

TOWARDS AUTOMATED COLOUR GRADING

F. Pitié, A.C. Kokaram and R. Dahyot

University of Dublin, Trinity College, Ireland,
fpitie@mee.tcd.ie, anil.kokaram@tcd.ie, dahyot@mee.tcd.ie

Keywords: colour grading, colour transfer, N-dimensional pdf matching

Abstract

This article proposes an original method for grading the colours between different images or shots. The first stage of the method is to find a one-to-one colour mapping that transfers the palette of an example target picture to the original picture. This is performed using an original and parameter free algorithm that is able to transform any N-dimensional probability density function into another one. The proposed algorithm is iterative, non-linear, has a low computational cost. Applying the colour mapping on the original picture allows reproducing the same ‘feel’ as the target picture, but can also increase the graininess of the original picture, especially if the colour dynamic of the two pictures is very different. The second stage of the method is to reduce this grain artefact through an efficient post-processing algorithm that intends to preserve the gradient field of the original picture. 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

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 match the colour between frames by tuning parameters. This is delicate task since the change in lighting conditions induces a very complex change of illumination. The method presented in this paper however succeeds in automating this painstaking process even when the lighting conditions have dramatically changed as shown in figure 7.

The first stage of colour grading is to transfer the full colour statistics of a target picture example to the original picture. 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 one-to-one mapping function that transforms the source samples into a new ensemble of samples that exhibits the same pdf as the target samples. Finding a one-to-one mapping for one dimensional (1D) samples is a simple task that can be solved in the same way as grayscale picture equalisation [3]. The difficulty resides in extending the method to higher dimensions. 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 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 or pixels processed. Moreover working in 1D considerably reduces the memory load required by the algorithm. This implies that it is possible to have more bins in the histogram and thus achieve higher accuracy in the colour mapping.

The one-to-one colour mapping to the original picture makes the transformed picture having the same ‘feel’ that the picture example but it might also produce some grain artefacts on parts of the picture. This can be understood if the mapping is done from a low dynamic range to a high dynamic range. The resulting mapping is stretched and thus enhances the noise level of the picture, which makes the transformed picture appear as a grainy. The second step to colour grading is therefore to reduce this artefact. The method proposed is to use a variational approach to preserve the gradient of the original while preserving also the colour transfer characteristic. Preserving the gradient of the original picture especially protects the flat areas and more generally results in the exact aspect of film grain/noise than the original image.

The overall method, which includes as a first step finding the colour mapping and then reducing the possible grain artefact, is simple and gives attractive results.

2 Related Work

Transfer of Colour Statistics. One popular method [12] 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 $\lambda\alpha\beta$ which has been proposed to account for human-perception [13]. 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 particular in [9] 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 a uniform distribution. This method can be seen as being related to warping theory which is explicitly used in [6] where the transfer of the 2D chromatic space is performed directly by using a 2D-biquadratic warping.

Without having to invoke image warping, a natural extension of the 1D case is to treat the mapping via linear programming [7]. This method allows for almost any kind of colour transformation. The major disadvantage of the method is its computational cost which 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 [12] 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 [5]. The extreme case concerns the difficult task of colouring grayscale pictures. Retrieving the missing chrominance channels values can be done by taking advantage of similarities between spatial neighbourhoods of the two pictures [15, 4].

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. By contrast our method is able to map shades of green to shades of orange in figure 5.

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 colour distributions. Thirdly, the proposed method for reducing grain artefact is unique and results in



Figure 1: Example of Colour transfer using Reinhard [12] Colour Transfer. The transfer fails to resynthesise the colour scheme of the target image. To be successful the method would require human interaction.

high quality picture. 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 specifically addressed in this community previously.

3 Colour Transfer through N-Dimensional pdf Transfer

This section considers the problem of finding the one-to-one colour mapping t that transforms the colour pdf of the original picture to the colour pdf of the target picture. The method can be generalised for any pdf transfer from N -dimensional continuous Random Variable. Denote $f(x)$ and $g(y)$ the pdf of X and Y , the original and target N -dimensional continuous Random Variables respectively. In colour transfer, the samples encapsulate the three colour components $x_i = (r_i, g_i, b_i)$. The problem is to find the mapping function t that will transform f to g .

Dimension $N = 1$ (Grayscale Pictures) This is a well known problem [3] which offers a simple solution. The mapping implies the following variational relation:

$$f(x) dx = g(y) dy \quad (1)$$

Integrating the previous equation results in:

$$\int_{x_0}^{x_1} f(x) dx = \int_{t(x_0)}^{t(x_1)} g(y) dy \quad (2)$$

$$t(x) = \mathcal{C}_Y^{-1}(\mathcal{C}_X(x)) \quad (3)$$

where \mathcal{C}_X and \mathcal{C}_Y are the cumulative pdfs of X and Y . This can be easily solved using discrete lookup tables.

Dimension $N \geq 2$. (Colour Pictures) To motivate the algorithm, consider the three dimensional case of colour transfer. A simple approach to the problem would be to use the one dimensional scheme *separately* on the chrominance channels of the the original picture. For instance, the histogram of the blue chrominance should be identical to the one of the

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 one rotation matrix R
 - 4: rotate the samples: $\mathbf{x}_r \leftarrow R\mathbf{x}^{(k)}$ and $\mathbf{y}_r \leftarrow R\mathbf{y}$
 - 5: project the samples on all axis i to get the marginals f_i and g_i
 - 6: for each axis i , find the 1D transformation t_i that matches the marginals f_i into g_i
 - 7: 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.
 - 8: rotate back the samples: $\mathbf{x}^{(k+1)} \leftarrow R^{-1}\mathbf{x}_r$
 - 9: $k \leftarrow k + 1$
 - 10: **until** convergence on all marginals for every possible rotation
 - 11: The final one-to-one mapping t is given by: $\forall j, \mathbf{x}_j \mapsto t(\mathbf{x}_j) = \mathbf{x}_j^{(\infty)}$
-

target picture. Most often, the ‘feel’ of the resulting picture is not right but an improvement is still noticeable. Applying again this technique offers no improvement since the histograms of the channels are now identical. The main point of the paper is to show that re-applying this simple technique but in a different colour space actually results in a new improvement. It is even possible, by applying iteratively this simple operation on a number of different colour spaces, to make the original picture present exactly the same colour statistics as the target picture.

Description of the Algorithm. 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 . This operation is equivalent to a colour space conversion. In a second step the samples of both distributions are projected on the new axis which gives the marginals $f_1 \dots f_N$ and $g_1 \dots g_N$. Then it is possible using equation (3) 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 but transforms the pdf $f^{(k)}$ into $f^{(k+1)}$. As shown in the appendix, the average Kullback-Leibler divergence between both pdf’s actually decreases at each iteration. The algorithm converges in practice 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.

Kick Start Using PCA. Although the method can cope with any kind of transformation, it might be desirable as a

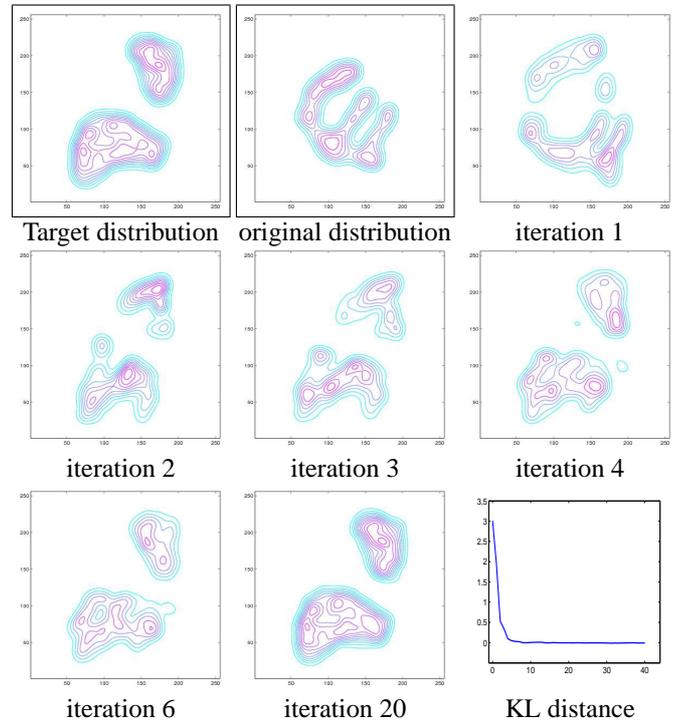


Figure 2: Example of a two dimensional pdf transfer. The experimental measure of the Kullback-Leibler distance at each iteration illustrates the convergence of the original distribution to the target one by the transfer method.

first step to use a different technique for registering global transformations between the two distributions. In particular, the proposed method is not naturally designed to find rotation between the original and target dataset. A simple solution is to use the Principal Component Analysis of each of the images to find the rotation transformation between the original samples and the target samples:

$$X^* = U_Y U_X^{-1} X \quad (4)$$

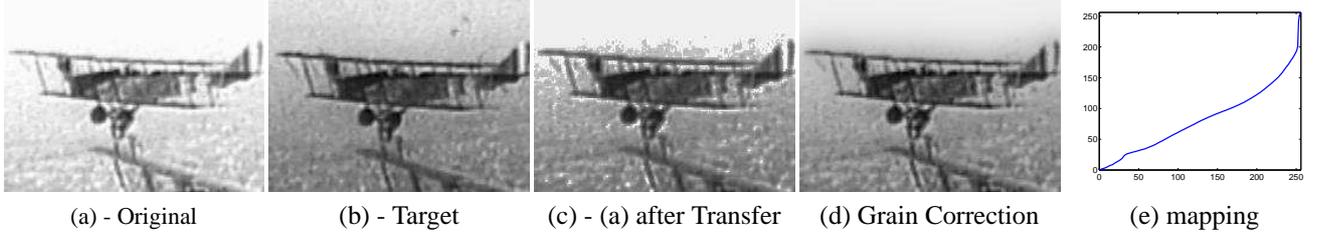


Figure 3: Result of grain reducing. The two consecutive archive frames (a) and (b) suffer from extreme brightness variation (this artefact is known as flicker [10]). The corresponding mapping transformation (e) is over-stretched, which results in an increased level of noise on the mapped original frame (c). The proposed grain artefact reducer is able to reproduce the noise level of the original picture. The top of the original picture is saturated and cannot be retrieved but the algorithm succeeds in preserving the soft gradient.

where $U_X = [\mathbf{u}_X^1, \dots, \mathbf{u}_X^N]$ and $U_Y = [\mathbf{u}_Y^1, \dots, \mathbf{u}_Y^N]$ are the $N \times N$ orthogonal matrices containing the eigenvectors of the covariance matrices of X and Y . To obtain a correct transformation, the eigenvectors in U_X and U_Y have to correspond. This is achieved by ordering them with respect to the magnitude of the corresponding eigenvalues and making sure that they do not point in opposite directions. i.e.

$$\forall i \leq N, \quad {}^t \mathbf{u}_X^i \cdot \mathbf{u}_Y^i > 0 \quad (5)$$

This operation cannot handle non-linear transformations but offers an interesting speed up if it is used as a first step, before the actual pdf transfer algorithm.

4 Reducing Grain Noise Artefacts

The colour mapping to the original picture transfers correctly the target colour palette to the original picture but it might also produce some grain artefacts as shown in figure 3 and 6. When the content differs or the dynamic range of both pictures is too different, the resulting mapping function can be stretched on some parts (see figure 3-d), and thus enhances the noise level. This can be understood by taking the simple example of a linear transformation t of the original picture X : $t(X) \rightarrow aX + b$. The overall variance of the resulting picture is changed to $\text{var}\{t(X)\} = a^2 \text{var}\{X\}$. This means that a greater stretching ($a > 1$) produces a greater noise.

The solution proposed here to reduce the grain is to run a post-processing algorithm that forces the level of noise to remain the same. The idea is to adjust the gradient field of the picture result so that it matches the original picture. If the gradient of both pictures are similar, the level of noise will be the same. Matching the gradient of a picture has been addressed in different applications like image stitching [8] or high dynamic range compression [2], and it can be efficiently solved using a variational approach.

Denote $I(x, y)$ the original picture¹. Let $t : I \rightarrow t(I)$ be the

¹to simplify the coordinates are omitted in the expressions and I, J, ψ, ϕ , etc. actually refer to $I(x, y), J(x, y), \psi(x, y)$ and $\phi(x, y)$.

colour transformation. The problem is to find a modified image J of the mapped picture $t(I)$ that minimises on the whole picture range Ω

$$\min_J \iint_{\Omega} \phi \cdot \|\nabla J - \nabla I\|^2 + \psi \cdot \|J - t(I)\|^2 \quad (6)$$

with Neumann boundaries condition $\nabla J|_{\partial\Omega} = \nabla I|_{\partial\Omega}$ to match the the gradient of J with the gradient of I at the picture border. The term $\|\nabla J - \nabla I\|^2$ forces the image gradient to be preserved. The term $\|J - t(I)\|^2$ ensures that the colours remain close to the target picture. Without $\|J - t(I)\|^2$, a solution of equation (6) will be actually the original picture I .

The weight fields ϕ and ψ affect the importance of both terms. The weight field ϕ emphasises that only flat areas have to remain flat but that gradient can change at object borders:

$$\phi = \frac{\phi_0}{1 + \|\nabla I\|} \quad \text{with } \phi_0 \text{ constant} \quad (7)$$

The weight field ψ accounts for the possible stretching of the transformation t . Where ∇t is big, the grain becomes more visible:

$$\psi = \frac{\psi_0}{1 + \|(\nabla t)(I)\|} \quad \text{with } \psi_0 \text{ constant} \quad (8)$$

$(\nabla t)(I)$ refers to the gradient of t for the colour I .

Numerical Solution. The minimisation problem in equation (6) can be solved using the variational principle which states that the integral must satisfy the Euler-Lagrange equation:

$$\frac{\partial F}{\partial J} - \frac{d}{dx} \frac{\partial F}{\partial J_x} - \frac{d}{dy} \frac{\partial F}{\partial J_y} = 0 \quad (9)$$

where

$$F = \phi \cdot \|\nabla J - \nabla I\|^2 + \psi \cdot \|J - t(I)\|^2 \quad (10)$$

from which the following can be derived:

$$\phi \cdot J - \text{div}(\psi \cdot \nabla J) = \phi \cdot t(I) - \text{div}(\psi \cdot \nabla I) \quad (11)$$

This is an elliptic partial differential equation. The expression $\text{div}(\psi \cdot \nabla I)$ can be approximated using standard finite differences [14] by:

$$\text{div}(\psi \cdot \nabla I) \approx \sum_{l \in \mathcal{N}_i} \frac{\psi_l + \psi_i}{2} (I_l - I_i) \quad (12)$$

where \mathcal{N}_i corresponds to the four neighboring pixels of i .

This yields in a linear system which can be solved by standard iterative methods like SOR, Gauss-Seidel with multigrid approach. Implementations of these numerical solvers are widely available and one can refer for instance to the Numerical Recipes [11].

5 Results

The remaining pages show some results from the Colour Grading technique proposed in the paper. In Figure 5 for instance, the original mountain picture in (a) is used as a target colour scheme for the sea-side image in (b) the third column. The result of the transfer appears in the second column.

Examples of Colour Grading for matching lighting conditions are presented in 7. On the first row, the colour properties of the sunset are used to synthesise the ‘evening’ scene depicted at sunset. On the second row, the colour grading allows correction of the change of lighting conditions induced by clouds. Even when using the grain artefact reducer, an 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 8 displays examples of colour restoration of faded movies. The idea is similar to colour grading as the idea is 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 technique for colour grading. The technique is based on an exact transfer of colour pdf of the target picture, which is guaranteed to converge at low computation costs. The possible grain artefacts are removed in a second step. The overall technique is simple, easily to implement and works for a large variety of scenarios, even when the example picture is very different from the processed images.

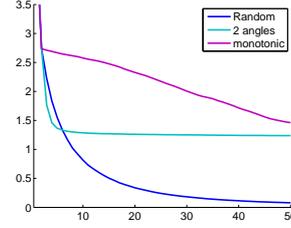


Figure 4: Average convergence rate for a two dimensional pdf transfer for 100 distributions. The random rotation selection strategy (in solid blue) eventually converges whereas strategy involving a monotonic increase of the rotation angle or a strategy involving an alternation between two rotation angles eventually stalls.

Acknowledgements

We 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.

A Convergence Issues.

Figure 4 shows the convergence of the algorithm, based on the average measure of the Kullback-Leibler divergence at each iteration for 100 different distributions. Considering a random set of rotation directions is sufficient to converge to the solution. The optimal rotation sequence is however still a subject of research.

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] R. Fattal, D. Lischinski, and M. Werman. Gradient domain high dynamic range compression. In *Proceedings of the 29th annual conference on Computer graphics and interactive techniques (SIGGRAPH '02)*, pages 249–256, New York, NY, USA, 2002. ACM Press.
- [3] R. C. Gonzalez and R. E. Woods. *Digital Image Processing*. Addison Wesley, 1992.
- [4] 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.

- [5] J. Jia, J. Sun, C-K. Tang, and H-Y. Shum. Bayesian correction of image intensity with spatial consideration. In *8th European Conference on Computer Vision (ECCV)*, 2004.
- [6] L. Lucchese and S. K. Mitra. a new Method for Color Image Equalization. In *IEEE International Conference on Image Processing (ICIP'01)*, 2001.
- [7] J. Morovic and P-L. Sun. Accurate 3d image colour histogram transformation. *Pattern Recognition Letters*, 24(11):1725–1735, 2003.
- [8] P. Pérez, M. Gangnet, and A. Blake. Poisson image editing. *ACM Trans. Graph.*, 22(3):313–318, 2003.
- [9] E. Pichon, M. Niethammer, and G. Sapiro. Color histogram equalization through mesh deformation. In *IEEE International Conference on Image Processing (ICIP'04)*, 2003.
- [10] F. Pitié, R. Dahyot, F. Kelly, and A. C. Kokaram. A new robust technique for stabilizing brightness fluctuations in image sequences. In *2nd Workshop on Statistical Methods in Video Processing In conjunction with ECCV 2004*, Prague, Czech Republic, May 2004.
- [11] W. Press, S. Teukolsky, W. Vetterling, and B. Flannery. *Numerical Recipes in C: The Art of Scientific Computing*. Cambridge University Press, New York, NY, USA, 1992.
- [12] E. Reinhard, M. Ashikhmin, B. Gooch, and P. Shirley. Color transfer between images. *IEEE Computer Graphics Applications*, 21(5):34–41, 2001.
- [13] 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.
- [14] J. Weickert, B. ter Haar Romeny, and M. Viergever. Efficient and Reliable Schemes for Nonlinear Diffusion Filtering. *IEEE Transactions on Image Processing*, 7(3):398–410, march 1998.
- [15] T. Welsh, M. Ashikhmin, and K. Mueller. Transferring Color to Greyscale Images. In *Proceedings of ACM SIGGRAPH*, pages 227–280, San Antonio, 2002.

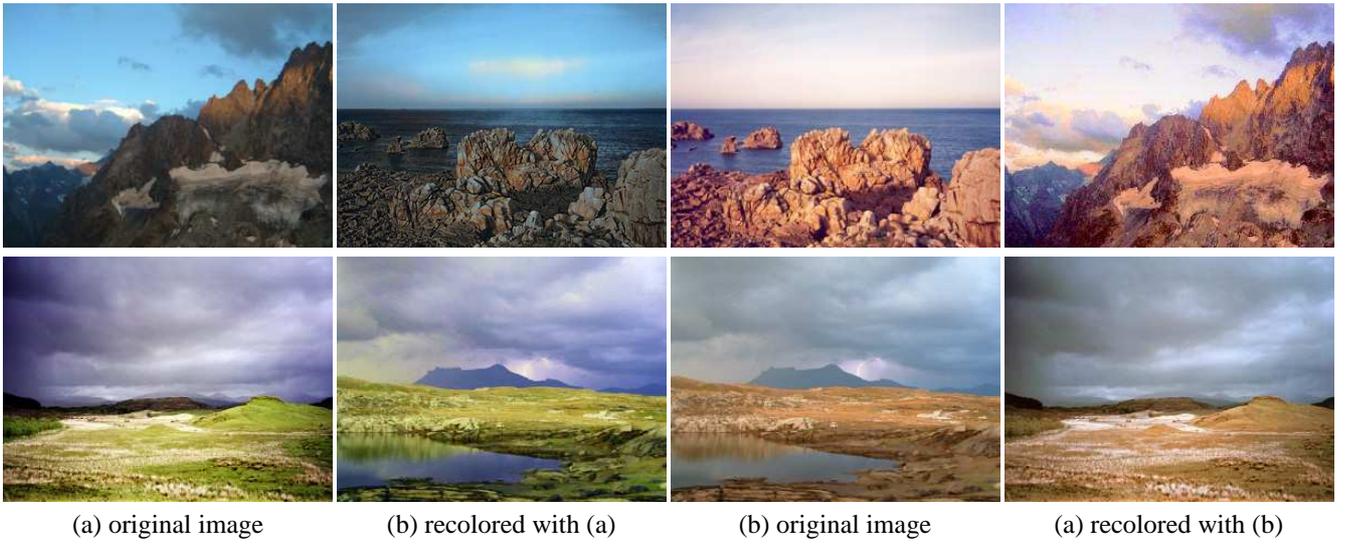


Figure 5: Examples of Colour Grading: the second column is obtained by transferring the colour pdf from the image of the first column (a) to image of the third column (b). The last column is obtained by transferring the colour pdf from image (b) to image (a).



Figure 6: Result of artefact grain reducing for colour picture. See how the details of the picture are preserved, while the spurious graininess in sky has been washed out.

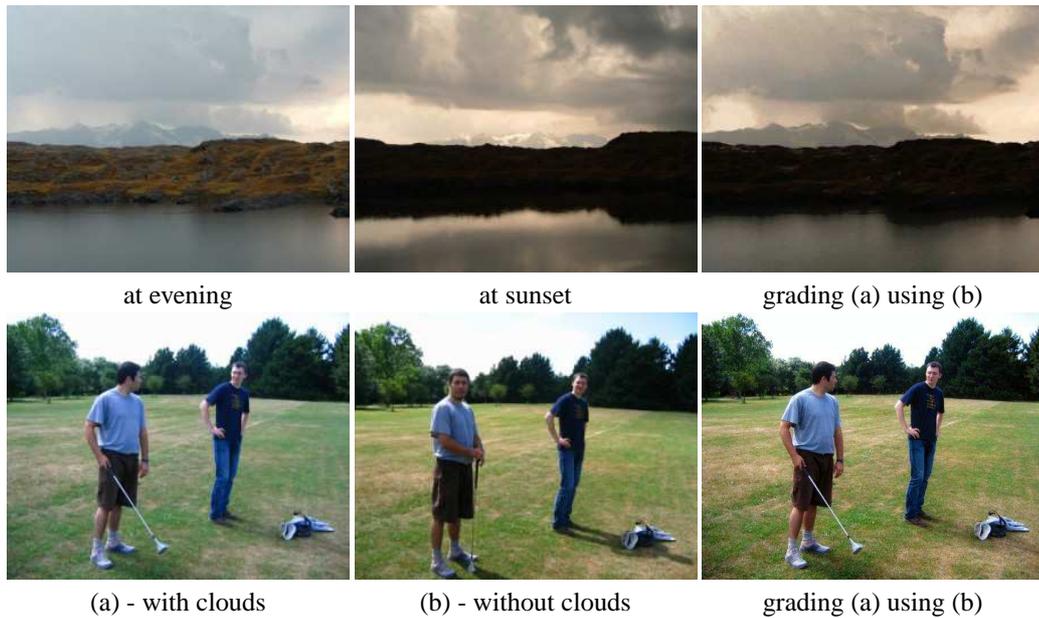


Figure 7: Examples of Colour Grading for matching lighting conditions. On the first row, the colour properties of the sunset are used to synthesise the ‘evening’ scene depicted at sunset. On the second row, the colour grading allows to correct the change of lighting conditions induced by clouds.



Figure 8: Example of Colour Grading 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 a pub ambiance.