scotland image. For this image pair, a zoom shows color mapping artifacts. For instance, the whitish appearance of the house should be preserved, and banding/noise artifacts are visible in the sky. On the contrary, our method generates a visually plausible and artifact-free result. Images are best viewed in color and on the electronic version.



**Fig. 4.** Results obtained by state-of-the-art color transfer techniques compared to our method for the Manarola/fruits pair. While state-of-the-art techniques generate inconsistent color mapping (reddish houses and halo sky), our method leads to a visually plausible and artifact-free result. Images are best viewed in color and on the electronic version.



**Fig. 5.** Results obtained by state-of-the-art color transfer techniques compared to our method for the Burano/moscow pair. While state-of-the-art techniques generate inconsistent color mapping (sky halo, incoherent color on the water), our method leads to a visually plausible and artifact-free result. Images are best viewed in color and on the electronic version.

artifacts, similarly to the result obtained by [18, 20], with the advantage that we do not need to rely on post-processing image regularization that blurs the output image. For the challenging test pair of Fig. 4, only our method is able to adapt the low-saturated colors of the original image (Manarola, Italy on a cloudy day) to the colorful palette of the example image (fruits and vegetables) while keeping a natural and convincing result. Note that state-of-the-art methods produce results with lower color dynamics where all the houses and the rock are reddish. Finally, for Fig. 5, the state-of-the-art methods produce color aberrations on the sky and houses, leading to unnatural colors. Color transfer evaluation is not at all straightforward, since it depends on subjective aesthetic preferences. In addition, the purpose of color transfer can be diverse, for example content homogenization or artistic content colorization. However, we argue that color artifacts are not accepted as good results. In Fig.3, the halos in the sky in [18, 20], the uniform reddish color transfer for the fruits in Fig. 4 or the inconsistent and unnatural color transfer (e.g. Fig. 5) are examples of not tolerable color artifacts. We claim that non plausible results lead to the introduction of new structures in the images. Thus, the perceptual metric Structural Similarity (SSIM) [26] is used to assess the artifacts of color transfer, as already proposed in [2, 28] and more recently in [10]. Since SSIM was employed with the goal of evaluating the capability of the method to produce an artifact-free result, we computed the SSIM between the luminances of the input and the output images, not taking the color into account. In Tab.1, we compare our method with state-of-the-art in terms of artifact generation. Results show that in all cases, our method was able to transfer the color palette while preserving the geometric structure of the image.

**Table 1.** Comparison of the SSIM measure [26] between input and output images for different color transfer methods, corresponding to figures 3, 4 and 5. A SSIM value of 1 denotes that no artifacts have been generated after color transfer. Our method creates no artifacts compared to other techniques.

	Reinhard [21]	Pouli [19]	Pitié+Rabin [18, 20]	Papadakis [17]	Ferradans [5]	Our
Fig. 3	0.98	0.87	0.96	0.86	0.68	<b>0.99</b>
Fig. 4	0.84	0.56	0.94	0.91	0.87	<b>0.98</b>
Fig. 5	0.83	0.76	0.91	0.78	0.80	<b>0.98</b>

## 4.2 Comparison to Local Patch-based Color Transfer

In Fig. 6, we compare our results with the method of HaCohen *et al.* [7], which performs color transfer with a transformation based on non-rigid patch correspondences between the input and the example images. Both methods lead to reasonable results, and the objective comparison of methods is difficult. Although [7] has better recovered the specularity of the dress, our result is less saturated (arguably more natural) on the skin, on the hair and on the background. The methods were compared on a larger set of images (results are visible in the supplementary material) and both methods produce comparable results. However, while [7] assumes that the scenes are visually similar which is a very restrictive hypothesis, our framework is generic and can be used without any *a priori* of the scene. In particular, the method in [7] is ineffective with the images of section 4.1. To sum up, our method is suitable for all types of images and in the case in which the images are very similar, our method is as good as the state-of-the-art method specifically tailored for this specific case.

## 4.3 Video Color Transfer

The color transfer method presented in this paper can also be used for example-guided video color transfer. The extension of our technique to video is straightforward. We estimate the color transformation between the example image and a key frame of the video sequence and we encode it in a LUT (Lookup Table). In fact, the LUT is the result of the 3D thin plate spline interpolation. Finally, we apply the LUT to all frames in the video sequence. Unfortunately, videos cannot be shown on this version and we invite the reviewer to look at the supplementary material where videos are visible. Results show that we obtain consistent color mappings without color flickering.

## 4.4 Constrained Color Transfer

Fig. 8 shows an experiment with a challenging test pair in terms of semantic color associations. Note that when color transfer is performed without saliency constraint, the result is semantically inconsistent. We have also tested state-of-the-art methods [18, 20, 19, 17, 5] and the color mapping was not semantically





**Fig. 6.** Results obtained with the color transfer technique of [7] compared to our method. Although [7] uses the restrictive hypothesis of spatial correspondences, we obtain a visually plausible result with a highly generic method.

correct. On the contrary, when the saliency constraint is used, the birds in the images are detected as salient and their colors are matched accordingly. Finally, in Figure 9, we illustrate two cases of color transfer constrained by face detection.

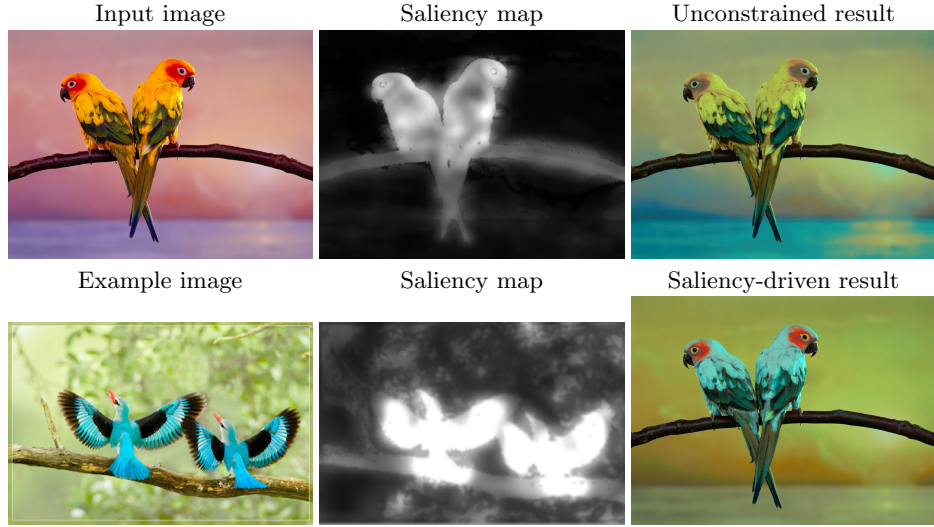
## 5 Conclusion

In this work, we have proposed a color transfer method that is based on global illuminant matching and optimal transport color transfer. All in all, the proposed color transfer method is automatic, does not require the input and example images to be visually similar and does not create visual artifacts as other state-of-the-art algorithms do. Our results present no visible artifacts since we limit changes in the luminance channel and regularize discontinuities in the color mapping through thin plate splines interpolation. We have also shown how semantic constraints can easily be considered in our framework and that our method can be applied successfully on video color transfer. The SSIM metric [26] was used to objectively assess our method compared to other techniques. Since SSIM is a metric limited to the evaluation of artifacts, future work will comprehend on one hand an extensive subjective evaluation of our method and on the other hand the assessment of color transfer quality through objective evaluation criteria.






**Fig. 7.** Video color transfer. Top row shows two possible example images. Second row, from left to right: the input frame chosen to perform the color transfer, and this frame recolored using the example images. For each example, the color transformation is computed and stored as a 3D Look-up Table (LUT). Bottom rows show some video frames before and after color transfer using the LUT computed on the reference frame. The two corresponding videos can be seen in the supplementary material.



**Fig. 8.** Illustration of saliency constraint: without the saliency constraint, the colors of the birds are not correctly transferred; while in the saliency-driven color transfer both the birds and the background are assigned to the expected colors.



**Fig. 9. Top part: Skin tones mapping.** From left to right: input and example images containing faces. The result without incorporating face constraint leads to undesirable reddish skin tones, while the result with the face constraint has mapped efficiently the skin tones. **Bottom part: Skin tones fidelity.** From left to right: input and example images. The result without incorporating face constraint leads to non plausible reddish skin tones, while the result with the face constraint has ensured skin tones fidelity.

## References

1. Bookstein, F.: Principal warps: thin-plate splines and the decomposition of deformations. *IEEE Transactions on Pattern Analysis and Machine Intelligence* **11** (1989) 567–585
2. Chiou, W.C., Chen, Y.L., Hsu, C.T.: Color transfer for complex content images based on intrinsic component. In: *IEEE Multimedia Signal Processing*. (2010) 156–161
3. Cusano, C., Gasparini, F., Schettini, R.: Color transfer using semantic image annotation. In: *IS&T/SPIE Electronic Imaging, International Society for Optics and Photonics* (2012) 82990U–82990U-8
4. Delon, J., Desolneux, A., Lisani, J.L., Petro, A.: A nonparametric approach for histogram segmentation. *IEEE Transactions on Image Processing* **16** (2007) 253–261
5. Ferradans, S., Papadakis, N., Rabin, J., Peyré, G., Aujol, J.F.: Regularized discrete optimal transport. In: *International Conference on Scale Space and Variational Methods in Computer Vision* (2013) 428–439
6. Freedman, D., Kisilev, P.: Object-to-object color transfer: Optimal flows and smp transformations. In: *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*. (2010) 287–294
7. HaCohen, Y., Shechtman, E., Goldman, D.B., Lischinski, D.: Non-rigid dense correspondence with applications for image enhancement. *ACM Trans. Graph.* **30** (2011) 70:1–70:10
8. HaCohen, Y., Shechtman, E., Goldman, D.B., Lischinski, D.: Optimizing color consistency in photo collections. *ACM Trans. Graph.* **32** (2013) 1–10
9. Huo, J.y., Chang, Y.l., Wang, J., Wei, X.x.: Robust automatic white balance algorithm using gray color points in images. *IEEE Trans. on Consum. Electron.* **52** (2006) 541–546
10. Hwang, Y., Lee, J.Y., So Kweon, I., Joo Kim, S.: Color transfer using probabilistic moving least squares. In: *Proceedings of the IEEE Computer Society Conference on Computer Vision and Pattern Recognition (CVPR)*. (2014) 3342–3349
11. Jiang, J., Gu, J.: An exemplar-based method for automatic visual editing and retouching of fine art reproduction. In: *Color and Imaging Conference*. (2013) 85–91
12. Le Meur, O., Le Callet, P., Barba, D., Thoreau, D.: A coherent computational approach to model bottom-up visual attention. *IEEE Pattern Analysis and Machine Intelligence* **28** (2006) 802–817
13. Lindner, A., Shaji, A., Bonnier, N., Süssstrunk, S.: Joint Statistical Analysis of Images and Keywords with Applications in Semantic Image Enhancement. In: *Proceedings of the 20th ACM international conference on Multimedia*. (2012) 489–498
14. Moroney, N., Fairchild, M.D., Hunt, R.W.G., Li, C., Luo, M.R., Newman, T.: The CIECAM02 Color Appearance Model. In: *Color and Imaging Conference*. (2002) 23–27
15. Murray, N., Skaff, S., Marchesotti, L., Perronnin, F.: Toward automatic and flexible concept transfer. *Computers & Graphics* **36** (2012) 622–634
16. Neumann, L., Neumann, A.: Color style transfer techniques using hue, lightness and saturation histogram matching. In: *Computational Aesthetics in Graphics, Visualization and Imaging*. (2005) 111–122
17. Papadakis, N., Provenzi, E., Caselles, V.: A variational model for histogram transfer of color images. *IEEE Transactions on Image Processing* **20** (2011) 1682–1695

18. Pitié, F., Kokaram, A.C., Dahyot, R.: Automated colour grading using colour distribution transfer. *Comput. Vis. Image Underst.* **107** (2007) 123–137
19. Pouli, T., Reinhard, E.: Progressive color transfer for images of arbitrary dynamic range. *Computers and Graphics* **35** (2011) 67 – 80
20. Rabin, J., Delon, J., Gousseau, Y.: Removing artefacts from color and contrast modifications. *IEEE Transactions on Image Processing* **20** (2011) 3073–3085
21. Reinhard, E., Adhikhmin, M., Gooch, B., Shirley, P.: Color transfer between images. *IEEE Computer Graphics and Applications* **21** (2001) 34 –41
22. Rubner, Y., Tomasi, C., Guibas, L.J.: The earth mover’s distance as a metric for image retrieval. *International Journal of Computer Vision* **40** (2000) 99–121
23. Sheikh Faridul, H., Stauder, J., Kervec, J., Tremeau, A.: Approximate cross channel color mapping from sparse color correspondences. In: *The IEEE International Conference on Computer Vision (ICCV) Workshops*. (2013)
24. Tai, Y.W., Jia, J., Tang, C.K.: Local color transfer via probabilistic segmentation by expectation-maximization. In: *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*. (2005) 747–754
25. Viola, P., Jones, M.J.: Robust real-time face detection. *International journal of computer vision* **57** (2004) 137–154
26. Wang, Z., Bovik, A.C., Sheikh, H.R., Simoncelli, E.P.: Image quality assessment: From error visibility to structural similarity. *IEEE Transactions on Image Processing* **13** (2004) 600–612
27. Wu, F., Dong, W., Kong, Y., Mei, X., Paul, J.C., Zhang, X.: Content-based colour transfer. *Computer Graphics Forum* **32** (2013) 190–203
28. Xu, W., Mulligan, J.: Performance evaluation of color correction approaches for automatic multi-view image and video stitching. In: *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*. (2010) 263–270