


Color Style Transfer Techniques using Hue, Lightness and Saturation Histogram Matching

László Neumann[†] and Attila Neumann[‡]


[†]Grup de Gràfics de Girona, Universitat de Girona, and Institució Catalana de Recerca i Estudis Avançats, ICREA, Barcelona, Spain

[‡]Institute of Computer Graphics and Algorithms, Vienna University of Technology, Austria

[†]neumann@ima.udg.es, [‡]aneumann@cg.tuwien.ac.at



(a)



(b)

Figure 1: Autumn forest, (Neumann, L., Valley of Szalajka, Hungary, 1980, 6x6 cm slide) (a) before and (b) after color style transfer using the style image Fig. 10.

Abstract

We present new methods which transfer the color style of a source image into an arbitrary given target image having a different 3D color distribution. The color transfer has a high importance ensuring a wide area of applications from artistic transformation of the color atmosphere of images until different scientific visualizations using special gamut mappings.

Our technique use a permissive, or optionally strict, 3D histogram matching, similarly to the sampling of multivariable functions applying a sequential chain of conditional probability density functions. We work by order of hue, hue dependent lightness and from both dependent saturation histograms of source and target images, respectively. We apply different histogram transformations, like smoothing or contrast limitation, in order to avoid some unexpected high gradients and other artifacts. Furthermore, we use dominance suppression optionally, by applying sub-linear functions for the histograms in order to get well balanced color distributions, or an overall appearance reflecting the memory color distribution better. Forward and inverse operations on the corresponding cumulative histograms ensure a continuous mapping of perceptual attributes correlating to limited derivatives.

Sampling an appropriate fraction of the pixels and using perceptually accurate and continuous histograms with minimal size as well as other gems make this method robust and fast.

Keywords: Color Style Transfer, Style Cloning, 3D Histogram, Sequential Sampling, Perception, Color Histogram Matching, Gradient Histogram Matching, Luminance Preservation, Hue-Identity Preservation, Spectrum Zooming

1. Introduction

1.1. Image Style Transfer

It is a mathematically and aesthetically interesting problem, how to transform a color image into a specified color style, defined by the color world of a given source or 'style' image. This problem addresses two important related fields, namely *style transfer* approaches, and *histogram matching*. Latter is also called histogram specification depending on the application area.

In the introduction, we give first a short survey of the existing style transfer techniques, next we review the histogram matching method and his application areas, and finally we write about our new method.

Style modification methods, based on sample images, are introduced in a general form in the Image Analogies methods [HJO*01], [Hut04]. These methods are applied not only to colors, but also to color textures. These methods result in very impressive pictures by a given brush-style, following arbitrary textures and in other applications. It applies an image pair (A, A') , describing the analogy rule for the target image B , instead of using one single style-image A as our approach does. Wide application possibilities of analogy methods include also the color transfer, although this direction has not been studied yet. Difficulty of this approach is to have or create appropriate (A, A') analogy pairs.

A simple and fast image re-coloring method is given by Greenfield et al [GH03], by using palette color associations. The first considerably fast and robust color transfer solution has been introduced by Reinhard et al [RAGS01]. This method applies the logarithmic $l\alpha\beta$ color space introduced by Ruderman et al [RCC98]. It has three opponent, highly decorrelated color channels. Reinhard's method [RAGS01] fits average values and variances separately for each orthogonal color channel. Calibration of the parameters can be illustrated by stretching or compressing of shifted ellipsoids having axis in the 3 color channel directions. In this way we can not approximate an arbitrary shaped gamut, especially if it contains discontinuities. This method can introduce different new hues in different diagonal directions by additive mixture of opponent color channels. However, the basic ideas and some nice results of this simple method have inspired us to search an alternative approach. This paper has inspired some gray scale colorizing methods too, using the advantages of the $l\alpha\beta$ space [WAM02] [VVD*03].

1.2. Histogram Specification and Equalization

The other pillar of our paper, beyond the style transfer approaches are the histogram matching techniques, called also histogram specification in the special literature. There is a source histogram given by an image or by just a curve in the case of gray scale images. The problem is as follows: pixels of an arbitrary given image are to be transformed into another one, which has exactly the specified histogram. A special case is the prescribed constant histogram distribution, which results in the well known histogram equalization. This problem as well as the optimal histogram matching by a monotone gray level transformation, with a really correct solution, can be a deep mathematical problem even in 1D case [CW78].

There are different approaches for color images in the literature of histogram equalization. 1D histogram equalization can be also applied for color images in a global or a local adaptive way only for the luminance channel, like hue preserving methods do it by using an adaptive neighborhood [BCRV01]. Equalization of saturation is used to exceed realizable *RGB* intensities [WHM95]. A similar goal is the extension of used gamut with 'histogram explosion' [PJ94], which preserves hue using nearly the full *RGB* extent, without causing clipping for multispectral images. Other histogram equalization approach for multispectral images is given by a pro-channel equalization and graph theory based segmentation [CCI01]. The nearest approach to our goal in the equalization literature of 3D case, or case of arbitrary N dimensions, gives a nice extension of 1D histogram equalization [PNS03], using appropriate deformation of a 3D Voronoi diagram.

1.3. Color Histogram Matching

Our goal is, of course, to specify 3D color histograms by a given *style* image. A 3D histogram is a finite approach of the continuous color distribution of the gamut. Color histogram matching can be used for quite different application areas, like e.g. fast image retrieval, where images, being similar to a given input image, are searched in a large set of images. Perceptually weighted histograms are used for color-based retrieval [LP98] or a low-dimensional distance measure between color distributions [HSE*95]. Histogram matching and distances of color distributions are often applied efficiently in pattern recognition problems too. 3D histogram matching is an important technique also in the field of gamut mapping techniques [SM01].

The best existing solutions of color histogram matching [MSS02], [MS03] ensure an exact source histogram after transformation. These approaches, beyond 3D histogram matching in the space of *CIECAM2002* color appearance model, try to preserve the original colors or minimize their changes. The price of this requirement is a slow iterative method, using the minimal earth mover's distance (EMD)



Figure 2: Colorful target image. (Kitaj, R.B., *The Oak Tree*, 1991, Oil on canvas)




Figure 3: A style image in the yellow-red hue interval. (Nemcsics, A., *Europe*, 2003, Oil on canvas)

technique, which needs a preliminary color clustering before the iteration. It needs a color quantization, clustering and an *EMD* histogram difference metric for providing a transformation *LUT* between original and target histograms.

Since the applied different color bins have different sizes, a color can quickly change to the color of the next bin, result-




Figure 4: This is transformed image of Figure 2 using the color style of Figure 3. Color matching is not quite 'strict', thereby some hues occur, like some bluish areas, which are missing from the style image. However, total appearance is basically preserved. 'Permissive' matching works with smoothed histograms, resulting in somewhat wider gamut, in order to avoid gradient artifacts.

ing in undesired artifacts. In order to avoid this phenomenon, a randomization is introduced, which is similar to a dithering. This technique could be realized by a quasi-random approach as well, like in some dithering techniques, e.g. by using a threshold matrix. But, all of these randomizations can not be considered as a perceptual solution, only as a good compromise, which fulfils the mathematical requirements. In order to improve the results, as well as to decrease the unexpected and unpredictable high-gradient effects, Morovic et al [MS03] applies a local Fourier filtering technique after the above described histogram matching.

1.4. Our method

Our motivation was to introduce a simple and fast technique to 'cloning' the color world of a given picture. We wanted to avoid the coloristic artifacts of very simple approaches, which approximate the gamut only by some parameters. On the other hand, we wanted to find a way, which is able to reproduce exactly the source gamut, but has some options to easily manipulate the source and/or target gamuts to ensure special effects. We apply the basic perceptual attributes: hue, lightness and saturation, instead using opponent color channels.

The introduced method is similar to the random sampling

of multivariable functions. Namely, it is a sequential chain of sampling by the variables, using conditional probability density functions, corresponding to some pre-tabulated cumulative distribution functions. The order of color coordinates, or perceptual attributes influences the style transfer, except in deterministic *strict* case with a very high resolution 3D histogram. Other important fact is that the hue is a circular or angularly periodic variable, while other channels, like lightness, saturation, or the coordinates of the opponent channels are not circular variables.

We can use arbitrary color spaces. However, nearly uniform spaces save memory by reducing the sufficient size of the histogram. There can not be greater steps than the perceptually noticeable ones, in case of 'strict' or exact matching. Using the original 24 *bpp* non gamma corrected *RGB* images in the non-perceptual approach, an array of size 2^{24} i.e. at about 17 million data would be required, which is not an issue nowadays. But, instead of *RGB* channels, we can use the *CIE Luv* or the different generations of *CIE Lab* systems, or *CIECAM02* system [MS03], which is used in the color appearance literature, and his orthogonal version with the *Jab* coordinates. These systems require some millions data for describing the 3D histogram.

It looks to be a better solution to use a cylindrical system of the perceptual attributes of the *Hue*, the *Lightness*, and the *Chroma*. Using this approach, we can construct a kind of hue-preserving histogram matching, where equal hues remain equal after the transformation. However, this method is somewhat more memory consuming, because the cylindrical arrangement uses unnecessarily small hue steps near to the neutral axis in order to ensure an acceptable maximal step at the highest radius representing the most saturated colors.

Preserving original hues, which is a quite different problem, is also used in the gamut mapping. We ensure another kind of hue preservation, which is a kind of invariance. Nearly all of the original hues will be more or less changed after our style transformation, but all of the colors, having the same hue on the original image, will have the same hue after the matching step too. An optimal angular rotation in the cylindrical system, i.e. the best hue matching can be an additional first automatic or interactive optimization step before the real 3D matching.

Luminance has a special importance in the human vision and recognition. E.g. the pseudo-colored gray-scale images look significantly better acceptable or more realistic, if the original luminance is preserved. Other perceptually important attributes are the local color changes, simultaneous contrast effects, which are characterized by the image gradient values at different spatial frequencies according to the spatial vision models. We can ensure a monotone and contrast limited mapping of luminance during the histogram matching, and thereby a limited change of gradients, but we can not control and preserve deeper structure of 3D gradient field by using only pixel-wise histogram matching ap-

proaches. An interesting solution is to use multiscale histograms [BD04]. Another promising generalization of the method is applying histogram matching to the gradient fields of luminance and/or saturation.

2. One Dimensional Histogram Matching

2.1. Principle of Matching

This simple case is ideal to demonstrate the basic techniques and problems. Let us be given two grayscale images with the same size. Pixels can be characterized by single scalar luminance values; therefore we have 2 images, described by 2 series of luminance values. The first image is to be modified, while the second image defines the *style* by its luminance distribution.

The core of the *strict* version of histogram matching corresponds to a substitution of the values of the first ordered luminance series by the second ordered ones. Of course their luminance values will not be explicitly represented in the arrays, since histograms are used just as auxiliary structures with certain resolutions, i.e. realizing some simplifications. Transformations use cumulated histograms, i.e. integrals, where the simplified histograms form non-negative grade function, therefore their integrals consist of monotone non-decreasing linear intervals. Fig. 5 (a) with Fig 5 (b) illustrate such cumulated histogram functions.




Figure 5: (a) Matching the domain of the histogram function of the original image to $[0,1]$ by its normalized cumulated histogram function. (b) Matching the interval $[0,1]$ to the domain of the histogram function of the style image by its normalized cumulated histogram function.

Fig 5 (left) illustrates matching the domain of the histogram of the original image with the dependent variable, which is identical to the $[0, 1]$ interval due to a normalization, and fig 5 (right) illustrates matching $[0, 1]$ with the domain of histogram of the *style* image, i.e. which defines a new style for the original image. Finally, the original histogram on fig 5 (left) will become identical to the histogram of Fig 5 (right).

This simple replacement ensures that the transformation is defined everywhere as well as monotony of matching T , i.e. if $i \geq j$, then $T(i) \geq T(j)$. In case of the second function is the identity function, matching realizes histogram equalization. It corresponds to the case, when luminance series

of the style image form a constant histogram, i.e. the style image has an equalized histogram.

This T transformation corresponds to a function from the original histogram domain to the other

$$x \rightarrow T(x) = {}^{(2)}F^{-1}({}^{(1)}F(x)) \quad (1)$$

Where ${}^{(2)}F$ denotes the cumulated histogram function of the style image, and ${}^{(2)}F^{-1}$ is its inverse function, while ${}^{(1)}F$ belongs to the original one.

However, ${}^{(2)}F$ has no inverse in some practical cases, when its histogram has 0 value within an interval, which can occur in non simplified theoretical infinite case too. This case is illustrated also on fig 5 (right). Considering this case, the following general form can be used in any case:

$$x \rightarrow T(x) = \min_t ({}^{(1)}F(t) \geq {}^{(2)}F(x)) \quad (2)$$

This T transformation can be applied not only for the original image, but also for any images. It corresponds to the earlier mentioned *image analogy*, where a T transformation, bringing I_1 to I_2 , is applied for I_3 , i.e. $T(I_1) = I_2$, and $T(I_3)$ is wanted. Note, that we can extend our 3D histogram matching method in the same way too.

2.2. Lack of Spatial Coherences

This histogram matching method ignores any spatial coherences, gradients, neighborhoods and other, topological characteristics. Although, global structure, edges and a lot of other image descriptors will be saved automatically, but magnitude of changes can not be regulated or predicted perfectly at this kind of transformation. This disadvantage follows just from the poor control of spatial correlations.

A contrast limited histogram equalization or smoothing of histograms can be applied in order to limit some too large gradients, but as an indirect consequence, some details disappear in some parts of the image. However, these simple tricks help often. Our solution uses similar gems as well.

2.3. Complicated Exact Solution of Optimal Matching

1D problem looks very simple. However, in a typical case, having two different sizes of pictures with N_1 and N_2 pixels, an exact solution would be complicated even though. Problem is as follows. There is a given target picture and a style image, a transformed image has to be found so that the sum of absolute errors between histogram of the transformed image and that of the style image is minimal. This problem is equivalent to placing N_1 linearly ordered objects of different sizes one by one into N_2 linearly ordered boxes of assorted sizes, such that the accumulated error of space underpacked or overpacked in the boxes is minimized; the placement function is monotonic, which ensures a polynomial time solution to this problem. A tree search algorithm for optimal histogram matching is presented by Shi-Kuo at al

[CW78], which has time complexity of $O(N_1 \times N_2)$. Furthermore, if the monotone property is dropped, then the problem becomes NP-complete, even if it is restricted to $N_2 = 2$.

The aforementioned algorithm demonstrates the mathematical deepness of this problem. Nonetheless, our method yields a fast and practical solution for it.

3. 3D Histogram Matching

3.1. The 3D Histogram

After selecting a 3D color space, an image generates a 3D histogram in the same way as it happens at the grayscale images' 1D histograms:

$$f(x) = |\{pixel : value(pixel) \in I\}| \quad \text{if } x \in I \quad (3)$$

where I denotes a 'bin' of the histogram, which might mean also just a point of the domain, but practically I forms an interval. Therefore f is an interval wise constant function. In 3D case, I corresponds to a *brick*, i.e. a direct product of 1D intervals. In 1D case p probability density function (*pdf*) with its F cumulated distribution function (*cdf*) are defined as

$$p(x) = \frac{f(x)}{\int f(t)dt} \quad \text{and} \quad F(x) = \int_{t \leq x} p(t)dt \quad (4)$$

The same definition can be applied in a 3D space, remarked that any integral concerns to the 3D space, and $t \in x$ means $t_i \in x_i$ at each dimension ($i = 1, 2, 3$). F is monotone too, i.e. monotone at any axis aligned lines of the domain.

The problem in 3D case is similar to the 1D one: which T transformation of the original pixels' color shall be used, that the histogram of the transformed image shall be exactly or approximately identical to the *style* image's one. 3D histogram equalization is a special case of this problem.

Despite of the simple domain of the 1D grayscale, which forms always an interval, shape of a 3D gamut can alter from a brick. For the sake of a simpler handling but not restricting its generality, we can complete the domain of the histogram to a brick, defining f for 0 at the new points. In principle, final result of our method will stay within the real gamut, i.e. it stays interpretable, while the occasional approximation errors, depending on the accuracy of the finite histogram, can be treated by some usual clipping method.

Our approach is essentially different from the existing other approaches: it is not a pixel-wise reordering of the style image and also not an analytical 3D histogram stretching, but it aims a correct 3D histogram matching, also with possibilities of its simplification and correction. Therefore we can obtain a fast, robust, flexible and simple algorithms, solving even such problems as matching an image containing just a couple of colors with a color rich one.

3.2. Sequential Sampling in the 3D space

Unfortunately, inversion in formula (1) or (2) cannot be applied for multidimensional functions. In this paper we deduce the problem of matching multidimensional cdfs to a sequence of matching 1dimensional ones, as it has been defined in formula (1) or (2). It is similar to the sequential sampling technique, but its purpose is slightly different.

Firstly, we take an order of the dimensions. Then, considering a certain value of a given dimension, the higher order dimensions have to be fixed, whilst the values of the lower order dimensions are to be integrated. In this way the following functions are defined by the original 3-variable function f , i.e. the histogram function:

$$\begin{aligned} f^1(x) &= \iint_{z,y} f(x,y,z) \, dzdy \\ f_x^2(y) &= \int_z f(x,y,z) \, dz \\ f_{x,y}^3(z) &= f(x,y,z) \end{aligned} \quad (5)$$

f^1 with f_x^2 look like a comb with its bars, as well as f_x^2 with $f_{x,y}^3$ at each given value of x . Now, each afore defined function is 1 dimensional, therefore 1D versions of formulae (3) and (4) can be applied for them, obtaining p^1 , p_x^2 and $p_{x,y}^3$ pdfs, and F^1 , F_x^2 and $F_{x,y}^3$ cdfs.

As we have mentioned, 3D version of formula (3) results in p and F 3D functions, since our goal is to make a matching on F 3D cdfs belonging to the *original* and the *style* images. 3D function p is a normalized histogram, showing probabilities of selection from its domain. However there is an interesting connection between the p pdf and the afore introduced p^1 , p_x^2 and $p_{x,y}^3$ pdfs. Let us recognize that

$$\begin{aligned} p^1(a) &= P_p(x=a) \\ p_a^2(b) &= P_p(y=b \mid x=a) \\ p_{a,b}^3(c) &= P_p(z=c \mid y=b \ \& \ x=a) \end{aligned} \quad (6)$$

Where P_p denotes probability, relying on the probabilities of 'lying into a certain bin', defined by pdf p , and at the right side of the second and third rows conditional probabilities are written. Considering connection between conditional probabilities, we get

$$\begin{aligned} p^1(a) \cdot p_a^2(b) \cdot p_{a,b}^3(c) &= \\ P_p(x=a) \cdot P_p(y=b \mid x=a) \cdot P_p(z=c \mid y=b \ \& \ x=a) &= \\ = P_p(x=a \ \& \ y=b \ \& \ z=c) &= \\ = p(a,b,c) \end{aligned} \quad (7)$$

$p^1(a) \cdot p_a^2(b) \cdot p_{a,b}^3(c) = p(a,b,c)$ makes meaning of p^i and use of F^i better understandable. Let us consider now the family of $(^1)F^1$, $(^1)F_x^2$, $(^1)F_{x,y}^3$, functions belonging to the original image, and $(^2)F^1$, $(^2)F_x^2$, $(^2)F_{x,y}^3$, are belonging to the style image, as it has been in the 1D case too. Transformation T with $(x',y',z') = T(x,y,z)$ is defined by applying

equation (2) sequentially:

$$\begin{aligned} x' &= \min_t ((^1)F^1(t) \geq (^2)F^1(x)) \\ y' &= \min_t ((^1)F_x^2(t) \geq (^2)F_x^2(y)) \\ z' &= \min_t ((^1)F_{x',y'}^3(t) \geq (^2)F_{x',y'}^3(z)) \end{aligned} \quad (8)$$

This sequential definition results in an appropriate matching transformation. It corresponds to the sequence of conditional probabilities in eq. (7), and it works dimension wise in a 1D way.

As it can be seen this method is simple and easily modifiable, does not need human interactions, even though controllable by introducing parameters, as it will be presented in this section. Its simplicity is relying on its deduction to the 1D matching problem.

3.3. Finite Histogram Resolution

The above explained method can be applied in its original form, named *strict* version, which results in a theoretical matching of two 3D functions. Note, that even this version gives a solution for the problem of very different histograms, since transformation T always exists and it is unambiguous.

The algorithm must work practically with finite arrays and some interpolation between the selected values of the domain, which is realized in our case by integrating functions, therefore it is realized by *trilinear approximations*. This means direct approximation on the first function family, and an indirect approximation on the second one. In addition to this practical simplification, it also realizes a simple smoothing, which is useful in case of very different shapes of histograms.

The following problem can be controlled also by selecting appropriate sizes. There can occur such values at the second, i.e. inverted histogram, which are approximated occasionally by such values, the subspace of which is empty, therefore its subordinate function is even not defined. This problem can be solved technically by introducing linear i.e. equalized functions as *default for these undefined cases*, but the under-represented subspaces mean still problem. This can be controlled e.g. by properly selected sizes, realizing smoothing as mentioned above.

We touched the problem of the size of histogram already in the Introduction. We have also seen above, that a low resolution or small histogram is an implicitly realized filtering. It also means a *storage reduction*.

In any cases, when an image contains only low saturations, a given range of luminance or just an interval of hues, we can found the *minimal box* of these coordinates. It is enough to build the histogram in the minimal boxes for source as well as for the target image, which restriction is proven memory saving or more accurate too.

4. Generalizations of the Basic Method

4.1. Histogram Smoothing by Convolution

We introduce two completion of the original algorithm, named also '*strict method*', comparing it to its modified versions. The first basic modification is a filtering or smoothing by an arbitrarily predefined filter profile, being used as a convolution kernel. 3 one dimensional filters can be applied sequentially, resulting in their Descartes product, or an arbitrary real 3D kernel can be defined as well, e.g. a Gaussian one with an ellipsoid.

Computational cost of this operation is a multiplication of the histogram size with the filter size, what needs a careful selection of them. In fact this operation is combined with the approximation filter, defined by the sizes of the histogram, and there is no need to use a much exact approximation if a large filter is applied afterwards. Therefore well balanced sizes make this '*smooth version*' of the method also efficient.

4.2. Histogram Suppression

The second modification aims to modify importance of the size of the color patches on an image. It is called '*dominance suppression*', corresponding to importance of sizes, since size and its importance are not equal. Smaller colorful areas get often higher importance, while large unsaturated areas get less importance. This modification is performed on the normalized histogram function, instead of on its cdfs, and its combination with the direct or indirect filtering makes its effect more realistic by extending its modification onto groups of nearby colors. It can be realized by a simple function, controlled by parameter, expressing its effect. The function tends to the identity function at small values and to a constant at great values. Increasing the parameter, latter effect starts to be realized even for small values.

$$x_{new} = \begin{cases} 1 - e^{-p \cdot x} & \text{if } p \neq 0 \\ x & \text{if } p = 0 \end{cases} \quad (9)$$

Dominance suppression can be realized also by any operation on the *pdfs*. One such an operation is limitation of their maximum value, which corresponds to some contrast limitation of the image. It is realized practically by a cutting operation and a renormalization, or their successive iteration in order to define its absolute maximum value. A *pdf* is the derivative of its corresponding *cdf*, therefore limit of the *pdf* means limit of derivative of the *cdf*, and results in a more equalized histogram.

We have to mention that convolution and dominance suppression operations can be executed by both orders, but their results will be different.

4.3. Saving and Changing Perceptual Attributes

We cannot save the original hue or luminance in the 3D color histogram matching. Our method preserves only the *hue* – *identity*, and the *monotony of the luminance and chroma*.



(a)



(b)



(c)

Figure 6: (a) Martian landscape with the rock 'Wopmay'. Hue differences are hardly visible. (b) Martian landscape after a histogram matching, which has not changed the original luminance, but hues and chroma values. This is an example of 'spectrum zooming'. Picture can be easier understood and different minerals, materials can be easier distinguished. (c) Style image 'Kristóf', having been applied to histogram matching for target image (a).

Ignoring this criteria, we can obtain interesting artistic images. We can create e.g. a negative image using a negative luminance, or better said a constant minus the original lightness. Similarly, complementary hues can be taken instead of the original ones. A somewhat less trivial variable, the relative lightness occurs in the *TRIA mapping*, which preserves the overall appearance by changing the original luminance [MR01]. For a *CRT*, colors are linearly mixed in a triangle from the black, white and the most saturated 'limit-color', which depends on the hue. Luminance of the limit-color can strongly change depending on the hue, see e.g. cases of yellow and blue colors. We can get colors outside of display or printer gamut by applying different histogram transformations. This problem will never occur using relative luminance, namely the relative position of color is preserved in the new triangle of the new hue. There are a lot of other and exotic possibilities using different transfer functions of hue, luminance and chroma, or above described ratios of white, black and limit-color.

A practical application is obtained by preserving the original hue exactly, when problem is reduced to a *2D* histogram matching. Similarly, we can preserve the original luminance, while the other two attributes will be matched. Luminance is highly important in human vision. Its unchanged value ensures a deep perceptual invariance, as it can be seen on figures 6 (a) and 6 (b). The style image was fig 6 (c). We can fix 2 color coordinates as well, that reduces the problem to a classical *1D* histogram matching applying on a certain perceptual attribute. We can enhance appearance and visibility of some details by applying e.g. luminance or saturation matching, or equalization which is a special case of matching.

Every histogram matching method, based on single pixel-wise functions, will ignore spatial coherences, or just reduce some artifacts applying histogram transformations. However, there are some possibilities to use variables containing spatial information too. We can match histograms of absolute values of the gradient (G_L) of luminance, while hue (H) and chroma (C) are used in a classical way, and the direction of the gradient vectors stays unchanged. After histogram matching, we get a gradient map, and the best approximate image can be reconstructed from it by different techniques. An efficient numerical method is the multigrid Poisson solver [MAW02]. This kind of style cloning is a new approach, which grasps deeper and reflects better the total image appearance, than only single-pixel histogram techniques.

Similarly chroma gradients (G_C) can be cloned too. There are some possible combinations: (H, G_L, C) , (H, L, G_C) , (H, G_L, G_C) or *2D* matching by fixing hue in the previous cases. Luminance and chroma gradients, luminance and color edges, only-luminance and only-color edges have a high importance in automatic recognition of paintings and photographs [LC03]. A painting contains several only-color

changes in a typical case. We can simulate some important attributes of pictorial style or photographic style by using the above described kind of matching in a different manner to the image analogies. We have described here just the basic ideas, problem of matching gradients and hybrid attributes requires further investigations.

5. Special Applications

5.1. Histogram Equalization

We have studied two important special areas. The first one is the histogram equalization. It is widely used in *1D* context. *3D histogram equalization* works for any input image in a similar way: the specified *3D* histogram has uniform distribution. Thereby we do not need an explicit source image, but we can imagine a fictive one or also generate images with the *3D* uniform distribution. The result of the equalization has often a strange appearance. The small differences of basic perceptual attributes can appear as fully different colors in hue, luminance and saturations. The technique works often as a natural noise amplifier. However, it has useful applications [PNS03] too. We illustrate *3D* histogram equalization with figures 7 and 8. This technique is a useful tool to analyze artworks, textiles and to find forgery or restored parts. The equalized image looks more realistic when fixing hues of the target image, which means a *2D* equalization for *luminance* and *chroma*.

5.2. Spectrum Zooming

Another interesting application area is the *spectrum zooming*, by our terminology. Original hues are often undistinguishable, if an image contains colors only in a small interval of hues or wavelengths of the visible spectra. The Martian landscape (fig.6 (a)) as a good example looks being nearly monochrome before the style transfer. Spectrum zooming can be used e.g. in the undersea imaging, where in deep water over some times ten meter distance is only a small bluish 'window' of visible spectra for image capturing: $430nm - 490nm$ or even shorter interval. However, multispectral photographs can be taken using narrow band filters in this small spectral interval too. We get visually quite similar bluish images after an absorption correction using the depth map [NGFN04]. Even the least and undistinguishable differences can be amplified by color style cloning into the desired color style. We can recolorize this ambient in this way, and different materials, minerals can be easily recognized on the processed image. This is a practical area of spectrum zooming.



Figure 7: Test image for color histogram equalization. (Neumann, L., Lake Baláta, Hungary, 1980, 6x6 cm slide)



Figure 8: Color histogram equalized version of Figure 7. This example is just a test. Real application area is Image Analysis, visualization of small color differences, like here some vertical black lines can be seen, due to slight error of slide scanner, which could not be visible without emphasizing them.

6. Results and further investigations

We introduced new color style cloning methods based on matching of 3D histograms of hue, luminance and saturation. The new method is fast and robust. It has *strict* and a different *permissive* versions and a lot of special and practical options and application areas.

We developed before the histogram matching techniques some extensions of Reinhard's method. Figure 9 (b) demonstrates the result of original Reinhard's [RAGS01] method applied to the very colorful target image (fig 9 (a)). The style image is a forest image containing mostly greenish color and somewhat yellow, furthermore a pastel and light bluish-lilac foggy background (fig 10). This testbed really shows the possible problems because segmentation is not applied. The referred method is equivalent with shifting and compressing of the 3 orthogonal axes of an 'ellipsoid-gamut'. We applied optimal stretching along the 6 half-axis directions in different ways, thereby there are less 'wrong' colors and also these ones have a lower saturation level on fig 9 (c). We have achieved a further improvement by applying exact 1D histogram matching on the 6 half-axis instead of simple stretching on them (fig 9 (d)). Colors are further improved by this kind of method, but undesired gradient effects occurred as well, in particular saturated blue colors on the boat. Reinhard et al [RAGS01] use image segmentation to obtain satisfying results. Segmentation is a step which deeply exploits spatial coherences of the image.

The color histogram matching technique introduced in this paper is a significantly different approach. It maps the

arbitrary source gamut to the arbitrary target one, while colors with same hues of target image will have the same hues after the transformation. However, the 'strict' version results in some cases in unwanted gradient or noise effects. It reproduces e.g. the original color histogram of the forest image exactly (fig 11), but the unexpected and undesired noise-like parts occurs on this hard test image-pair 'perfectly'. To avoid or reduce the unwanted gradient effects we introduced the *permissive version* of histogram matching applying a convolution-like smoothing of the histograms or a contrast limitation of the cumulated distribution functions. The method of *histogram suppression* manipulates the source histogram to suppress some dominant but unwanted or boring colors filling large areas, or to emphasize important characteristic, saturated, but under-represented colors. Fig. 12 (c), Fig. 12 (a), and Fig. 12 (b) illustrate this technique.

The introduced technique has some really special application areas. The color histogram equalization has rather scientific than aesthetic importance. It is a color visualization technique making the originally small color differences visible. It makes possible the detection of the restored parts of images or forgeries. Other interesting application area is the spectrum zooming for images with nearly monochromatic appearance (fig. 6 (a)). In fig. 6 (b) we show beyond this technique the luminance preservation option. The real appli-




Figure 9: (a) Target image 'ski'. This well known colorful test image with the style image 'forest' (Fig. 10) is an ideal testbed of histogram matching. (b) The ski-forest transfer by the original method of Reinhard et al. [RAGS01]. Some yellow, blue, brown and lilac colors are evidently missing from the source image, e.g. colors of boot or helmet. This transfer, as well as the further ones, do not apply image segmentations. They are global transformations. (c) Ski-forest transfer with 6 half axis compression. The result is somewhat better, disturbing colors are less saturated, 'diagonal-color' problem occurs less. (d) 1D histogram matching is applied on all of 6 ellipsis half-axes. Disturbing 'new colors' are quite unsaturated. Boots looks nearly gray, they were bluish using the two previous transfers. Unfortunately, they conserve the blue color on a small area. It looks like an unexpected and not predictable gradient effect. All of the half axis transformations can not overstep some limitations, which follows the nature of this simple method.



Figure 10: Style image 'forest' contains mostly greenish colors, some yellow and brown areas, furthermore bright and unsaturated bluish-lilac foggy background.



Figure 11: 3D histogram of this image and Fig. 10 are exactly the same, due to the strict version of the histogram matching. Comparing with methods illustrated by Fig. 9, the color world has dramatically changed as it has been expected. However, ignoring spatial coherences, the same histogram in itself looks to be not enough for a true style cloning. Problems are caused specially by unpredictable noise and gradient effects. The result can be improved using image segmentation and 3-dimensional smoothing of the color histogram.

cation area of the spectrum zooming is the undersea imaging or e.g. the near infrared (NIR) range.

Although we presented fast new methods with different options, the real 'image appearance transfer' is a perceptually deep and complex problem, a real challenge, which is only partially solved. A common problem of each pixel-wise matching method is the lack of spatial information, thereby local changes, like gradient effects of the transformed target image can not be predicted. Different image styles, like photographs and paintings, have characteristically different luminance and chroma gradient histograms. One of the possible new approaches is the histogram matching of the gradients. This approach is new and it needs further investigations, which will be reported in a joint paper.

Acknowledgements

This work was partially supported by the Spanish Government by project number TIC2004-07451-C03-01 and by the European Union within the scope of the RealReflect project IST-2001-34744, "Realtime Visualization of Complex Reflectance Behaviour in Virtual Prototyping".

References

- [BCRV01] BUZULOIU V., CIUC M., RANGAYAN R. M., VERTAN C.: Adaptive-neighborhood histogram equalization of color images. *Journal of Electronic Imaging* 10 (April 2001), 445–459.
- [BD04] BAE S., DURAND F.: *Statistical Analysis and Transfer of Pictorial Style*. Tech. rep., MIT Computer Science and Artificial Intelligence Laboratory, USA, <http://sow.lcs.mit.edu/2004/proceedings/Bae.pdf>, 2004.
- [CCI01] CHITWONG S., CHEEVASUVIT F., INTAJAG S.: Color image enhancement based on local area histogram equalization. In *Fourth International Conference on Signal Processing, Proceedings of SPIE* (2001).
- [CW78] CHANG S.-K., WONG Y.-W.: Optimal histogram matching by monotone gray level transformation. *Communications of the ACM* 21 (October 1978), 835–840.
- [GH03] GREENFIELD G. R., HOUSE D. H.: Image Recoloring Induced by Palette Color Associations. *Journal of WSCG 11* (2003), 189–196.
- [HJO*01] HERTZMANN A., JACOBS C. E., OLIVER N., CURLESS B., SALESIN D. H.: Image Analogies. In *Proceedings of ACM SIGGRAPH 2001* (2001), pp. 327–340.
- [HSE*95] HAFNER J., SAWHNEY H. S., EQUITZ W., FLICKNER M., NIBLACK W.: Efficient Color Histogram Indexing for Quadratic Form Distance Functions. *IEEE Transactions on Pattern Analysis and Machine Intelligence* 17 (July 1995), 729–736.



(a)



(b)



(c)

Figure 12: (a) A style image containing a dominant color on a larger area. We are interested often just in transferring 'important colors', similarly to our 'memory image', which contains colors non area proportionally. (b) Fig. (a) cloned by histogram suppression, which operation suppressed the dominant beige background color of style image (c). Additionally the original luminance has been preserved.

- [Hut04] HUTTUNEN J.: Image Analogies. In *Seminar on Computer Graphics, Helsinki Spring 2004* (2004), pp. 111–500.
- [LC03] LEYKIN A., CUTZU F.: Differences of Edge Properties in Photographs and Paintings. In *Proceedings of ICIP 2003* (2003).
- [LP98] LU G., PHILLIPS J.: Using perceptually weighted histograms for colour-based image retrieval. In *Fourth International Conference on Signal Processing* (October 1998), pp. 12–16.
- [MAW02] MILLAR C., ASENOV A., WATLING J. R.: Excessive Over Relaxation Method for Multigrid Poisson Solvers. *Journal of Computational Electronics 1* (October 2002), 341–345.
- [MR01] MOROVIC J., RONNIER L. M.: Evaluating gamut mapping algorithms for universal applicability. *Color Research and Application 26* (2001), 85–102.
- [MS03] MOROVIC J., SUN P.: Accurate 3D image colour histogram transformation. *Pattern Recognition Letters 24* (2003), 1725–1735.
- [MSS02] MOROVIC J., SHAW J., SUN P.: A fast, non-iterative and exact histogram matching algorithm. *Pattern Recognition Letters 23* (2002), 127–135.
- [NGFN04] NEUMANN L., GARCIA R., FERRER J., NEUMANN A.: *High-Resolution Depth and Reflectance Map Estimation in Underwater Imaging*. Tech. rep., TR-186-2-04-14, Vienna University of Technology, December 2004.
- [PJ94] PHILLIP A. M., JEFFREY J. R.: Enhancement of multispectral images using histogram explosion. In *Proceedings on Applications of Digital Image Processing Conference XVII* (1994), pp. 132–142.
- [PNS03] PICHON E., NIETHAMMER M., SAPIRO G.: Color histogram equalization through mesh deformation. *ICIP 2* (2003), 117–120.
- [RAGS01] REINHARD E., ASHIKHMIN M., GOOCH B., SHIRLEY P.: Color Transfer between Images. *IEEE Computer Graphics and applications, special issue on Applied Perception* (2001), 34–41.
- [RCC98] RUDERMAN D. L., CRONIN T. W., CHIAO C. C.: Statistics of Cone Responses to Natural Images: Implications for Visual Coding. *J. Optical Soc. Of America 15* (1998), 2036–2045.
- [SM01] SUN P., MOROVIC J.: The influence of image histograms on cross-media colour image reproduction. In *PICS 2001* (2001), pp. 363–367.
- [VVD*03] VIERA L., VILELA R., DONASCIMENTO F., FERNANDES J. F., CARCERONI R., ARUAJO A.: Automatically choosing the source image in the coloring process of grayscale images. In *Anais do IV Workshop em Tratamento de Imagens, NPDI / DCC / ICEx / UMGF* (June 2003), pp. 77–89.
- [WAM02] WELSH T., ASHIKHMIN M., MUELLER K.: Transferring color to grayscale images. In *Proceedings of ACM SIGGRAPH 2002* (2002), pp. 277–280.
- [WHM95] WEEKS A. R., HAGUE G. E., MYLER H. R.: Histogram equalization of 24-bit color images in the color difference (C-Y) color space. *Journal of Electronic Imaging 04* (Januray 1995), 15–22.