

N-Dimensional Probability Density Function Transfer and its Application to Colour Transfer

François Pitié , Anil C. Kokaram and Rozenn Dahyot
University of Dublin, Trinity College
Dublin 2, Ireland
fpitie@mee.tcd.ie

Abstract

This article proposes an original method to estimate a continuous transformation that maps one N-dimensional distribution to another. The method is iterative, non-linear, and is shown to converge. Only 1D marginal distributions are used in the estimation process, hence involving low computation costs. As an illustration this mapping is applied to colour transfer between two images of different contents. The paper also serves as a central focal point for collecting together the research activity in this area and relating it to the important problem of Automated Colour Grading.

1. Introduction

The principle of example-based rendering is probably the simplest and most effective approach to rendering realistic images. The idea is to transfer the statistics of a target dataset (the example) to a source dataset. After transfer, the re-rendered source dataset has the same look and feel as the target one. The recent breakthrough in texture synthesis [4] is probably one of the most significant examples of this idea. By simply sampling from the neighbouring distribution of a seed example texture, it is possible to synthesise bigger pictures that look similar. The success of realistic rendering thus relies on the transfer of real data statistics.

The idea of transfer of statistics encompasses an entire range of possibilities from the simple match of the mean and variances of both datasets to the exact transfer of the whole probability density function (pdf) of the samples. This paper considers the general problem of finding a continuous mapping that transforms the source samples into a new ensemble of samples that exhibits the same pdf as the target samples. Finding a mapping for one-dimensional (1D) samples is a simple task that can be solved in the same way as grayscale picture equalisation [5]. The difficulty resides in

extending the method to higher dimensions.

There are a wide range of applications for the notion of exact transfer of pdf's for multidimensional datasets. In the transfer of colour between two images, it is known as example-based colour transfer [14]. In digital restoration [12] the idea is to recolour paintings that have been faded by smoke, dust etc. If the target pdf is uniform, the process can also be used for colour image equalisation for scientific data visualisation [13] or simply used for non-realistic rendering. In particular, a major problem in the post production industry is matching the colour between different shots possibly taken at different times in the day. This process is part of the large activity of *film grading* in which the film material is digitally manipulated to have consistent grain and colour. The term *colour grading* will be used specifically to refer to the matching of colour. Colour grading is important because shots taken at different times under natural light can have a substantially different 'feel' due to even slight changes in lighting. Currently these are fixed by experienced artists who manually tune parameters. This is a delicate task and the method presented in this paper succeeds in automating this painstaking process even when the lighting conditions have dramatically changed, as shown in figure 4.

This paper proposes a simple method that performs an exact transfer of a pdf, based on the iterative use of the one-dimensional pdf transfer. The method is guaranteed to converge to the target pdf and a mathematical proof is presented. The advantage of the method is that the pdf transfer operates in 1D which means that the overall algorithm has a linear computational complexity of $\mathcal{O}(M)$, where M is the number of samples processed. This implies also that it is possible to have more bins in the histogram and thus achieve higher accuracy in the colour mapping.

2. Related Work

Transfer of Colour Statistics. One popular method [14] matches the mean and variance of the target image to the source image. The transfer of statistics is performed separately on each channel. Since the RGB colour space is highly correlated, the transfer is done in the less correlative colour space $l\alpha\beta$ which has been proposed to account for human-perception [15]. But the method is limited to linear transformations and some example-based recolouring scenarios actually require non-linear colour mapping. Figure 1 shows exactly this problem and the method fails to transfer any useful statistics.

The problem of finding a non-linear colour mapping is addressed in [13] for colour equalisation (like for grayscale histogram equalisation). The paper proposes to deform the tessellation meshes in the colour space to fit to the 3D histogram of uniform distribution. This method can be seen as being related to warping theory which is explicitly used in [10] where the transfer of the 2D chromatic space is performed directly by using a 2D-biquadratic warping.

Without having to invoke image warping, an extension of the 1D case is to find the mapping via linear programming [11] and the popular Earth-Mover distance. The major disadvantage of the method is that 1) the mapping is not continuous and 2) pixels of same colours may be mapped to pixels of different colours, which require random selection. Furthermore the computational cost becomes intractable if a very fine clustering of the colour space is desired.

Dealing with Content Variations. One important aspect of the colour transfer problem is the change of content between the two pictures. For example the target picture may present more sky than the source picture. Since all colour transfer algorithms are sensitive to the variations of the colour cluster sizes, they risk overstretching the colour mappings and thus producing unbelievable renderings.

To deal with this issue a simple solution [14] is to manually select swatches in both pictures and thus associate colour clusters corresponding to the same content.

One automated solution is to invoke the spatial information of the images to constrain the colour mapping [7]. The extreme case is colouring grayscale pictures. Retrieving the missing chrominance channels values can be done by taking advantage of similarities between spatial neighbourhoods of the two pictures [16, 6].

Another automated method is to restrict the variability on the colour mapping. For example in [1], the pixels of both images are classified in a restricted set of basic colour categories, derived from psycho-physiological studies (red, blue, pink...). The colour transfer ensures that blue-ish pixels remain blue-ish pixels. This gives a more natural transformation but limits the range of possible colour transfers.



Figure 1. Example of Colour transfer using Reinhard [14] Colour Transfer. The transfer fails to resynthesise the colour scheme of the target image. To be successful the method would require human interaction.

By contrast our method is able to map shades of green to shades of orange in the first row of figure 4.

Novelty. The work presented in this paper addresses many of the shortcomings of previous efforts in this area. Firstly, it is computationally attractive as it uses just 1D pdf matching in an iterative scheme which is shown to converge. Secondly, the method is completely non-parametric and is very effective at matching arbitrary distributions. The paper also serves as a central focal point for collecting together the research activity in this area and relating it to the important problem of Automated Colour Grading which has not been previously specifically addressed in this community.

3. N-Dimensional pdf Transfer

Denote $f(x)$ and $g(y)$ the pdf of X and Y , the original and target N -dimensional continuous Random Variables respectively. For example in colour transfer, the samples x_i of X encapsulate the three colour components $x_i = (r_i, g_i, b_i)$. The goal is to find a continuous mapping function t that transforms f in g .

Dimension $N = 1$. This is a well known problem [5] which offers a simple solution:

$$t(x) = C_Y^{-1}(C_X(x)) \quad (1)$$

where C_X and C_Y are the cumulative pdfs of X and Y . This can be easily solved using discrete lookup tables.

Dimension $N \geq 2$. The idea is to reduce the problem from N dimensions to the 1-dimensional case. For example consider a 1D axis. The projections of the N -dimensional samples for X and Y are computed along the axis. Matching these two marginals using the previous 1D pdf matching scheme results in a 1D mapping. This mapping can be applied along the axis to transform the original N -dimensional samples. The new distribution of the transformed original

Algorithm 1 pdf transfer algorithm

- 1: **Initialisation** of the data set source \mathbf{x} and target \mathbf{y} . For example in colour transfer, $\mathbf{x}_j = (r_j, g_j, b_j)$ where r_j, g_j, b_j are the red, green and blue components of pixel number j .
 $k \leftarrow 0, \mathbf{x}^{(0)} \leftarrow \mathbf{x}$
 - 2: **repeat**
 - 3: take a rotation matrix R and rotate the samples: $\mathbf{x}_r \leftarrow R\mathbf{x}^{(k)}$ and $\mathbf{y}_r \leftarrow R\mathbf{y}$
 - 4: project the samples on all axis i to get the marginals f_i and g_i
 - 5: for each axis i , find the 1D transformation t_i that matches the marginals f_i into g_i
 - 6: remap the samples \mathbf{x}_r according to the 1D transformations. For example, a sample (x_1, \dots, x_N) is remapped into $(t_1(x_1), \dots, t_N(x_N))$, where N is the dimension of the samples.
 - 7: rotate back the samples: $\mathbf{x}^{(k+1)} \leftarrow R^{-1}\mathbf{x}_r$
 - 8: $k \leftarrow k + 1$
 - 9: **until** convergence on all marginals for every possible rotation
 - 10: The final one-to-one mapping t is given by: $\forall j, \mathbf{x}_j \mapsto t(\mathbf{x}_j) = \mathbf{x}_j^{(\infty)}$
-

samples is proved to be actually closer to the target distribution than before transformation. The main point of this paper is to observe that iterating this simple procedure for different set of axes is sufficient to eventually transform the original samples into a distribution that is identical to the target distribution.

The algorithm operates as follow. At iteration k , the samples involved are the iterated samples of $X^{(k)}$ (the transformed image) and the samples from Y (the target image). The first step of the iteration is to change the coordinate system by rotating *both* the samples of $X^{(k)}$ and the samples of Y . In a second step the samples of both distributions are projected on the new axes which gives the marginals $f_1 \dots f_N$ and $g_1 \dots g_N$. Then it is possible using equation (1) to find for each axis i the mappings $t_i : \forall j x_j \rightarrow t_i(x_j)$ that transfers the marginals from f_i to g_i respectively. The resulting transformation t maps a sample (x_1, \dots, x_N) into $t(x_1, \dots, x_N) = (t_1(x_1), \dots, t_N(x_N))$. The iteration is completed with a rotation of the samples by R^{-1} to return in the original coordinate system.

The iteration leaves the samples from g unchanged and transforms the pdf $f^{(k)}$ into $f^{(k+1)}$. The algorithm is shown to converge to $f^{(\infty)} = g$ if the operation is repeated for enough different rotations (taking random rotations is sufficient to converge). The full algorithm is presented in a separate figure on this page and is simple to implement as it requires no extra parameters.

The outline of the method presents some resemblance with an iterative algorithm of gaussianization of data presented in [9]. However this method investigates a much wider range of transformations as it proposes to match not only gaussian distributions but any kind of distribution. Moreover a mathematical proof of the method is presented, whereas in [9] the validity of the method is limited to numerical experiments.

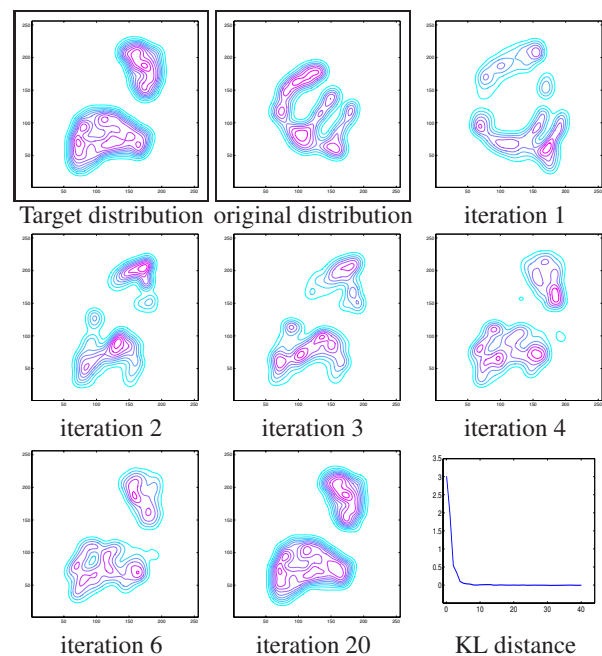


Figure 2. Example of 2D pdf transfer. Note the decrease of the measure of the Kullback-Leibler distance.

Softening the pdf Transfer. The algorithm proposed performs an exact transfer of the pdf. To account for the possible change of content in the data and thus avoid excessive stretching of the mapping functions, we introduce a way of reducing the magnitude of the stretching. The modification is done on the 1D pdf matching scheme. Instead of considering the marginals f_k and g_k , the pdf transfer will actually match $f_k^{(1/m)}$ and $g_k^{(1/m)}$, with $m \in [1; 6]$. This operation reduces the relative variations of the marginals and thus the

mapping function get closer to the identity mapping if m increases. Such a modification slows down the convergence of the algorithm but it is often required since the content of the pictures rarely match exactly.

4. Proof of the Method

Proof. Consider the operation T which describes one iteration of our algorithm (ie, rotate, transform and rotate back the samples): the target distribution g is of course stationary for T since $T(Y) = Y$ and T transforms $f^{(k)}$ in $f^{(k+1)}$ with $X^{(k+1)} = T(X^{(k)})$.

The Data Processing Theorem [8, pp. 18-22] — which is fundamental in Information Theory — states that no statistical processing of the data can increase the Kullback-Leibler distance. Thus the Kullback-Leibler distance between $f^{(k)}$ and g decreases for each iteration k , and this for every possible rotation and initial distributions.

$$D(f^{(k)} \| g) \geq D(T(f^{(k)}) \| T(g)) \quad (2)$$

$$D(f^{(k)} \| g) \geq D(f^{(k+1)} \| g) \quad (3)$$

The sequence $D(f^{(k)} \| g)$ is monotonically non-increasing non-negative and must therefore have a limit. The limit is actually 0 if the stationary distribution g is unique (see [3, pp. 33-35]). It is then sufficient to prove that g is the unique stationary distribution to show the convergence of the sequence $f^{(k)}$ to $f^{(\infty)} = g$.

This can be done by invoking a geometrical argument. Let h be another stationary distribution. By definition of the algorithm, for every rotation of the axis, the projections (or marginals) of h match the projections of g . Thus the N -dimensional Radon Transforms of h and g are identical. Since the Radon transform admits a unique inverse, we have $h = g$ and thus g is the only stationary distribution. \square

Numerical Experiences. The measure of the KL distance is performed via the kernel density approximation of the density. We implemented a kernel density estimation with variable bandwidth to account for the sparseness of samples. A clear outline of the bandwidth selection is available in [2]. The numerical Kullback-Leibler distance or relative entropy can be computed as follow:

$$\mathcal{D}_{KL}(f \| g) = \frac{1}{N} \sum_i \ln \left[\frac{\sum_j K\left(\frac{x_i - x_j}{h}\right)}{\sum_j K\left(\frac{x_i - y_j}{h}\right)} \right] \quad (4)$$

where K is the Epanechnikov kernel. As expected by the theory, the KL distance decreases with the iteration n (see figure 2).

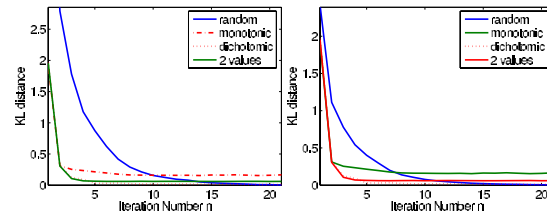


Figure 3. Convergence rate depending on the rotation selection strategy. If the random strategy (in solid blue) performs poorly in the first iterations it eventually outperforms our deterministic strategies.

Choice of the Rotation Matrices. From the proof we can state that the algorithm has to take marginals for every possible rotation, where there is an infinite number of possible rotations. The simplest way of addressing this is to spread the angle values over the whole angular spectrum and sample the angles of the rotation matrix from a uniform distribution.

The convergence rate of the random selection strategy has been tested against 3 deterministic strategies. The first strategy is a monotonic increase of the rotation angles, the second strategy is a ping-pong between two fixed rotations (0 and $\pi/4$) and the third strategy explores angles in a dichotomous way ($0, \pi/4, 0, \pi/8, \pi/4, 0, \pi/16, 2\pi/16 \dots$). The third method tries to emphasise the harmonics of the angular spectrum. Figure 3 shows that, on average over 50 experiences, the random sampling of rotation outperforms eventually the deterministic schemes. In particular the monotonic scheme seems almost to stall.

5. Application to Image Recolouring

The last two pages show some results from the colour transfer technique proposed in the paper. Figure 4 displays some examples of colour transfer. For instance the original Alpine mountain picture in (b) is used as a target colour scheme for the Scottish landscape image in (a) the first column. The result of the transfer appears in the last column. The figure presents also examples of Colour Transfer for matching lighting conditions. On the second row, the colour properties of the sunset are used to synthesise the ‘evening’ scene depicted at sunset. On the last row, the colour transfer allows correction of the change of lighting conditions induced by clouds. A unavoidable limitation of colour grading is the clipping of the colour data: saturated areas cannot be retrieved (for instance the sky on the golf image cannot be recovered). A general rule is to match pictures from higher to lower range dynamics.

The figure 5 displays examples colour restoration of

faded movies. The idea is similar to colour grading as we want to recreate different atmospheres. The target pictures used for recreating the atmosphere are on the second row.

6. Conclusion

This paper has proposed an original algorithm for transferring N -dimensional pdfs. The method is guaranteed to converge at low computation costs. We have shown the efficiency of the algorithm for colour transfer. Future works will look at its application to other areas of computer vision involving pdfs matching.

Acknowledgements

The authors would like to acknowledge the helpful discussions with Bill Collis at The Foundry[®]. This work was supported in part by HEA Project TRIP and GreenParrotPictures[®].

References

- [1] Y. Chang, K. Uchikawa, and S. Saito. Example-based color stylization based on categorical perception. In *Proceedings of the 1st Symposium on Applied perception in graphics and visualization (APGV)*, pages 91–98. ACM Press, 2004.
- [2] D. Comaniciu, V. Ramesh, and P. Meer. The Variable Bandwidth Mean Shift and Data-Driven Scale Selection. In *IEEE Int. Conf. Computer Vision (ICCV'01)*, pages 438–445, Vancouver, Canada, 2001.
- [3] T.M. Cover and J.A. Thomas. *Elements of Information Theory*. Wiley, 1991.
- [4] A. A. Efros and T. K. Leung. Texture synthesis by non-parametric sampling. In *IEEE International Conference on Computer Vision*, pages 1033–1038, Corfu, Greece, September 1999.
- [5] R. C. Gonzalez and R. E. Woods. *Digital Image Processing*. Addison Wesley, 1992.
- [6] Y. Ji, H-B. Liu, X-K Wang, and Y-Y. Tang. Color Transfer to Greyscale Images using Texture Spectrum. In *Proceedings of the Third International Conference on Machine Learning and Cybernetics*, Shanghai, 2004.
- [7] J. Jia, J. Sun, C-K. Tang, and H-Y. Shum. Bayesian correction of image intensity with spatial consideration. In *European Conference on Computer Vision (ECCV)*, 2004.
- [8] S. Kullback. *Information Theory and Statistics*. Wiley, New York, 1959.
- [9] J.-J. Lin, N. Saito, and R. A. Levine. An iterative non-linear Gaussianization algorithm for resampling dependent components. In *Proc. 2nd International Workshop on Independent Component Analysis and Blind Signal Separation*, pages 245–250, 2000.
- [10] L. Lucchese and S. K. Mitra. a new Method for Color Image Equalization. In *IEEE International Conference on Image Processing (ICIP'01)*, 2001.
- [11] J. Morovic and P-L. Sun. Accurate 3d image colour histogram transformation. *Pattern Recognition Letters*, 24(11):1725–1735, 2003.
- [12] M. Pappas and I. Pitas. Digital Color Restoration of Old Paintings. *Transactions on Image Processing*, (2):291–294, Feb. 2000.
- [13] E. Pichon, M. Niethammer, and G. Sapiro. Color histogram equalization through mesh deformation. In *IEEE International Conference on Image Processing (ICIP'04)*, 2003.
- [14] E. Reinhard, M. Ashikhmin, B. Gooch, and P. Shirley. Color transfer between images. *IEEE Computer Graphics Applications*, 21(5):34–41, 2001.
- [15] D.L. Ruderman, T.W. Cronin, and C.C. Chiao. Statistics of Cone Responses to Natural Images: Implications for Visual Coding. *Journal of the Optical Society of America*, (8):2036–2045, 1998.
- [16] T. Welsh, M. Ashikhmin, and K. Mueller. Transferring Color to Greyscale Images. In *Proceedings of ACM SIGGRAPH*, pages 227–280, San Antonio, 2002.



Figure 4. Examples of Colour Transfer. On the first row the Scottish landscape is recoloured to match the palette of the Alpine scenery. On the second row, the colour properties of the sunset are used to synthesise the 'evening' scene depicted at sunset. On the third row, the colour transfer allows to correct the change of lighting conditions induced by clouds.




Figure 5. Example of Colour Transfer for Image and Video Restoration. It is possible to recreate different atmospheres. Here an old faded film is transformed to match the colour scheme of a movie from the 70's and an pub ambience.