

Exact Affine Histogram Matching by Cumulants Transformation

Andrea Fusiello

DPIA - Università di Udine - Via Delle Scienze, 208 - Udine, Italy
name.surname@uniud.it

Dedicated to the memory of Egle Poggi (1936–2022)

Abstract. This paper proposes an exact solution to histogram matching under affine transformation, i.e., it shows how to retrieve an unknown affine transformation between the colors of two images. The key is the use of the third central moment (skewness) in addition to covariance and mean. These three moments (a.k.a. cumulants) are sufficient to determine all the d.o.f. of an affine mapping of the RGB space.

Keywords: Color Matching, Color Transfer, Colour Mapping

1 Introduction

Histogram matching is the transformation of the colors of one image (source) so that its histogram matches a specified target one (see Fig. 1). The well-known histogram equalization method is a special case in which the target histogram is a uniform distribution. It can be used to normalize two or more images, when they were acquired with different sensors (or from a sensor whose response changes over time), atmospheric conditions or global illumination. Affine transformations are particularly relevant for color histograms of images taken under varied illuminant conditions are related by an affinity [9].




Fig. 1. Source (left) and target (right) images. The goal is to transform the source so as to match the histogram of the target. Image from the UPenn dataset [19]

Histogram matching, a.k.a. color transfer, has been well studied in the literature, see [12] for a review. In their seminal paper, Reinhard et al. [13] match the mean and standard deviation of each axis separately after converting the source and target images into the decorrelated colour space $l\alpha\beta$ [15]. They write:

While it appears that the mean and standard deviation alone suffice to produce practical results, the effect of including higher moments remains an interesting question.

This paper shows empirically that an affine transformation is completely and exactly determined by the first three *multivariate* moments.

In recent years, the problem of color transfer evolved into "style transfer" (see e.g., [6]) where algorithms based on spatial (local) color mappings can handle applications such as time-of-day [18], weather and season change [4], painterly stylization [1] and transfer of artistic edits [17].

Motivated by the reduction of seam artefacts in mosaicking from aerial images [16], in this work we aim at recovering the *exact* affine mapping between the colors of two images (assuming that such mapping exists). Previous work in histogram matching fail to reach this goal: Pitié et al. [11] obtain an *approximate* non parametric map; others [20, 10] arbitrarily fix some degrees of freedom of a general affinity (this point will be clarified in Sec. 3), or restrict the solution to a subset of affine mappings [9, 13].

2 Background

In single-channel (gray-scale) images, the brightness values can be considered as the samples of a univariate random variable X , whose probability density function (pdf) is the normalized histogram. For any function g , also $Y = g(X)$ is a random variable.

2.1 Q-Q plots

Let F_X and F_Y be cumulative distribution function (CDF) of X and Y , respectively. The functions F_Y^{-1} and F_X^{-1} are the *quantile functions*, i.e. $F^{-1}(\alpha)$ is the α -th quantile of F . If g is monotone increasing, then $F_Y(y) = F_X(g^{-1}(y))$, and also $F_Y^{-1}(u) = g(F_X^{-1}(u))$. This means that the quantile functions transform according to g , hence the graph of g is described by the pairs $(F_Y^{-1}(\alpha), F_X^{-1}(\alpha))$ for all $\alpha \in [0, 1]$. This graph, where the quantiles of two variables are plotted one against the other is called the *Q-Q plot*. So, for gray-scale images, histogram matching can be accomplished by means of the Q-Q plot, that allows to find the exact transformation that makes the quantiles of X and Y to match. For example, if g is an affine map, the Q-Q plot will be a line, with a given (positive) slope and intercept.

In the case of color (RGB) images the underlying random variable is multivariate, the CDF is $\mathbb{R}^3 \rightarrow \mathbb{R}$ and the quantile function is not defined. The Q-Q plot trick can only be applied to each channel separately, but this would yield a very special affinity, whose linear part is represented by a (positive) diagonal matrix. Note that, since any matrix A can be decomposed with SVD as $A = UDV^T$, if one transform X with V^T and Y with U^T , then the residual transformation is the diagonal D and can be recovered with three Q-Q plots. The problem is that these U and V are unknown. In fact, Pitié et al. [11] iteratively apply random rotations to the color space and solve using Q-Q plots on each channel separately.

2.2 Cumulants

Instead of exploiting the quantiles of the two probability distributions, let us turn our attention to its *moments*, in particular to the first raw moment – the mean, and the second and third central moments, a.k.a. variance and skewness. These three moments coincide with the first three *cumulants* of the probability distribution (see e.g. [8]), and so they will be collectively referred to in the following. The cumulants of a probability distribution provide an alternative description of the distribution to those given by its moments: any two probability distributions whose moments are identical will have identical cumulants, and vice versa. As a matter of fact, the first three cumulants are also moments, but fourth and higher-order cumulants are not.

Using the Kronecker product, the first three cumulants of X write:

- Average $\kappa_1(X) = \mathbb{E}[X] = \mu_X$
- Covariance $\kappa_2(X) = \mathbb{E}[(X - \mu_X) \otimes (X - \mu_X)]$
- Skewness $\kappa_3(X) = \mathbb{E}[(X - \mu_X) \otimes (X - \mu_X) \otimes (X - \mu_X)]$

The third cumulant of a d -dimensional random vector $X = (X_1, \dots, X_d)^\top$ is an $d^2 \times d$ matrix, containing at most $d(d+1)(d+2)/6$ distinct elements. This is analogous to $\kappa_2(X)$, that is $d \times d$ and has at most $d(d+1)/2$ distinct elements, being symmetric. An elimination matrix G_d $d(d+1)(d+2)/6 \times d^3$ can be defined [7] that extracts only the unique elements from $\text{vec } \kappa_3(X)$.

Given a random sample from X of size N , the third sample cumulant is computed as follows:

$$\kappa_3(X) = \frac{1}{N} \sum_{i=1}^N (\mathbf{x}_i - \mathbf{m}) \otimes (\mathbf{x}_i - \mathbf{m})^\top (\mathbf{x}_i - \mathbf{m}) \quad (1)$$

where $\mathbf{x}_i \in \mathbb{R}^d$ is the i -th sample and $\mathbf{m} \in \mathbb{R}^d$ is the sample mean.

It is particularly useful to our goals to be able to determine how the cumulants change under a generic affine transformation $Y = AX + t$, where A is a $d \times d$ non-singular matrix and \mathbf{t} is a $d \times 1$ vector:

$$\kappa_1(Y) = A\kappa_1(X) + \mathbf{t} \quad (2)$$

$$\kappa_2(Y) = A\kappa_2(X)A^\top \quad (3)$$

$$\kappa_3(Y) = (A \otimes A)\kappa_3(X)A^\top. \quad (4)$$

The first two relationships are well-known, while the third is proved in [3]. Using the "vec trick" [14], we get the following equivalent formulae:

$$\text{vec } \kappa_1(Y) = (I_3 \otimes A) \text{vec } \kappa_1(X) + \mathbf{t} \quad (5)$$

$$\text{vec } \kappa_2(Y) = (A \otimes A) \text{vec } \kappa_2(X) \quad (6)$$

$$\text{vec } \kappa_3(Y) = (A \otimes A \otimes A) \text{vec } \kappa_3(X). \quad (7)$$

3 Method

Given two multivariate random variables X and Y we want to find the affine transformation that match the first three cumulants of Y with those of X .

The translation can be recovered from Eq. (2) once A has been determined, so let us concentrate on Eq. (3) and Eq. (4) and on the linear transformation A .

Let us start with the second cumulant (covariance), and follow [10]. Since $\kappa_2(X)$ and $\kappa_2(Y)$ are positive definite, let $\Sigma_X \Sigma_X^\top = \kappa_2(X)$ and, likewise, $\Sigma_Y \Sigma_Y^\top = \kappa_2(Y)$ (see Lemma 1 in the Appendix). Substituting this into (3) yields:

$$\Sigma_Y \Sigma_Y^\top = (A \Sigma_X)(A \Sigma_X)^\top \quad (8)$$

from which a solution is derived as:

$$A = \Sigma_Y \Sigma_X^{-1}. \quad (9)$$

Since the factorization of a psd matrix as per Lemma 1 is not unique, we arbitrarily set

$$\Sigma_X = U D^{1/2} \text{diag}(1, 1, \det(U)) \quad (10)$$

where

$$\kappa_2(X) = U D U^\top \quad (11)$$

is the spectral decomposition of $\kappa_2(X)$, with ordered eigenvalues. It is easy to see that $\Sigma_X \Sigma_X^\top = \kappa_2(X)$. The same applies to Σ_Y likewise.

Then, as observed by [10], the general solution, that matches the second order cumulants of X and Y , is given by:

$$A = \Sigma_Y Q \Sigma_X^{-1} \quad (12)$$

where $Q Q^\top = I$. This leaves three d.o.f. in the choice of Q , that have been fixed according to different considerations, in the literature.

For example, in the principal components (PCA) method [20, 9], a rotation is sought that align the principal axis of $\kappa_2(X)$ with those of $\kappa_2(Y)$, and Q is chosen among four possible matrices, namely:

$$\text{diag}(1, 1, 1), \quad \text{diag}(1, -1, -1), \quad \text{diag}(-1, 1, -1), \quad \text{diag}(-1, -1, 1). \quad (13)$$

In the Monge-Kantorovitch (MK) solution [10]:

$$A_{\text{MK}} = \hat{\Sigma}_X^{-1} \left(\hat{\Sigma}_X \kappa_2(Y) \hat{\Sigma}_X \right)^{\frac{1}{2}} \hat{\Sigma}_X^{-1} \quad (14)$$

where $\hat{\Sigma}_X = \kappa_2(X)^{\frac{1}{2}}$ and $\hat{\Sigma}_Y = \kappa_2(Y)^{\frac{1}{2}}$. This corresponds to

$$Q = \Sigma_Y^{-1} A_{\text{MK}} \Sigma_X. \quad (15)$$

Keeping in mind that $\hat{\Sigma}_X = U D^{1/2} U^\top$ (and likewise for $\hat{\Sigma}_Y$) it can be proven that Q is indeed an orthogonal matrix.

None of the previous methods, however is able to exactly recover the affine transformation, when the two histograms are affinely related. Our method instead achieve exact recovery by exploiting the third order cumulant κ_3 in Eq. (7). In particular, we fix the remaining three d.o.f. by solving the non-linear system of 10 equations:

$$G_3\kappa_3(Y) = G_3(A \otimes A \otimes A) \text{vec } \kappa_3(X) \quad \text{with } A = \Sigma_Y Q \Sigma_X^{-1} \quad (16)$$

in the least squares sense with the Levenberg-Marquardt method.

Experiments reported in the next section show that in this way the affinity is always recovered, upon convergence of the numerical method.

If we restrict to A with positive determinant, then $Q \in SO(3)$ and so we can parametrize it with the Euler angles ω, ϕ, ϑ . This assumption can be relaxed at the price of searching for the minimum in $-SO(3)$ as well, since $O(3)$ has two connected components.

The Matlab code that implements our method is available on the web at <http://www.diegm.uniud.it/fusiello/demo/ahm/>

4 Experiments

In order to validate the method in a controlled setting, we applied random affine transformations (with positive determinant) to the reference image shown in Fig. 1 and then run our method to retrieve the affine transformation matrix. Convergence is only local, so we started the minimization in a neighborhood of the ground truth.

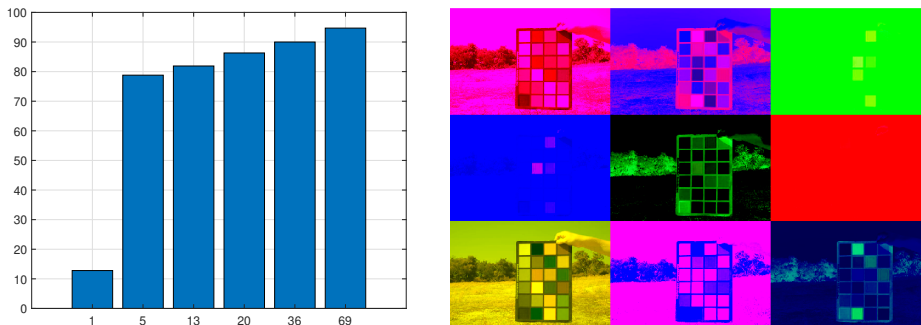


Fig. 2. Convergence rate as a function of the number of initial points (left); the first bar corresponds to the MK solution, while the second includes the four PCA solutions. Sample of images mapped according to the random affine transformation (right).

The error was measured as the Frobenius norm of the difference between the output matrix and the ground truth one. Upon convergence it was of the order of $1e - 10$ (or less). To be more realistic we also started our method from multiple

points by picking N regularly spaced samples in $SO(3)$, in addition to the MK solution and the four PCA solutions. These samples are the centers of an equal area partitioning of the unit sphere in \mathbb{R}^4 [5], which is a double cover of $SO(3)$ (antipodal points represent the same rotation) via the quaternion representation of rotations. Figure 2 shows the convergence rate as a function of N for 1000 random affine transformations. These experiments indicate empirically that the method is able to recover the correct affine transformation, up to numerical errors, when suitably initialized.

In the real experiments that follow the underlying transformations are more regular than the random ones (see Fig. 2 for a sample), although not necessarily affine, so we initialized instead the method with the MK solution, that provides a reasonable starting point, thanks to the minimization of the transportation cost. Our solution is nevertheless different from MK, as one can see in Fig. 3.



Fig. 3. Comparison with results reported in [11] (IDT) and [10] (MK)

In our first real experiment we shoot eight pictures of an indoor scene with a Sony $\alpha 6300$ camera, while varying the temperature of the LED light source from 3200K to 5500K. The image at 5500K was chosen as the target one and all the others have been transformed with the affinity found by our algorithm to match the target histogram. The results are shown in Fig. 4.

A similar experiment has been conducted outdoor (Fig. 5). A series of nine photos have been taken during the day, one has been selected as target and all the others have been transformed to match its histogram.

Fig. 6 shows the result of our method on a portion of the dataset that originally motivated this research [16], where the effect of the illumination change is particularly prominent.



Fig. 4. Indoor scene with varying light temperature. The top image is the target, then from the second row onward images are: source, output and checkerboard comparison.



Fig. 5. Outdoor scene in different time of the day. The top image is the target, then from the second row onward images are: source, output and checkerboard comparison.




Fig. 6. Top row: two images of the same area under different illumination, source (left) and target (right) Bottom row: output (left), i.e., source image transformed to match the target, and checkerboard comparison (right) between target and output (a small misalignment is present).

In the last experiment we tried also to apply our algorithm to color transfer, although this was not our focus. As a global method, it achieved reasonable results on some of images used in [6] (Fig. 7), while it failed on more challenging cases where local operations (and a segmentation) are required, e.g. changing the color of the eyes, or lighting the windows of a building.

5 Conclusions

In this paper we have shown how to retrieve an unknown affine transformation between the colors of two images (formally, between two probability distributions) thanks to the use of the third cumulant, which allows to fix the three d.o.f. that remain after matching mean and covariance. This study was motivated by image mosaicking, but can also find application in color transfer or in histogram normalization [9]. Future work will study the extension of this approach to more general classes of transformations.

Appendix

Lemma 1 [2] A real matrix M is positive semi-definite (psd) if and only if it can be decomposed as $M = \Sigma \Sigma^\top$.

The decomposition is not unique: if $M = \Sigma \Sigma^\top$ also $M = \Sigma Q Q^\top \Sigma^\top = (\Sigma Q)(\Sigma Q)^\top$ with $Q Q^\top = I$. Requiring Σ to be psd as well makes the factorization unique. In this case Σ is referred to as the *square root* of M .

Acknowledgements

Images shown in Figs. 4 and 5 are courtesy of Leonardo Fusiello, Fig. 6 are courtesy of Helica srl. The MATLAB implementation of the recursive zonal equal area sphere partitioning is due to P. Leopardi.



Fig. 7. Some examples of color transfer on images from [6]. The output is the affine transform of the source to match the target palette.

References

1. Gatys, L.A., Ecker, A.S., Bethge, M.: Image style transfer using convolutional neural networks. In: IEEE Conference on Computer Vision and Pattern Recognition. pp. 2414–2423 (2016)
2. Horn, R.A., Johnson, C.R.: Matrix Analysis. Cambridge University Press, Cambridge; New York, 2nd edn. (2013)
3. Kollo, T., von Rosen, D., Hazewinkel, M.: Advanced Multivariate Statistics with Matrices. Springer (2005)
4. Laffont, P.Y., Ren, Z., Tao, X., Qian, C., Hays, J.: Transient attributes for high-level understanding and editing of outdoor scenes. ACM Transactions on Graphics (proceedings of SIGGRAPH) **33**(4) (2014)
5. Leopardi, P.: A partition of the unit sphere into regions of equal area and small diameter. Electronic Transactions on Numerical Analysis **25**, 309–327 (2006)
6. Luan, F., Paris, S., Shechtman, E., Bala, K.: Deep photo style transfer. In: Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition (July 2017)
7. Meijer, E.: Matrix algebra for higher order moments. Linear Algebra and its Applications **410**, 112–134 (2005)
8. Muirhead, R.J.: Aspects of Multivariate Statistical Theory. Wiley-Interscience (2005)
9. Pei, S.C., Tseng, C.L., Wu, C.C.: Using color histogram normalization for recovering chromatic illumination-changed images. J. Opt. Soc. Am. A **18**(11), 2641–2658 (Nov 2001)
10. Pitié, F., Kokaram, A.: The linear monge-kantorovitch linear colour mapping for example-based colour transfer. In: 4th European Conference on Visual Media Production. pp. 1–9 (2007)
11. Pitié, F., Kokaram, A., Dahyot, R.: N-dimensional probability density function transfer and its application to color transfer. In: IEEE International Conference on Computer Vision. vol. 2, pp. 1434–1439 Vol. 2 (2005)
12. Pitié, F.: Advances in colour transfer. IET Computer Vision **14**, 304–322 (09 2020)
13. Reinhard, E., Ashikhmin, M., Gooch, B., Shirley, P.: Color transfer between images. IEEE Comput. Graph. Appl. **21**(5), 34–41 (sep 2001)
14. Roth, W.E.: On direct product matrices. Bulletin of the American Mathematical Society **40**(6), 461 – 468 (1934)
15. Ruderman, D.L., Cronin, T.W., Chiao, C.C.: Statistics of cone responses to natural images: implications for visual coding. J. Opt. Soc. Am. A **15**(8), 2036–2045 (1998)
16. Santellani, E., Maset, E., Fusiello, A.: Seamless image mosaicking via synchronization. ISPRS Annals of Photogrammetry, Remote Sensing and Spatial Information Sciences **IV-2**, 247–254 (2018)
17. Shih, Y., Paris, S., Barnes, C., Freeman, W.T., Durand, F.: Style transfer for headshot portraits. ACM Trans. Graph. **33**(4) (jul 2014)
18. Shih, Y., Paris, S., Durand, F., Freeman, W.T.: Data-driven hallucination of different times of day from a single outdoor photo. ACM Trans. Graph. **32**(6) (2013)
19. Tkačik, G., Garrigan, P., Ratliff, C., Milčinski, G., Klein, J., Seyfarth, L., et al.: Natural images from the birthplace of the human eye. PLOS ONE **6**(6), 1–12 (06 2011)
20. Trussell, H., Vrhel, M.: Color correction using principle components. In: Image Processing Algorithms and Techniques II. SPIE Proceedings, vol. 1452, pp. 2–9 (1991)