Example-Based Image Manipulation

Erik Reinhard; Max Planck Institute for Informatics; Saarbrücken, Germany

Abstract

Many color-related image adjustments can be conveniently executed by exposing at most a small number of parameters to the user. Examples are tone reproduction, contrast enhancements, gamma correction and white balancing. Others require manual touch-ups, applied by means of brush strokes. More recently, a new class of algorithms has emerged, which transfers specific image attributes from one or more example images to a target. These attributes do not have to be well-defined and concepts that are difficult to quantify with a small set of parameters, such as the "mood" of an image, can be instilled upon a target image simply through the mechanism of selecting appropriate examples. This makes example-based image manipulation a particularly suitable paradigm in creative applications, but also finds uses in more technical tasks such as stereo pair correction, video compression, image colorization, panorama stitching and creating night-time images out of day-light shots.

General Appearance Transfer

Selecting an image to serve as an example for the manipulation of another image is a powerful technique to achieve complex adjustments. The key insight is that selecting an image from one's own private collection or from the millions of images freely available on the internet is effectively a creative, simple and enjoyable activity. Dependent on the task, the alternatives would involve either setting a potentially large set of user parameters to steer an image filter, or would require the user to painstakingly paint over pixels.

In an early example, filters can be "learned" from pairs of images. Assume that an image I_{in} is altered in some way, resulting in I_{out}, and that the images before and after applying this process are available. It would then be possible to learn a filter that mimics these alterations so that the filter can be applied to any number of new images. This idea is known as Image Analogies and can be applied to a variety of scenarios, ranging from very simple tasks such as blurring and embossing, to texture synthesis, super-resolution, texture transfer and image colorization [10]. An example of texture transfer is shown in Figure 1. In addition this technique can be used to learn artistic filters, allowing various drawing and painting styles to be applied to photographs. It has even been used to synthesize terrains for flight simulation.

This technique is generally applicable, as long as an example image pair can be found. In many cases there may be an advantage in having a somewhat less general technique which applies a specific transfer to an image based on the input of a single example image. The type of problem that can then be solved takes the form of making an image look more like another image in some predetermined fashion.

An example of this is color transfer, in which an image is made to look more like another image in terms of color content [19]. This particular problem has received considerable attention in the last decade, as it allows a range of problems to be solved. These include letting rendered images attain a more natural color palette, matching stereo image pairs, creating night-time images from day-time images, and stitching panoramas. Moreover, the basic algorithmic form of this class of algorithm is exceptionally straightforward, although many variants and alternatives are now known.

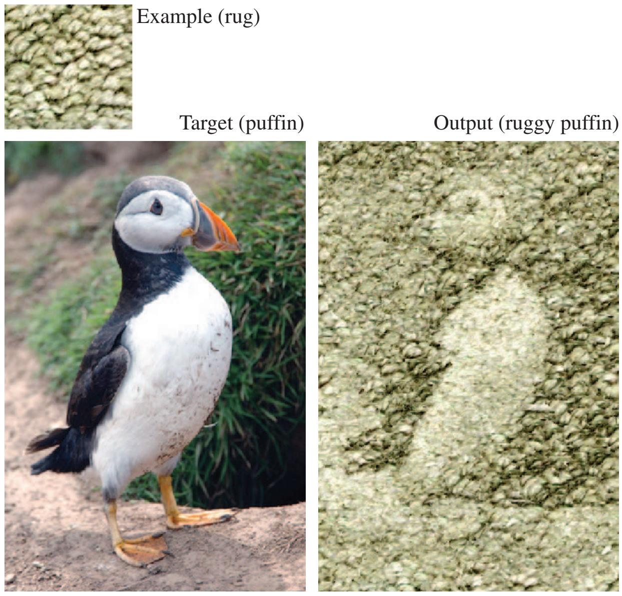

Figure 1. *Texture transfer using the image analogies framework [10].*

Targeting professional photographers specifically, tone management was proposed as a more encompassing technique to allow images to mimic a specific photographic style. Rather than transfer colors, the technique matches brightness, contrast as well as a sense of texture.

Color Transfer

Adjusting colors in an image is often necessary. In particular white balancing is a well-known post-processing technique. White balancing can be seen as adjusting the mean color in the image to a desired value. It leaves relationships between colors intact.

More dramatic and creative results can be obtained by not only changing the mean color of the image, but also manipulating how clusters of colors relate to each other within an image. In the case of color transfer, this is achieved by pushing the colors in one image to resemble those in another image.

Assuming that colored pixels are specified by tristimulus values in some three-dimensional color space, it is possible to identify clusters of colors in this 3-dimensional space, for instance by means of employing Gaussian mixture models [24, 28, 30]. If this is independently done on an image that is to be adjusted as well as on an example image which has a desirable appearance, then we end up with two sets of clusters. Clusters in the target image can then be warped to resemble those found in the example image. This adjusts pixel values in the target image such that the mood of the image will begin to resemble the example image. Alternatively, probability distributions can be transferred between pairs of images in three-dimensional color space [16, 17]. The choice of color space in each of these approaches is relatively unimportant as long as both images are defined in the same color space.

A much simpler solution relies on an important observation, which is that it is possible to decorrelate the pixel data in this three-dimensional color space. It was found that for natural images, a linear transform from cone space to a color opponent space, as implemented in the human retina, decorrelates the three color channels and in practice achieves a result that is close to independence [21]. This result was obtained by capturing a set of spectral images and converting them to LMS cone space. This image ensemble was then subjected to Principal Components Analysis (PCA), which decorrelates the color channels. It was found that the resulting axes closely correspond to the color opponent space encoded by the ganglion cells in the retina.

The implication for color transfer is that if the input constitutes a pair of natural images, then it would be possible to convert both of these images to a well-chosen color opponent space. The decorrelation achieved in this manner would then enable the above three-dimensional cluster matching problem to be simplified to three one-dimensional problems, without much chance of cross-talk between channels [19].

The simplest possible approach to achieve color transfer is then to convert both target and example images to a color opponent space, followed by computing the mean and standard deviation of all pixel values in each of the six color axes (three for each image). Color transfer can then be achieved by shifting and scaling the pixel values in the target image so that the means and standard deviations of this image are matched to the example image. Finally, the image is converted back to RGB for display or subsequent processing.

This method is more robust than could be expected on the basis of its simplicity. However, it does rely on the compositions of the example and target images to be largely similar. If this assumption is broken, or if significant correlations remain after conversion to a color opponent space, the results may become unexpected.

There are several ways in which this basic algorithm can be improved. First, if the compositions of target and example images are not matched, then it would be possible to select small regions in source and target images and compute statistics on those regions [19], a technique that can be improved by optimizing for global consistency [14]. Instead of selecting swatches, user input could also be guided by means of strokes [3].

Second, there are many color opponent spaces, of which an appropriate space could be selected. Further, it would be possible to subject both example and target images to active decorrelation, i.e. by running either Principal Components Analysis [1, 2, 11, 12, 29] or Independent Components Analysis [9] on them. However, it should be noted that there exists evidence that the CIELAB color space may well be more robust than applying PCA, leading to acceptable imagery in a larger number of cases [20].

A further refinement would be to transfer higher order moments, such as skew and kurtosis, to reshape the target image's histogram. However, from experience this produces very little extra effect relative to transferring the mean and standard deviation. Moreover, this would require optimisation, increasing computational complexity.

In the limit, it would be possible to make the histograms of the target image identical to the example image. However, we found that straightforward histogram matching often produces results that look too heavy-handed [18]. However, histograms do afford a good representation of the image to allow effective and robust color transfer, while still operating three times independently on the three different color channels in color opponent space.

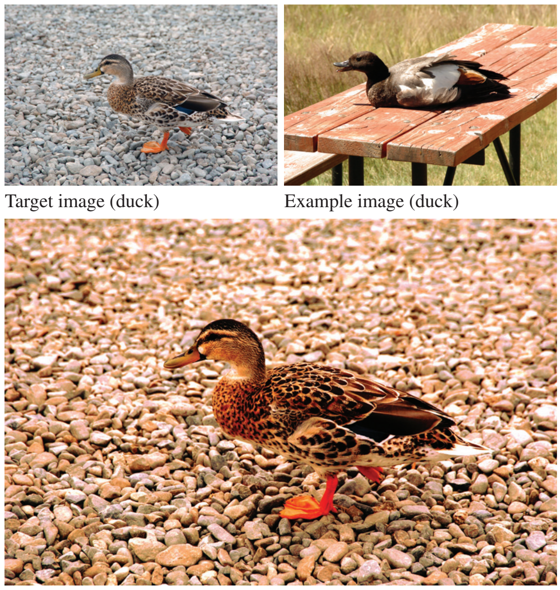

Target image (duck) Example image (duck)

Color transfer result (duck)

Figure 2. *Color transfer using progressive histogram reshaping [18].*

In particular, it is possible to compute histograms in each of the six color channels, followed by warping the histograms of the target image to resemble the histograms of the example image. This produces a new histogram which is then instilled upon the target image through histogram matching. An example result is shown in Figure 2. The method of warping is key to success here. If the histograms are only shifted and stretched, then the original color transfer algorithm would emerge [19]. If the histograms are warped to be identical to the example image's histograms, then the method would reduce to histogram matching.

However, if the warp sits somewhere in-between these extremes, a more robust algorithm can be created [18]. Warping a histogram part-way towards a second histogram can be achieved by first observing that images produce histograms that tend to have clusters of pixels with similar colors. Thus, we could select a cluster of pixels, for instance by taking the pixels belonging to a set of histogram bins that sit in-between two local minima in the histogram. The corresponding pixels would exhibit similar colors, and additionally would likely occupy connected regions in the image. Further, within those selected bins we can calculate the mean bin-count as well as its standard deviation. This can be done separately in the example and target images, allowing the bin-counts in the selected region of the target image to be adjusted to have the same mean and standard deviation as the corresponding region in the example image. This basic operation can be repeated for the entire histogram, assuming that each bin belongs to exactly one cluster.

This histogram manipulation can be extended in a hierarchical fashion. Effectively, the histogram can be blurred by some amount, leading to a new histogram with fewer local minima. The number of times the histogram is smoothed gives control over the amount of color transfer that will be applied, leading to a progressive algorithm [18] that proves to be more robust than both histogram matching and the original color transfer algorithm [19].

Color transfer was originally developed to allow rendered images to be post-processed, leaving them appear more natural [19]. Creative uses are apparent, for instance sunset images can be made more vivid by transferring colors of a suitably chosen second sunset image. It can also be used to create convincing night time images, although a simulation of subjective noise and loss of visual acuity may further enhance the simulation [25]. Conversely, it has also been used to create images with a convincing day-time color appearance from false color multi-band night-time imagery [22, 23, 26], as well as in remote sensing applications [13].

Further applications include the correction of stereo pair images. Creating stereo images and video requires a set of cameras that likely have a slightly different response and therefore produce images with somewhat different tonal characteristics. This can be corrected with color transfer techniques, matching one of the two images to the other. Similarly, photos comprising a panorama may have somewhat different color characteristics if they were taken with automatic settings. Correcting for these differences can be achieved with color transfer prior to stitching the images together.

Tone Management

Photographs made by professional photographers often have many other intentionally applied characteristics, such as contrast, sharpness and graininess. Such qualities can also be gleaned from example photographs, although this requires somewhat more involved machinery.

In particular, to control tonal balance as well as the amount of detail present in the output image, an example image with the required characteristics can be selected, followed by a decomposition into two layers, one with the overal tonal variation and the other containing fine details [5]. Computing the tonal variation from an image can be achieved by means of bilateral filtering [7], which is an example of an edge-preserving smoothing operator. As such, it effectively blurs the image without destroying sharp edges. The result of this procedure is called the base layer. The detail layer can then be computed as the ratio between the input image and the smoothed image.

In this decomposition, the base layer contains course scale tonal variation, which can be transferred from the example image to the target image by means of histogram matching. The detail layer contains all remaining high frequency detail. As high frequency detail is not exclusively located in the detail layer, both layers are used to compute a measure of textureness [5] by means of high pass filtering. This is done for both target and example images. The target image is then modified by applying histogram matching to the textureness maps.

The two resulting base and detail layers are then recombined in to an output image which is then processed further to manage gradients as well as overly dark or bright areas. Final post-processing filters may include sharpness management, film grain modelling and color adjustments.

In the case that a near-infrared (NIR) photograph as well as a conventional photograph of a scene is available, the conventional photograph can be enhanced with the texture and contrast information from the NIR image [31]. This has the effect of salvaging information lost in under- and over-exposed regions.

An alternative approach to tone management uses steerable pyramids to analyse images [4]. A steerable pyramid [8] is a hierarchical stack of band-pass filtered images, generated with anisotropic filters applied in a set of different orientations. The resulting coefficients can be analysed over different bands. The average magnitude of the coefficients in each of the bands can be plotted against scale, effectively producing a crude characterization of the image. Moreover, these statistics can be transferred from one image to another, allowing for instance particular photographic or painterly styles to be transfered [4].

Image Colorization

A related application area where example-based techniques have contributed is that of image colorization. Here, the task is to color a grey-scale image or video. This is an under-constrained problem in that color information requires three values per pixel, whereas only one value per pixel is given. To begin to solve this problem, it can be observed that in color images regions with a specific set of colors (grass, sky, etc.) also has a specific texture. Thus, textures and colors are correlated.

This means that in the grey-scale target image a given pixel's local environment can be statistically analysed. In the color image that serves as the example, a region with similar statistics can then be found [27]. The statistics used in this case consists of the mean luminance of a 5×5 pixel neighborhood. To limit the search space and make the algorithm tractable, a small set of sample points is chosen in the example image, located on a jittered grid. Retaining the luminance value in the target image, color information is added by transferring the color of the pixel selected in the example image.

An extension that improves robustness first segments the grey-scale target image which helps to disambiguate semantically different regions that may have similar luminance distributions [6]. As an example, it helps with coloring photographs of human faces, where the segmentation ensures that lips can be colored differently from the rest of the skin.

Material Assignment

All techniques and applications outlined so far relate to the processing of images and video. However, images can also be analysed to infer the material description of objects depicted in these images [15]. This is useful when a 3D model of a scene needs to be quickly augmented with sensible materials. In this approach an image or a video is analysed to infer a set of material descriptions that are likely to be present in the image/video. These materials are then assigned to objects in the 3D scene. This assignment is done via an optimization scheme that ensures that materials are assigned in a perceptually plausible manner.

Thus, technical directors are aided by a simple approach, requiring the selection of an image that conveys the mood of the scene. Its materials are extracted and then assigned to the 3D scene, which of course allows further edits to refine the overall appearance.

Summary

Whenever the appearance of a photograph or a video is not quite as the owner imagined it, there are several approaches available to change it, often involving adjustment of sets of parameters, or directly painting over pixels. A relatively new and powerful concept has recently emerged, which involves adjusting an image according to an exemplar. This allows color content to be created, adjusted or generally the mood or feel of an image can be altered to match a desired outcome. Rather than dragging sliders, the process involves the selection of an example image, after which the desired attributes are applied to the target image.

Aside from creative applications, this approach has found many technical uses which are listed throughout this paper, ranging from correcting stereo pairs to remote sensing and image fusion. Although color transfer was devised as a toy technique that allows holiday snapshots to be improved by snatching the color mood from renowned paintings [19], by now all manner of variations have found their way into practical applications.

References

[1] A Abadpour and S Kasaei. A fast and efficient fuzzy color transfer method. In *Proceedings of the 4th IEEE International Symposium on Signal Processing and Information Technology*, pages 491–494, Washington, DC, USA, 2004. IEEE Computer Society.

[2] A Abadpour and S Kasaei. An efficient PCA-based color transfer method. *Journal of Visual Communication and Image Representation*, 18:15–34, 2007.

[3] X An and F Pellacini. User controllable color transfer. *Computer Graphics Forum*, 29:263–271, 2010.

[4] S Bae and F Durand. Statistical analysis and transfer of pictorial styles. In *Proceedings of the 2004 MIT Student Oxygen Workshop*, Ashland, MA, 2004.

[5] S Bae, S Paris, and F Durand. Two-scale tone management for photographic look. *ACM Transactions on Graphics*, 25(3):637–645, 2006.

[6] T Chen, Y Wang, V Schillings, and C Meinel. Grayscale image matting and colorization. In *Proceedings of the Asian Conference on Computer Vision (ACCV 2004)*, pages 1164–1169, 2004.

[7] D Comaniciu and P Meer. Mean shift: a robust approach toward feature space analysis. *IEEE Transactions on Pattern Analysis and Machine Intelligence*, 24(5):603–619, 2002.

[8] W T Freeman and E H Adelson. The design and use of steerable filters. *IEEE Transactions on Pattern Analysis and Machine Intelligence*, 13(9):891–906, 1991.

[9] M Grundland and NA Dodgson. The decolorize algorithm for contrast enhancing, color to grayscale conversion. Technical Report UCAM-CL-TR-649, University of Cambridge, 2005.

[10] A Hertzmann, C E Jacobs, N Oliver, B Curless, and D Salesin. Image analogies. In *Proceedings of SIGGRAPH 2001*, pages 327–340, 2001.

[11] H Kotera. A scene-referred color transfer for pleasant imaging on display. *IEEE International Conference on Image Processing*, 2:5–8, 2005.

[12] H Kotera, T Morimoto, and R Saito. Object-oriented color matching by image clustering. In *Proceedings of the 6th Color Imaging Conference*, pages 154–158, 1998.

[13] Zhenhua Li, Zhongliang Jing, Xuhong Yang, and Shaoyuan Sun. Color transfer based remote sensing image fusion using non-separable wavelet frame transform. *Pattern Recognition Letters*, 26:2006–2014, 2005.

[14] Q Luan, F Wen, and Y-Q Xu. Color transfer brush. In *Proceedings of Pacific Graphics*, pages 465–468, 2007.

[15] C H Nguyen, T Ritschel, K Myszkowski, E Eisemann, and H-P Seidel. 3D material style transfer. *Computer Graphics Forum (Proceedings Eurographics 2012)*, 2(31), 2012.

[16] F Pitié, AC Kokaram, and R Dahyot. N-dimensional probability density function transfer and its application to colour transfer. In *ICCV '05: Proceedings of the 2005 IEEE International Conference on Computer Vision*, volume 2, pages 1434–1439, Washington, DC, USA, 2005. IEEE Computer Society.

[17] F Pitié, AC Kokaram, and R Dahyot. Automated colour grading using colour distribution transfer. *Computer Vision and Image Understanding*, 107(2):1434–1439, 2007.

[18] T Pouli and E Reinhard. Progressive color transfer for images of arbitrary dynamic range. *Computers and Graphics*, 35(1):67–80, 2011.

[19] E Reinhard, M Ashikhmin, B Gooch, and P Shirley. Color transfer between images. *IEEE Computer Graphics and Applications*, 21:34–41, September/October 2001.

[20] E Reinhard and T Pouli. Colour spaces for colour transfer. In R Schettini, S Tominaga, and A Tremeau, editors, *IAPR Computational Color Imaging Workshop*, volume LNCS 6626 of *Springer Lecture Notes in Computer Science*, pages 1–15, 2011.

[21] 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 A*, 15(8):2036–2045, 1998.

[22] S Shi, L Wang, W Jin, and Y Zhao. Color night vision based on color transfer in yuv color space. In *Proceedings of the SPIE*, volume 6623, page 66230B, 2007.

[23] S Sun and H Zhao. Perceptual evaluation of color night vision image quality. In *Information Fusion, 10th International Conference on*, pages 1–7, 2007.

[24] Y Tai, J Jia, and C Tang. Local color transfer via probabilistic segmentation by expectation-maximization. In *IEEE Computer Society Conference on Computer Vision and Pattern Recognition (CVPR 2005)*, volume 1, pages 747–754, Washington, DC, USA, 2005. IEEE Computer Society.

[25] W B Thompson, P Shirley, and J Ferwerda. A spatial post-processing algorithm for images of night scenes. *Journal of Graphics Tools*, 7(1):1–12, 2002.

[26] A Toet. Natural colour mapping for multiband nightvision imagery. *Information Fusion*, 4(3):155–166, 2003.

[27] T Welsh, M Ashikhmin, and K Mueller. Transferring color to greyscale images. *ACM Transactions on Graphics*, 21:277–280, 2002.

[28] Y Xiang, B Zou, and H Li. Selective color transfer with multi-source images. *Pattern Recognition Letters*, 30:682–689, 2009.

[29] X Xiao and L Ma. Color transfer in correlated color space. In *VRCIA '06: Proceedings of the 2006 ACM international conference on Virtual reality continuum and its applications*, pages 305–309, 2006.

[30] S Xu, Y Zhang, S Zhang, and X Ye. Uniform color transfer. In *IEEE International Conference on Image Processing*, pages III – 940–943, 2005.

[31] X P Zhang, T Sim, and X P Miao. Enhancing photographs with near infrared images. In *IEEE Conference on Computer Vision and Pattern Recognition*, pages 1–8, 2008.

Author Biography

Erik Reinhard received his Ph.D. in Computer Science from the University of Bristol in 2000, having worked on his Ph.D. for three years at Delft University of Technology, prior to a further three years in Bristol. Following a post-doctoral position at the University of Utah (2000-2002) and assistant professor at the University of Central Florida (2002-2005), he returned to Bristol as a lecturer in January 2006 to become senior lecturer in 2007. In 2012 he started work as senior researcher at the Max Planck Institute for Informatics in Saarbrücken. Erik founded the prestigious ACM Transactions on Applied Perception, and has been Editor-in-Chief since its inception in 2003, until early 2009. He is lead author of two books: 'High Dynamic Range Imaging: Acquisition, Display, and Image-Based Lighting' and 'Color Imaging: Fundamentals and Applications'. He was programme chair of various workshops and symposia, most recently the Eurographics Workshop on Rendering (2011). He was keynote speaker for the Eurographics Conference in 2010, as well as the Computational Color Imaging Workshop in 2011. His interests are in the application of knowledge from perception and neuroscience to help solve problems in graphics, image processing and color science.