

A Survey of Color Mapping and its Applications

H. S. Faridul^{1,2}, T. Pouli¹, C. Chamaret¹, J. Stauder¹, A. Tremeau², E. Reinhard¹

¹Technicolor Research & Innovation, France

²University of Saint Etienne, France

Abstract

Color mapping or color transfer methods aim to recolor a given image or video by deriving a mapping between that image and another image serving as a reference. This class of methods has received considerable attention in recent years, both in academic literature and in industrial applications. Methods for recoloring images have often appeared under the labels of color correction, color transfer or color balancing, to name a few, but their goal is always the same: mapping the colors of one image to another. In this report, we present a comprehensive overview of these methods and offer a classification of current solutions depending not only on their algorithmic formulation but also their range of applications. We discuss the relative merit of each class of techniques through examples and show how color mapping solutions can and have been applied to a diverse range of problems.

Categories and Subject Descriptors (according to ACM CCS): I.4.8 [IMAGE PROCESSING AND COMPUTER VISION]: Scene Analysis—Color

1. Introduction

Color is an integral part of our visual world and one of the main features of images used in art, photography and visualization for relaying information, or conveying a specific mood. By modifying the colors in an image, it is possible to alter the overall emotion of a scene, to simulate different illumination conditions or to achieve different stylistic effects.

In many cases, color manipulation may also be necessary to reduce differences between images for further processing. When stitching a panorama for instance, consecutive images may have slight color variations, hindering the stitching process. Similarly, differences between individual camera sensors may lead to small changes across a stereo pair that could affect viewer comfort. In another scenario, when processing video content, color edits applied to one frame often need to be replicated to subsequent frames of a sequence.

Both in creative and in more practical scenarios such as these, editing the color content of images requires skilled and extensive user input, while the tools available to non-expert users tend to not offer adequate control.

Color mapping or color transfer is a class of techniques that aims to provide a simple way of achieving complex color changes in images by allowing the color palette and possibly other properties of an image to be altered using a second image as a reference, as shown in Figure 1 (note




Figure 1: Color mapping methods at a high level can be described as a function that maps the colors of a reference image to the input.

that the terms color transfer and color mapping are used interchangeably throughout this report). The user can select a reference image whose colors are preferred and modify the original image such that it acquires the palette of that reference.

Color mapping has received a lot of attention within the computer graphics, computer vision and image processing communities in recent years, both because of its conceptual simplicity and the wide variety of solutions that it can employ. Applications of color mapping vary from making the appearance of renderings more realistic, to tonemapping and panorama stitching, with examples even within security and medical imaging. Yet, it is not clear which methods are best suited for which purposes or even what can really be achieved with this class of solutions. Despite the large number of available methods, there are still open and interesting

research questions and challenges in this area that could help color mapping or color transfer solutions reach their full potential.

In this state of the art report, we aim to provide a comprehensive review of color mapping methods and their applications. The objectives of this report are twofold: first, we aim to categorize existing color mapping techniques according to their algorithmic formulation. Second, we discuss the main application areas where color transfer has been employed and analyze the suitability of different methods in the context of different applications.

1.1. Goals and Challenges

At a high level, the goal of color transfer solutions is always the same, namely to change the colors of a given input image to match those of a reference. Yet, very different solutions have emerged depending on the type of input expected or the specific requirements of each application.

If the two images contain similar scenes, or even the same scene captured under different conditions, a desired mapping is conceptually easy to define: matching objects and regions between images should obtain the same colors. For example, multiple-view camera grids might have color differences due to a lack of color calibration or due to the use of automatic camera settings [CMC10, DN09]. These can be corrected by transferring the colors of one image to the other(s). Similarly, images intended for panorama stitching may be captured with varying camera settings, leading to exposure or color differences [BL07]. Although in these examples the differences between image pairs are likely to be small, it is the quantity of data that necessitates the use of color mapping rather than manual adjustments here.

If a more general mapping is our goal, such as for instance transferring the color palette of a particular painting of a given artist to another image, an almost infinite space of solutions is available, many of which could be considered successful. Unlike the first case, where the goal is to achieve color consistency between similar views, the objective here is to transfer the overall color palette from one image to another. For instance, the style of a particular film may be transferred to other content [BSPP13, XADR13] or a photograph may be used to improve the appearance of rendered content [RAGS01].

In either of the above scenarios, color transfer methods take the same general form. Once the input and reference images are selected, color correspondences are determined between them, defining a mapping between colors in the target image and the selected reference. The mapping is then used to modify the input image.

In the specific case that images have semantic similarities, it is possible to determine exact feature correspondences between objects that are present in both images. To that end,

correspondences by geometric local features are often leveraged in those cases [DN09, HSGL11, FST12].

On the other hand, if the source and reference images depict different scenes, automatic feature correspondences are not attainable, and therefore user-assisted [LLW04, WHCM08] or statistically determined [XM06, PR11] correspondences may be used. In the former case, the user can indicate region correspondences using strokes [LLW04, WHCM08, AP10] or swatches [RAGS01]. Alternatively, a mapping may be determined directly in the color distributions with no specific considerations of image structure [XM06, PKD07, PR11] or a form of pixel clustering or segmentation may be employed to provide an initial grouping of “like for like” pixels [ICOL05, WYW*10].

Color correspondences between images typically only determine how a subset of the image colors should be transformed. The second challenge in color mapping is how to model the complete transformation from the input to the reference palette. Simpler changes, such as remapping of the illumination or overall intensity changes can be achieved through simpler, linear models [RAGS01, TGTC02], whereas more complex changes are likely to require non-linear models [HSGL11].

Although more complex models are able to encode larger changes in color palette between the reference and input images, they also risk creating inconsistencies or discontinuities in the processed image. If the gradient in a smoothly changing region is for instance increased, banding artefacts may appear [XM09]. To avoid this, regularization constraints may be imposed on the transfer [HSGL11] or image regions may be selected based on a soft segmentation that avoids strong boundaries [WYW*10].

Most color transfer algorithms take as input two images, one exemplar and one image to be modified according to the exemplar. Sometimes it may be possible to find an image that is already transformed in a desirable manner. It is then possible to apply machine learning techniques to encode the transform, and subsequently apply the same transform to other images [HJO*01]. Although this is a generally applicable technique, suitable not only for color but also for other image modalities such as texture, in most applications discussed in this paper exemplars are easy to find, but transform pairs are more difficult to come by.

A final challenge related to color transfer, which is perhaps separate from the algorithmic considerations discussed so far, is that of evaluating color transfer results. Because of the under-constrained nature of this problem and therefore the wide range of plausible solutions given a pair of input images, a quantitative assessment of a color transfer result is a difficult task. This important issue is discussed in depth in Section 4, while algorithmic considerations are explored in the following section.

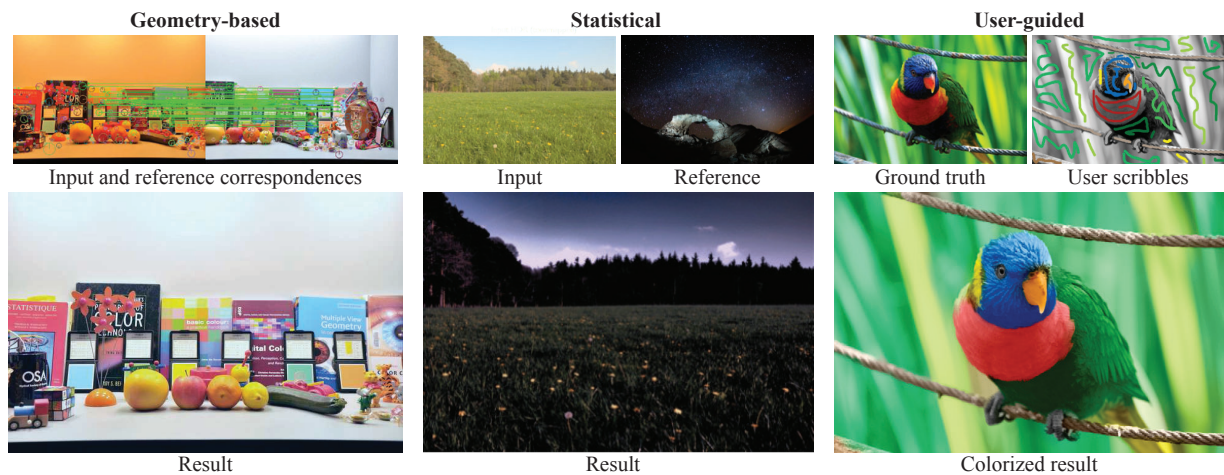


Figure 2: A broad classification of color transfer approaches. Geometry-based methods are discussed in Section 2.1, statistics-based algorithms are presented in Section 2.2 and user-assisted approaches are introduced in Section 2.3.

2. Classification of Color Mapping Methods

Here, we provide an overview and categorization of color mapping methods. The color mapping problem can be approached from widely different perspectives that stem from computer vision, image processing or computer graphics, leading to very different goals and requirements along with matching algorithmic solutions. We find that most color transfer algorithms broadly fit into one of three categories, namely geometry-based, statistical as well as user-assisted methods (see Figure 2). These are discussed in the following sections.

2.1. Geometry-based Methods

Although a precise definition is elusive, in some sense many color transfer methods aim to match the appearance between images. For some applications such as image stitching and stereo capture, this could be sets of images that were taken of the same scene. In this case, the transfer of aspects of color rendition from one image to another can be facilitated by searching for corresponding features that are depicted in both images. By actively finding correspondences between pairs of images, the color transfer algorithm can better ensure that features that occur in both images end up having the same colors.

2.1.1. Sparse Correspondence

To determine which features occur in multiple images, feature detection methods can be successfully employed [Lin98]. These include often-used algorithms such as Scale-Invariant Feature Transform (SIFT) [Low04] and Speeded-Up Robust Features (SURF) [BTVG06]. Such feature detection methods are applied to both input images, leading to two sets of feature vectors. Features from both sets can

then be matched to each other to find a set of candidate correspondences [BL97]. This candidate set can be refined by applying the Random Sample Consensus (RANSAC) algorithm [FB81], effectively rejecting outliers. Many variants to this basic structure are of course possible. Further, taking the colors directly from the matched feature points may lead to issues with robustness. Instead, it may be beneficial to consider neighborhoods around each feature point to derive the transfer function [FST12].

Once features are detected and a correspondence is determined, color differences can be compensated. For instance, Yamamoto et al. use SIFT to calculate spatial correspondences [YYF*07, YO08a] and subsequently apply a Gaussian convolution kernel to the images to detect corresponding colors. The use of blur kernels (or other forms of windowing [OHSG12]) improves robustness for instance against noise pollution. Corresponding colors are then encapsulated into look-up tables that are subsequently applied to the target frames. This and similar techniques [TISK10] could be applied to each of the three color channels independently. An example application where this technique could be used is multi-view video coding [YKK*07].

Alternatively, sparse spatial correspondences have equivalent color correspondences in a given color space. Such color correspondences can be seen as spanning a vector field, for instance in the CIELab color space [OHSG12]. As this vector field will be sparse, a denser but still smooth vector field can be derived to guide color transfers for colors that are not close to the ones for which an exact transform is defined. To create a dense vector field in color space, each color correspondence vector in the sparse field is interpolated using radial basis function interpolation. It was found that normalized Shepard basis functions are better suited for

this task than Gaussian or inverse quadratic basis functions [OHSG12].

2.1.2. Region-based Correspondence

Rather than rely on point-wise feature correspondences, it is possible to apply a color transfer method after matching regions in images, possibly improving the robustness of the results. To this end, the image is first segmented [WSW10, SJYC07], for instance with a mean-shift based image segmentation technique [CM02], followed by feature detection using an optical-flow based algorithm [WSW10].

Alternatively, the images can be decomposed using quad-trees [KMHO09]. A consistency measure is defined in terms of color. If a node in the quad-tree is deemed inconsistent, it is recursively split. Testing corresponding nodes in the quad-trees belonging to each image allows the definition of a consistent region-based set of mappings. Each pair of corresponding regions has its own mapping, which is effected by histogram matching. To reduce the number of such mappings, similar mappings are clustered [HTF01]. Finally, the quad-tree decomposition gives rise to blocky artefacts after the mapping, which requires an additional refinement process.

This method can handle differences in illumination, but assumes that the viewpoint is identical between the two images [KMHO09]. In case both the viewpoints and the illumination differ, the method can be augmented with feature detection using SIFT/RANSAC, as discussed above.

2.1.3. Dense Correspondence

The methods discussed so far in this section, each compute either sparse correspondences or region-based correspondence. Such approaches thus help to control the computational load. However, sparse feature matching does make the implicit assumption that the geometric differences between two images are relatively small, and that the transform between the two images is rigid.

However, the differences between images could be due to objects or humans moving and deforming. The viewpoints of two images of the same scene could also be significantly different, for instance in applications whereby images in photo collections are to be made more consistent [HSGL11, HSGL13]. Such cases are less well served by color transfer based on sparse feature matching.

This has given rise to techniques that compute dense correspondences. A well-known algorithm to compute dense approximate nearest neighbor correspondences is the Patch-Match algorithm [BSFG09]. This algorithm has been extended to include n nearest neighbors, search across different scales and rotations, and to match with the aid of arbitrary descriptors and distances. This algorithm is known as the Generalized PatchMatch algorithm [BSGF10].

In turn, this algorithm was extended to deal with tonal

differences across the images for which correspondences are sought [HSGL11]. Typically, dense correspondences are found only for parts of the images. These correspondences can be used to define a global parametric transfer curve for each of the three color channels (in RGB space, in this case). Such a global transfer function ensures that colors in the input image that do not have a counterpart in the reference image, can still be adjusted in a meaningful and coherent manner.

In addition, a matrix with a single degree of freedom is used to scale colors relative to gray. This captures saturation changes between images that cannot be accounted for by having three independent parametric curves. This matrix has the form:

$$\begin{bmatrix} s - w_r & w_g & w_b \\ w_r & s - w_g & w_b \\ w_r & w_g & s - w_b \end{bmatrix} \quad (1)$$

Here, optimization is used to find the scale factor s that best models the corresponding chroma values. This optimization can be repeated for two common estimates of gray: $(w_r, w_g, w_b) = (0.333, 0.333, 0.333)$ and $(w_r, w_g, w_b) = (0.299, 0.587, 0.114)$, the latter stemming from YUV color space.

With this approach, dense correspondences enable a more precise transfer of color than sparse correspondences (obtained with SIFT), as shown in Figure 3.

The method was recently extended to deal with multiple photographs such as collections of images of a single event [HSGL13]. In this case, a possibly large number of photographs need to be made consistent. This problem can be tackled by super-imposing a graph structure on the set of images. Images that have content in common will then be connected in the graph, whereas disjoint images will be connected via multiple vertices. This allows adjustments to one image to be propagated across the graph to affect nearby images more than images further along the graph. Once again, parametric curves can be used to carry out the transfer of color between images.

For rigid scenes and significant overlap between two views of the same scene, dense correspondences can also be computed by estimating global disparity. For instance, assuming that the two views form a stereo image pair, the global disparity can be calculated by first transforming both images into rank transform space [ZW94]. The rank transform is computed by considering a window around each pixel, and calculating the number of pixels within this window that have a value less than the value of the center pixel. The result is then used in an optimization that determines the optimal offset between the two rank transformed images:

$$\operatorname{argmin}_d \sum_x \sum_y |r_i(x, y) - r_r(x + d, y)| \quad (2)$$

where r_i and r_r are the input and reference images, respectively [CMC10]. This procedure defines how and where the



Figure 3: Input and reference images (top row) are processed with HaCohen’s color transfer algorithm [HSGL11] using dense correspondences (bottom left) as well as sparse correspondences computed using SIFT features (bottom right). Images used with the kind permission from the authors [HSGL11].

two frames overlap, leading to dense correspondences between the two frames. Block-based methods can also be used to determine the disparity between two views [DN09]

2.2. Transferring Statistical Properties

When direct correspondences between image features are not available, statistical properties are often used to define a mapping between the two input images. Statistical descriptors have been used extensively in visual computing as they provide a compact means of describing tendencies in images or image classes [PRC13].

Color information in images is typically encoded using three numbers per pixel, representing a point within a 3-dimensional color space. By reshaping the color distribution of the input image such that it approaches that of the reference, the colors can be transferred between the two images. To simplify the transfer process, statistical properties that describe this distribution can be used as a proxy instead.

In most imaging applications and indeed in our visual system [RKAJ08], images are encoded using RGB-like color spaces, where, broadly speaking, each channel encodes the amount of each primary present within a given color. Information encoded in the three channels is highly correlated, though; effectively changes in one channel affect the values in the other channels. Although this property is necessary for human vision to allow us to perceive color [WS00], it also means that manipulations of the color content of images in such a color space may have unpredictable effects.

Fortunately, the space used for manipulating the color content of images need not be the same as the color space spanned by human photoreceptors. In fact, the same color can be described in infinitely many ways by changing the set of primaries that it is defined upon. Mathematically, this can be achieved through a transformation that takes a given color from one set of coordinates to another. By shifting, scaling, and rotating the axes defining a color space, a different space can be constructed to achieve different goals [RKAJ08].

In the context of color transfer, the correlation between the channels of the RGB space is of course highly undesirable. By choosing a less correlated color space, many statistical-based methods can map colors between images by considering each channel separately, and therefore converting a potentially complex 3D problem into three separate 1D transfers [RAGS01, GD05], [XM06, SA07], [AK07, XM09], [PR10, PR11]. In these methods, the selection of color space plays a crucial role [RPer], as discussed in Section 2.2.2.

2.2.1. Per-Channel Transfer

One of the earliest color transfer techniques takes advantage of the decorrelation property of a perceptually-based color space known as $l\alpha\beta$ [RCC98] and transfers simple statistical moments (mean and standard deviation) between each channel of the two images. An example result is shown in Figure 5. To achieve this, the values of each channel of the input image I_i are shifted to a zero mean by subtracting the mean of the input from each pixel. They are then scaled by the ratio of standard deviations of both images such that they acquire that of the reference image, and finally, they are shifted to the mean of the reference by adding its mean instead of the input mean that was originally subtracted:

$$I_o = \frac{\sigma_r}{\sigma_i}(I_i - \mu_i) + \mu_r \quad (3)$$

Here, the subscripts i , r , and o correspond to the input, reference and output images and μ and σ are their respective means and standard deviations. Note that this process is applied on each channel of the image separately after converting to the $l\alpha\beta$ color space.

Despite its simplicity, this technique can be successful for a wide range of images, but the quality of the results relies on careful selection of the reference image. Additionally, be-

cause the $l\alpha\beta$ color space is constructed so that it decorrelates natural images (i.e. forest scenes, grass etc.), it cannot be guaranteed to successfully decorrelate other scene types.

The latter limitation can be overcome by computing a decorrelated color space for the given input images [GD05, XM06, AK07] using principal component analysis (PCA), and apply color transfer in the resulting color space. In the method by Reinhard et al. discussed previously, colors are transferred from the reference to the input image by first rotating the image data to $l\alpha\beta$, followed by a translation and scaling using the input image statistics, which removes the input color characteristics. The same process is then repeated in reverse to apply the statistics of the reference. Xiao and Ma follow a similar process but replace the rotation to $l\alpha\beta$ transformation to the axes defined by the principal components of each image [XM06]. Formally, this can be expressed as a series of affine matrix transformations:

$$I_o = T_r R_r S_r S_i R_i T_i I_i, \quad (4)$$

where T , R and S represent translation, rotation and scaling matrices for the input and reference images accordingly.

In both of these methods, it is possible to achieve increased control by using swatches that define specific mappings as will be discussed in Section 2.3.2. However, in the methods discussed so far, the transfer of colors between the two images relies on simple, global statistics. Such statistics can provide useful information about overall tendencies in a distribution but more information and ideally higher-order analyses are necessary to capture more subtle variations in images.

A more faithful representation can be achieved by considering the full distribution of values in each channel. Commonly, histograms are employed as a compact means of describing a probability distribution. For a given image I , we define its histogram H with B bins of width V as follows:

$$H = \{(h(1), v(1)), \dots, (h(B), v(B))\} \quad (5)$$

$$B = \left\lceil \frac{\max(I) - \min(I)}{V} \right\rceil \quad (6)$$

$$h(i) = \sum_{p=1}^N P(I(p), i), \quad i \in [1, B] \quad (7)$$

$$v(i) = \min(I) + (i - 1)V \quad (8)$$

$$P(I(p), i) = \begin{cases} 1 & i = \left\lfloor \frac{I(p) - \min(I)}{V} + 1 \right\rfloor \\ 0 & \text{otherwise} \end{cases} \quad (9)$$

The shape of the input histogram H_i can be matched to that of the reference H_r using a process known as *histogram equalization* or more generally, *histogram matching*. First, the cumulative histograms of the source and target are com-

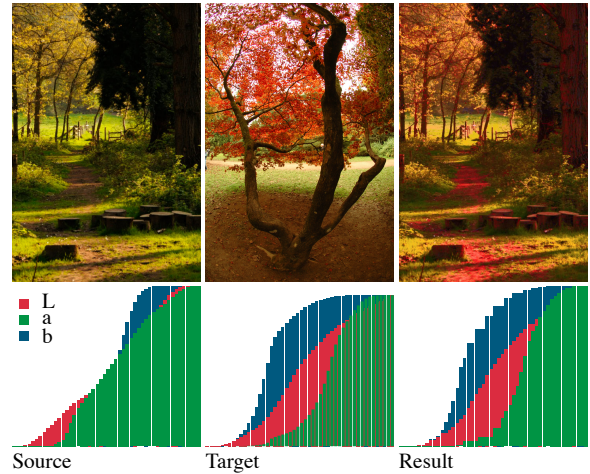


Figure 4: Pixel values in the input image are modified such that the histogram of each channel matches that of the reference image, changing the image appearance of the input to approach that of the reference. In this case the color mapping is done in the CIE Lab color space.

puted:

$$C_i(i) = \sum_{i=1}^B h_i(i) \quad (10)$$

$$C_r(i) = \sum_{i=1}^B h_r(i) \quad (11)$$

after which the input image can be matched to the reference according to the two cumulative histograms:

$$I_0(p) = v_r \left(C_r^{-1} \left(C_s \left(\frac{I(p) - \min(I) + 1}{V} \right) \right) \right) \quad (12)$$

Here, a cumulative histogram C is defined as a function mapping a bin index to a cumulative count. The inverse function C^{-1} acts as a reverse lookup on the histogram, returning the bin index (and therefore the intensity value) corresponding to a given count. This process can be easily extended to a color image by repeating the histogram matching process for each of the three channels to achieve a result as shown in Figure 4.

Using this process, the transformed image will acquire exactly the same distribution as the selected reference and will therefore have the same colors. However, the resulting image may be too harsh as the transfer can amplify artifacts that were previously invisible, indicating that higher-order properties of the image may need to be matched or preserved to achieve a successful result. Based on this observation, a wealth of color transfer methods have been developed, each relying on a different set of statistical properties of the im-




Figure 5: A comparison of three global, statistical color mapping methods.

ages to transfer the color palette between them without otherwise affecting the appearance of the image.

Since a full histogram transfer is likely to be too rigid, one possibility is to only partially match the two histograms by taking advantage of their local structure [SA07, PR11, PR10]. One recent solution achieves that by considering features of the histograms in different scales. An example result from this method is shown in Figure 5. By matching only coarse features of the histogram, a progressive match can be achieved between the two images [PR11, PR10], transferring only large scale color variations without forcing the exact distribution of the reference image.

This leads to an interesting observation: such an approach is possible because nearby pixels in images tend to be similar. This is a recurring property of natural scenes [PRC13] and obviates the need for using segmentation in order to determine corresponding groupings of pixels for color transfer, as will be discussed in Section 2.2.4.

Alternatively, to ensure that the histogram matching process does not introduce artefacts in the image, the transfer can be constrained by aiming to preserve the gradient distribution of the input [XM09] (see Fig 5 for an example result) or a combination of the input image colors and geometry [PPC11]. In the method by Xiao and Ma [XM09], the transfer between the two images can be naturally expressed as an optimization that minimizes the difference between the color histogram of the output and the reference, while keeping the gradient distribution of the output as close as possible to that of the input. Although this is a conceptually simple modification to histogram matching, it leads to visually smoother results. However, as the gradient constraint is global and carries no spatial information, in extreme cases

(such as when transferring between high dynamic range images), it can over-smooth results.

2.2.2. Color Spaces

Statistics-based color transfer algorithms either treat pixel data as a single 3D point cloud in a 3D color space. Color transfer is then a reshaping of this 3D point cloud to statistically match the point cloud of an example image. Several algorithms take a simpler route by mapping pixel values one color channel at a time. Of course, if a mapping is carried out for each channel separately, then the choice of color channels, i.e. the choice of color space, becomes very important.

The argumentation as to what constitutes a good color space for this type of problem is as follows [RPer]. Natural image statistics is the study of statistical regularities in images, and usually aims to help understand how the human visual system operates [HHH09, PCR10, PRC13]. Human vision has evolved in the presence of a natural environment, and is therefore likely to be specialized to observe and understand natural images.

This was tested by Ruderman et al. [RCC98] who converted a set of spectral natural images to LMS cones space, a color space that resembles the responses of the L, M and S cones in the human retina. This image ensemble was then subjected to Principal Component Analysis (PCA), showing that the three principal components obtained closely correspond to the opponent color channels found in the retina [RCC98]. As PCA yields components that are maximally decorrelated, the retina appears to decorrelate its input. While formally PCA only decorrelates, it was found that transforming natural images in this manner, the resulting color space (called $L\alpha\beta$) yields channels that are close to independent.

As $L\alpha\beta$ is a non-standard color space [RCC98], we give

the transformation from XYZ here. The first step is to transform to LMS cone space:

$$\begin{bmatrix} L \\ M \\ S \end{bmatrix} = \begin{bmatrix} 0.3897 & 0.6890 & -0.0787 \\ -0.2298 & 1.1834 & 0.0464 \\ 0.0000 & 0.0000 & 1.0000 \end{bmatrix} \begin{bmatrix} X \\ Y \\ Z \end{bmatrix} \quad (13)$$

The LMS values are then logarithmically compressed to reduce skew in the data before applying the linear transform to $L\alpha\beta$:

$$\begin{bmatrix} L \\ \alpha \\ \beta \end{bmatrix} = \begin{bmatrix} \frac{1}{\sqrt{3}} & \frac{1}{\sqrt{3}} & \frac{1}{\sqrt{3}} \\ \frac{1}{\sqrt{6}} & \frac{1}{\sqrt{6}} & -\frac{1}{\sqrt{6}} \\ \frac{1}{\sqrt{2}} & -\frac{1}{\sqrt{2}} & 0 \end{bmatrix} \begin{bmatrix} \log L \\ \log M \\ \log S \end{bmatrix} \quad (14)$$

This color space is characterised by a luminance channel L and two opponent color channels α and β , which encode yellow-blue and red-green chromaticities.

Of course, if the three color channels of an image can be made (near-) independent, then image processing can take place in each of the three channels independently [RAGS01]. The success of this approach, however, depends strongly on the choice of color space. A better decorrelation between the channels tends to yield a better visual quality of the results. See Figure 6 for an example.

Although $L\alpha\beta$ has been adopted as the space of choice for many color transfer algorithms [RAGS01, GH03],[Toe03, YC04],[LJYS05, XZZY05],[WHH06, LWCO*07],[ZWJS07, XM09, XZL09], related spaces such as CIELAB are also used [CUS04, WHCM08],[PR10, PR11]. Other color spaces implicated in color transfer are CIE LCh* [NN05], $YCbCr$ [KM07], Yuv [LLW04, WZJ*07], as well as color appearance models [MS03]. On the other hand, several authors suggest to compute a dedicated color space for each transfer, based on PCA [KMS98, AK04, Kot05, XM06, AK07].

To determine the relative merit of each color space, a study was performed [RPer], comparing the visual quality obtained with different color spaces using the same color transfer algorithms. Although the details of this study are beyond the scope of this report, the main conclusion is non-trivial. Their finding is that the best results are obtained when the CIELAB color space is used with illuminant E. This produces on average better results than all other spaces that were tested. Notably, although both CIELAB and $L\alpha\beta$ are color opponent spaces, CIELAB outperforms $L\alpha\beta$. Surprisingly, CIELAB with illuminant E also outperforms the use of per-ensemble color spaces, derived with principal components analysis. As such, this color space was recommended for use in statistics-based color transfer algorithms that apply transforms to each color channel independently.

2.2.3. Color Transfer as a 3D Problem

As we have seen, color spaces describe a 3-dimensional space. Pixels in an image are in that case given as a triplet, denoting a point within that space. In our discussion so far, we have focused only on techniques that manipulate each of the image channels independently. Although this description of color provides the obvious advantage of simplifying a potentially complex 3D problem to a set of three 1D problems, it cannot capture all subtleties in the color distribution of images.

To maintain local color information and interrelations, the 3D color distribution of the two images needs to be treated as a whole; this is done so that the input 3D distribution will be reshaped to match or approximate that of the reference. Unfortunately, translating processes such as histogram matching or histogram reshaping to more than one dimension is not straightforward and either requires an optimization-based solution or a way to simplify the problem to fewer dimensions.

In the latter case, it is possible to match the 3-dimensional probability distributions describing the two images through an iterative match of 1D projections. This can be achieved by repeatedly selecting random 1D projections [PKD05, PKD07], or by projecting onto known color properties such as hue, lightness and chroma [NN05].

In the method by Pitie et al., at each iteration step the 3D distributions of the input and reference images are rotated using a random 3D rotation matrix and projected to the axes of their new coordinate system [PKD05, PKD07]. Each 1D projection of the input is then matched to that of the reference and the data is transformed back to its original coordinate system. This process is repeated with different rotations until convergence occurs [PKD05, PKD07]. An example result and the effect of an increasing number of iterations can be seen in Figure 7.

2.2.4. Segmentation and Clustering

In the methods discussed so far in this section, the color distribution of the image is considered independently of the image structure. Although these methods are simpler algorithmically, the choice of image used as reference can define the quality and success of the result. If the reference image is too different in structure to the input, global methods such as the ones discussed, cannot guarantee that the color mapping will be successful.

To improve the coherence of color mapping results, the input and reference images can be segmented into regions, allowing for a local mapping between regions to be defined. Although it is possible to directly segment the image pixels using a binary segmentation, possibly in multiple scales [GH03], most methods opt for a fuzzy segmentation.

Such a softer segmentation can be achieved by determining dominant colors in the image and extracting segments




Figure 6: Decorrelation properties of color spaces. The top row shows pixel values for pairs of channels in an RGB color space, taken from the image at the top left. The bottom row shows the same pixels plotted in pairs of channels from the $L\alpha\beta$ space instead. Values in channel in RGB spaces tend to be good predictors of other channels, resulting in an almost diagonal distribution.

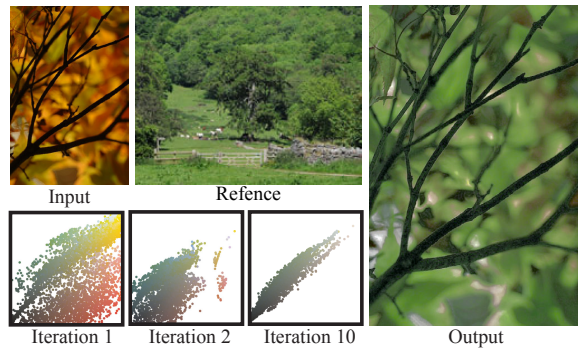


Figure 7: An example result by Pitie et al. [PKD05]. The color distribution of the input image is iteratively reshaped to match that of the reference. The distributions at different iterations are shown for red versus green channels.

for each of them. Yoo et al. [YPCL13] use the mean-shift algorithm [CM02] to that end, and rely on a simple statistical transfer [RAGS01] to transfer colors between corresponding regions.

More commonly, a soft segmentation according to dominant colors in the image is achieved by applying Gaussian Mixture Models (GMM), probabilistic models that aim to describe a given distribution using a mixture of Gaussian distributions [Rey09]. A pixel (x, y) is assigned a probability $P_i(x, y)$ of belonging to a given Gaussian distribution $i \in M$,

computed as:

$$P_i(x, y) = \frac{\exp\left(-\frac{(I(x,y)-\mu_i)^2}{2\sigma_i^2}\right)}{\sum_{j=1}^M \exp\left(-\frac{(I(x,y)-\mu_j)^2}{2\sigma_j^2}\right)} \quad (15)$$

where, μ_i and σ_i are the mean and standard deviation of the i th gaussian distribution respectively.

To segment an image into M regions, a set of M Gaussian distributions may be fitted using an optimization procedure known as Expectation Maximization (EM) [DLR77]. Once GMMs are determined for the input and reference image, a mapping can be constructed between the segmented regions. Typically, luminance information in each segment can be used to guide the color mapping [XZZY05].

To smooth the color mapping result, the EM procedure can be modified to allow a given pixel to fit within more than one GMM, resulting in more seamless transitions between adjacent regions [TJT05, TJT07]. As this method relies on a mapping being available between all regions of the two images, multiple references may be used to ensure that all input regions are assigned a corresponding reference [XZL09].

Instead of clustering images based on their content, segmentation can be guided according to more general color categories. Several studies within color vision and perception have determined that colors can be divided into a small number of almost universal categories [Ber69, Hei72, UB87]. In the context of color transfer, color categorization can be used to both segment images into coherent regions but also to determine correspondences between the input and reference images [CUS04, CSUN06, CSN07b].

In a similar vein, rather than using another image, it is possible to use a simple color palette as a reference [WYW*10]. In that case, no spatial information is available in order to guide how colors should be mapped. To determine a mapping between regions in the input and the colors in the given palette, Wang et al. analyzed a collection of images to establish typical mappings between colors and textures or types of objects (e.g. grass, sky etc.). By segmenting the input image and assigning an object class to each region, an appropriate range of colors could be selected from the given palette.

Alternatively, image content analysis may be used to extract semantically meaningful regions from the image [WDK*13]. Recent image analysis methods can automatically segment images according to the 3D structure of the depicted scene [HEH05, HEH07], while saliency analysis can determine which parts of the image are likely to be important and which serve as background [FWT*11]. Using this information, Wu et al. transfer colors between semantically similar regions (e.g. sky to sky, ground to ground etc.) [WDK*13].

2.3. User-assisted solutions

If the structure and content of the input image is too different from that of the reference or example image, many automatic methods will fail to find a successful mapping, requiring user input to guide the source and reference correspondences. Beyond the goal of achieving better results, user assistance is also necessary additional reasons. Previously mentioned classes of color mapping do not take into account semantic aspects: i.e., pixels associated with specific types of objects such as “sky”, “grass” or “face”, have a restricted range of plausible colors.

To incorporate user control into color mapping methods, one possibility is for the user to manually define layer mask [COSG*06, PR11] or strokes [WHCM08] for preserving regions from any color mapping. On the other hand, manual user interaction may be used to define region correspondences between images or videos [OHSG12, WAM02]. In addition, it may be the case that the user input itself serves as a rough reference or initialization for color mapping through colored scribbles [LLW04] and colored regions [SVS*10].

2.3.1. Stroke-based Approaches

A mapping between images can be defined by a user through a set of strokes that can be painted onto the input and reference images. Wen et al. for instance define correspondences between images, as well as regions that should not be adjusted, in this way [WHCM08]. Similarly, Cohen-Or et al. allow human interaction for avoiding color harmonization of tricky regions or for ensuring that similar regions are grouped together [COSG*06].

Although some approaches do not explicitly deal with




Figure 8: An example result using the stroke-based colorization method of Levin et al. [LLW04]. A greyscale image is supplemented with user-provided strokes (middle), which are propagated to create a colorized result (right). The ground truth is shown on the left.

color mapping they are closely related from a color modeling perspective. For example, Lischinski et al. [LFUS06] allow users to select manual image regions to perform local tonal adjustment. They apply different properties to strokes and brushes to better cope with their problem: *luminance brush* covers all pixels having a similar luminance than the annotated brush stroke, *over-exposure brush* represent all over-exposed pixels enclosed by a manual stroke-based region. In the same vein, An et al. [AP08] use a similar idea by propagating manual tonal adjustment toward spatially-close regions of similar appearance. Later, they extended the method to color transfer problem [AP10].

Instead of transferring colors from another image, strokes can also be used more directly to define the actual color that a region should obtain. The pioneering work of Levin et al. [LLW04] introduced the notion of strokes in the process of adding color to greyscale images, known as the *colorization* problem (discussed in detail in Section 3.1). Stroke-based approaches provide an easy manual way for efficiently guiding the mapping or for initializing the appearance of regions. The main idea behind such methods is to spatially diffuse information from the user provided stroke to the rest of the image.

To achieve that, once strokes are placed by the user, the algorithm by Levin et al. [LLW04] propagates the defined colors to the remainder of the image. The guiding premise is that neighboring pixels with similar luminance should also obtain similar colors, while nearby pixels whose luminance is different are likely to have different colors. This relation can be expressed through a weighting function known as the *affinity function* that assigns weights to neighboring pixels according to their luminance differences.

2.3.2. Swatch-based Approaches

This kind of approach is more related to the issues of finding correspondences between target and reference. Welsh et al. [WAM02] pointed out the issues related to region correspondences in the context of grayscale colorization. Mapping colors based on luminance statistics or structures is a restrictive assumption that makes the problem difficult. Consequently, they first introduced the swatch approach for guaranteeing

color mapping of critical regions. Thus, they applied color transfer or color mapping from the swatched region of colorized pictures to the corresponding one in the greyscale target. Then, instead of searching for texture/luminance correspondences between the target and the reference pictures, they restrict the search area to the target pictures, considering an internal picture correspondence is more probable than an external one.

While Farbman et al. and Li et al. employed local diffusion of annotation through diffusion map [FFL10] and radial basis function interpolation in a feature space [LJH10, OHS12] used rather a global approach for color balancing in formulating the problem as a global color space transformation independent from the image color space. Recently, Xu et al. [XYJ13] improved the required level of requested sample edits and applied their framework to different applications where manual editing is inescapable.

3. Applications

Color mapping methods first appeared as a means to improve the appearance of rendered content [RAGS01], but the flexibility of the general idea soon gave rise to a variety of other applications where automatic but guided color manipulation was necessary. The following sections discuss the main application areas where color mapping has been used.

3.1. Image Colorization

In the scenarios discussed so far, there is an implicit assumption that the target and reference content have the same dimensionality. However, many practical applications require the colorization of greyscale images or videos. In these cases, the input image to be colorized has by nature no color on it and therefore carries less information to guide the mapping of colors from the reference image; nevertheless it is not necessarily a gray-scale picture (meaning a natural image or a pattern), it could be any graphical creation [LRFH13] or a sketch, a cartoon with only delimited regions and without any luminance information [SVS*10].

Colorization of greyscale content is a process that has received attention early on in the history of image processing. The earliest computer-aided colorization method was developed by Markle in the 1980s and relied on a user for identifying regions where motion occurred between frames [Mar84]. More recently, automatic methods have been proposed to colorize greyscale content, either using a color image as a reference [WAM02, ICOL05, CHS08] or simple color palettes [LRFH13]. Alternatively, colorization can be controlled stroke-based user input [LLW04]. Although colorization methods follow similar principles to color-to-color mapping solutions discussed in Section 2, they merit their own discussion as they require additional strategies to deal with the difference in dimensionality between the input and reference images.

3.1.1. Exemplar-based Colorization

One of the earliest solutions to automatically colorize images was introduced by Welsh et al. [WAM02]. Similar to many of the color mapping methods discussed in Section 2.2, to determine a mapping between regions (and therefore colors) of the reference image and those of the input Welsh et al. rely on statistical information in the images. Their method operates in the decorrelated $L\alpha\beta$ color space and selects a reference color according to the similarity of the luminance distribution in a small neighborhood around each pixel. To improve spatial coherence of colorized results, Bugeau and Ta propose to use a Total Variation regularization after the colorization step [BT12].

Although this approach can lead to successful colorization results with a carefully chosen reference image, it often fails when the scene content and composition are not well matched. To improve on some of these limitations, Irony et al. [ICOL05] focused on the spatial aspect to improve its consistency in the colorization process. Given a grayscale image to colorize, they first determine for each pixel which example segment it should learn its color from. This is done automatically using a robust supervised classification scheme that analyzes the low-level feature space defined by small neighborhoods of pixels in the example image. Next, each pixel is assigned a color from the appropriate region using a neighborhood matching metric, combined with spatial filtering for improved spatial coherence. Each color assignment is associated with a confidence value, and pixels with a sufficiently high confidence level are provided as "micro-scribbles" to the optimization-based colorization algorithm of Levin et al. [LLW04], which produces the final complete colorization of the image.

Later, Charpiat et al. [CHS08] colorized greyscale images automatically without any manual intervention. Their method deals directly with multimodality (meaning that the same object may have different colors depending on the image, context etc.) and estimate, for each pixel of the image to be colored, the probability distribution of all possible colors, instead of choosing the most probable color at the local level. They also predict the expected variation of color at each pixel, thus defining a non uniform spatial coherency criterion. Then, they use graph cuts to maximize the probability of the whole colored image at the global level. They work in the CIELab color space in order to approximate the human perception of distances between colors, and use machine learning tools to extract as much information as possible from a dataset of colored examples.

Recently, Bugeau et al. [BTP13] proposed a variational framework to address the colorization problem which allows them to simultaneously solve the candidate selection and color regularization problems through a variational energy minimization. Their approach avoids the arbitrary pre-selection of a single candidate for each pixel. To achieve that, they design an energy minimization process that au-

tomatically select the best color for each pixel from a set of candidates, while ensuring the spatial coherency of the reconstruction, as it enforces neighbor pixels to have similar colors in the final result.

In all methods discussed so far, the selection of a reference image or images still rests upon the user. If the user selects an unsuitable image as a reference for the colorization, the result may not be plausible. To resolve this issue as well as to remove the need for user-provided input, several methods automatically select one or more reference images. Viera et al. [VdNJ*07] rely on content-based image retrieval methods to find a suitable reference image, which can be then used to colorize the greyscale input. To select an appropriate reference automatically they explore several image descriptors such as luminance, gradients and multi-scale analysis.

Taking this concept further, Liu et al. [LWQ*08] find multiple appropriate images online by attempting to register parts of each exemplar with the input. Once a set of references is selected, intrinsic images are computed from them using the method of Weiss [Wei01]. Intrinsic image decomposition separates illumination from reflectance and therefore allows the colorization method to consider the true color of each surface. As this method relies on registration of the input and the selected references, only images of known monuments are considered, restricting the applicability of their technique.

The idea of finding reference images from online sources is extended by Chia et al. [CZG*11], removing the strict reliance on feature correspondences. Although the main contribution of their method lies on the automatic selection of candidate images, the colorization process is also updated to rely on super-pixel based feature correspondences, allowing for a more flexible colorization.

If the goal of the colorization is to propagate user-selected colors from one frame to others in an animation sequence of video, as would be the case when converting greyscale video to color, it is possible to use geometric and structural cues to guide the color mapping process [SBŽ04]. In that case, rather than aim to find corresponding regions based on statistical properties, a much more direct registration between frames can be performed.

3.1.2. Stroke-based Approaches

This class of approaches takes advantage of strokes or scribbles that are manually provided by a user. Typically in such methods, the user draws simple lines on the image(s) to either define correspondences between the input and reference or to initialize the colorization process. The number of required scribbles is variable and depends on the efficiency of the considered methods and the expected performance, however an active research area within stroke-based colorization has been that of minimizing the amount of user input required.

Unlike the exemplar-based methods that start from the large variety of color present in a reference image, the color of the scribbles are the input information for the colorization problem in this case. Even though scribbles provide less information, such as texture or luminance statistics, they ensure locally and spatially the right color correspondence. The pioneering work of Levin et al. [LLW04], presented in Section 2.3, showed how color from scribbles can be propagated to the rest of the pixels through an optimization process. Figure 8 shows an example result using this method. This method was later refined to consider edges [HTC*05] and gradients in the image [LLD07] more explicitly, reducing potential bleeding artifacts across edges.

Yatziv and Chapiro [YS06] introduced a computationally simple, yet effective, approach for colorization. Their method is based on a reduced set of chrominance scribbles defined by a user. Based on the concepts of luminance-weighted chrominance blending and fast intrinsic distance computations, they substantially reduce the complexity and computational cost of previously reported techniques. The (geodesic) intrinsic distance gives a measurement of how “flat” is the flattest curve between two points.

Luan et al. [LWCO*07] were motivated to reduce the amount of user interaction necessary to produce complex, nuanced color images and to handle highly textured images as well as non-adjacent regions of similar texture. To achieve that, a user can draw strokes indicating regions that (roughly) share the same color, which are used to automatically segment the image. Further user input serves to define a color for a few pixels within each region, which is finally propagated to the rest of the segment through a piecewise-linear interpolation within regions and a soft blending at region boundaries.

In most stroke-based methods, the user defined strokes are propagated to the remainder of the image according to luminance information. Although this information can correctly determine continuous regions in natural images, this is not the case in manga (japanese cartoons), where regions often have distinctive patterns and textures to indicate shading, material or motion. Qu et al [QWH06] rely on a statistical approach to define representative features describing each region. Using these features, strokes can be propagated according to texture and pattern similarity in addition to luminance information, allowing for successful colorization of highly textured manga images.

3.1.3. Palette-based Approaches

This category of algorithms uses a set of colors (usually manually defined by a user) as an input to the colorization process, with no further spatial constraints (unlike stroke/scribble approaches). Sauvaget et al. [SVS*10] design a solution to colorize empty or pre-filled sketches. Their solution relies on Itten’s theory regarding color harmony combination to determine how colors should be distributed.

Itten's contrast models global picture harmony regarding the proportion of "basic" colors [Itt61]. The optimal color proportion is adjusted depending on the considered numbers of colors used in the image.

In the same vein, Lin et al [LRFH13] recently proposed a method for the automatic colorization of patterns according to a probabilistically selected color palette. Their approach to pattern coloring is data-driven: given a dataset of example colored patterns from the online community *COLOURlovers*, they learn distributions of color properties such as saturation, lightness, and contrast for individual regions and for adjacent regions. They predict these distributions using discriminative spatial features of the pattern, such as the size and shape of different regions. Finally, they use the predicted distributions together with the color compatibility model by O'Donovan et al. [OAH11] to score the quality of pattern colorings. They introduce three types of scoring and show how functions for unary, pairwise, and global scoring can be combined into one unified model using the framework of probabilistic factor graphs.

3.2. Artistic Color Mapping

Another family of methods within color image processing that requires color mapping is the rendering of artistic effects. In color harmonization [COSG*06, BUC*13], a subset of pixels is identified by means of harmony models as being out of the harmonious range of colors. Then, a strategy for re-mapping such pixels in the harmonious space is applied also involving segmentation algorithms. To some extent, art and non-photo-realistic rendering in computer graphics may be associated to color mapping techniques [Her10]. The color mapping can be performed such that it mimics the painting style of artists.

3.2.1. Color Harmonization

The concept of aesthetic ordering of colors and their separation into primary and secondary colors was first proposed by Goethe in his "Theory of Colors" [vGE40], where colors were arranged according to their hue, forming the color wheel shown in Figure 9a. More recently, Itten [Itt61], one of the main proponents of the Bauhaus movement, developed the concept of the color wheel further by focusing on relations between colors along the color wheel in order to derive harmonic combinations of colors. He established a set of templates determining relative positions within the color wheel, which included one, two or multiple color combinations, either adjacent or on opposite sides of the circle. These harmony templates were further expanded by Matsuda [Mat95] and analyzed by Tokumaru et al. [TMI02]. Some color harmony template examples are shown in Figure 9b.

Several color harmonization algorithms [COSG*06, HT09, TMW10, BUC*13] are based on




Figure 9: Left: the original color wheel proposed by Goethe [vGE40]. Right: The color harmony templates proposed by Matsuda [Mat95].

this theory to recolor or recompose color within an image. The first step consists of determining the harmonious template type and its rotation angle that are the closest to the original image by minimizing a cost function. Then, the colors are transformed so that the colors outside harmonious sectors are mapped inside a harmonious sector. Color segmentation has been identified as a crucial pre-processing before color mapping because visible artifacts can appear when two close colors, eventually associated to the same object, are mapped to two different sectors. Strategies differ in terms of cost function for template selection, color segmentation and color mapping.

A real-time color harmonization for video has been introduced by Sawant et al. [SM08] where a histogram splitting method is employed instead of a graph-cut approach to reduce computation cost. One dedicated template per group of frames is determined to guarantee temporal consistency. Tang et al. [TMW10] perform a foreground/background detection for their dedicated video implementation and apply the same template determination as Sawant et al. [SM08] for a coherent group of frames.

Recently, new harmonic templates using unsupervised machine learning on the Adobe Kuler database have been proposed, consisting of hundreds of thousands color palettes used for creative purposes [PM12]. To evaluate existing models of color harmony and compatibility, O'Donovan et al. [OAH11] analyze a large set of online color palettes both from Adobe Kuler and *COLOURlovers*. Based on their analysis, they construct a model to predict the color compatibility of a given set of colors.

3.2.2. Stylization

Non-photorealistic rendering (NPR) creates imagery that mimics the style of artists. It deals mainly with the concepts of line drawings, painting and cartoon illustration. Only the two last mentioned techniques are relevant to our survey. There is a real discussion on the plausibility of describing an artist's style or intent by a functional mapping such as mentioned in [Her10]. The grail would be to computationally capture the artist's style and reproduce or color map it

automatically on new images. Assuming that is possible to computationally model creative intent to some extent, a second dominant issue appears: how to objectively evaluate the performances of such rendering [Her10]. The NPR field provides an interesting and challenging ground for the aesthetic evaluation question. (see Section 4 for a discussion on evaluating color mapping results).

Pioneer work of Haeberli [Hae90] applied a painting style to a color image by simply sampling a random set of points and placing a brush strokes over them. Impressionist painting style is effectively performed in such simple way. Rendering has been improved by setting the strokes orientation to the normal of the image gradient [Lit97]. One step further, Hertzmann [Her98] proposed by changing adequately the painting parameters to mimic different painting style: impressionist, expressionist, pointillist, "colorist wash".

Other algorithms approach cartoon illustration by performing an effective line detection and drawing using bold edges and large regions of constant color [DS02, KKD09, KLC09].

Previous solutions have been extended to video dimensions leading to temporal issues related to consistency (flickering artifacts) and computational time. Hertzmann and Perlin [HP00] were pioneer in fixing consistency issues when painting video material. They apply painting effect only in regions where the source video is changing. Image regions with minimal changes, such as due to video noise, are also left alone, using a simple difference masking technique. By using optical flow, the painting surface flows and deforms to follow the shape of the world. Hays and Essa [HE04] rather focus on temporal consistency of brush strokes properties that include motion information. All brush stroke properties are temporally constrained to guarantee temporally coherent non-photorealistic animations. Bousseau et al. [BNTS07] worked on video water color effects. Their method involves also optical flow for guaranteeing consistency to texture operations and abstraction steps.

3.3. Video Color Mapping

It should come as no surprise that color mapping solutions are frequently employed for video processing, for instance as part of a larger system for video enhancement [BZS*07]. A crucial step in post-production of movies is a process known as *color grading*, where properties such as the luminance, hue and saturation of each clip are manually adjusted to achieve the desired look and feel. This is often a labor-intensive process as each scene and input source requires adjustment both for visual effect and for consistency. Using color mapping, it is possible to transfer the style of specific sample frames or existing content to the rest of the movie, however not all methods are suited for this scenario.

One of the main challenges faced when mapping colors to a video is ensuring temporal consistency between frames. To

map the colors between two images or video frames, a correspondence needs to be formed between the input and reference palette. In the case of video, the correspondence determined for one frame will not necessarily hold for the next, leading to temporal inconsistencies that might present themselves as flicker and sudden color or luminance changes. To resolve this, special considerations are necessary to enforce temporal coherence or to avoid temporal issues altogether. This however introduces a trade-off between adaptivity to changing content within the sequence and temporal stability, leading to many different solutions.

At the simplest case, a global mapping can be determined that is applied in the same way for each frame. For example, the mapping for a given input and reference image pair can be encoded in a look-up table and then applied to all subsequent frames in the same way. In the method by Hogervorst et al. [HT08, HT10], the look-up table serves as a predefined mapping between input and reference pixel values for the purpose of colorizing night-vision imagery so that it acquires a daytime appearance. Similarly, by learning a global mapping expressed in terms of luminance, hue and saturation changes from a collection of clips, particular styles can be applied to a given video [XADR13].

A global mapping used for all frames of a sequence is a simple solution of achieving temporal coherence and it also offers the additional advantage of low computational complexity for applying the mapping. However, it is not suitable when the aim is to change the mapping over the duration of the video sequence. More specifically, although a global mapping could be used to when the reference is a still image, different solutions are necessary when multiple references are used [WHH06] or when mapping colors between two videos [BSPP13]. In those cases, temporal smoothing is necessary.

Wang and Huang extend the statistical color transfer method by Reinhard et al. [RAGS01] to the temporal dimension and allow for multiple reference images to be used for different parts of the sequence to vary the transferred appearance over time [WHH06]. The different reference images are assigned to specific key frames in the sequence, for which the mapping can be determined according to the mean and standard deviation of the input and reference images. To ensure that the mapping is temporally coherent and that no artifacts appear for intermediate frames, the color statistics of the reference images can be interpolated using a linear or non-linear curve, obtaining a mapping for each frame of the sequence.

If the reference is a video sequence as well, it would be possible to determine a more dense mapping between frames. However, small temporal changes in the reference may not correspond to the content of the input, leading to artifacts. To resolve this issue, Bonneel et al. [BSPP13] opt for selecting a number of representative frames that best describe the style of the reference video, which can be used to

determine a per-frame mapping for each input frame. If this mapping is applied naively, flickering can occur in parts of the sequence where the mapping changes abruptly. Instead, similar to Wang and Huang [WHH06], the per frame mapping can be smoothed by interpolating between key frames, as they represent stable parts of the sequence.

This form of temporal smoothing can successfully reduce flickering artefacts, while allowing the color mapping to change over the duration of the sequence. At the same time, as the whole mapping needs to be pre-determined and smoothed before applying, such a solution is not suitable for online processing of video, as would be needed on the side of the display.

An alternative solution is that of temporal smoothing using a sliding window in the temporal dimension. In Section 2.2.4, methods relying on image segmentation to determine correspondences were discussed. When extending such methods to video, pixels belonging to the same object may be clustered into different regions in subsequent frames, which in turn may lead to flicker. In the color categorization based solution by Chang et al. [CSN07b], the anisotropic filtering step used to smoothly segment the image is extended to the temporal dimension, by smoothing the color category based segmentation across a few seconds.

Rather than map colors to those of a different video or image, a reference frame can be selected from the input video itself, serving as an ‘anchor’ for the remaining frames when dealing with tonal fluctuations in videos due to automatic adjustment of camera settings [FL11]. In this case, subsequent frames have much stronger correspondences since they are likely to represent the same scene, allowing for color correspondences to be determined using an approximate optical-flow computed from the luminance channel, as discussed in Section 2.1.

3.4. Stereoscopic Video Color Transfer

Stereoscopy (or stereo imaging) uses pairs of images/videos captured with slightly different viewpoints, to reproduce a sense of depth.

In human vision, various anomalies between imagery presented to the left and right eyes, including color differences, can create binocular rivalry [Bla01] and affect visual comfort [KT04]. In image processing, when disparity or depth is estimated from matching left and right views, color differences are used as a cost function and therefore should be zero for corresponding semantic elements [SS02]. Likewise, the calculation of feature correspondences could become more robust if images are matched in color first [FST11].

Unfortunately the colors in stereo pairs are often not well matched. In stereo video cameras, left and right frames are captured by the same sensor. The polarization filters and mirrors — used to separate left and right views — introduce

color differences. Additional color differences are caused by differences in optics (vignetting, color filters), sensor (sensitivity, uniformity) and processing (white balance, color encoding, image enhancement).

Color characterization as well as color calibration of the cameras involved in the capture can be used to prevent color differences. Color characterization involves models describing how light is transformed into the cameras’ RGB signal. Color calibration aims to directly correct the RGB signal with respect to a reference [TW05, KP07, KLL*12]. This procedure requires the placement of a color checker board in the scene. Its colors are acquired by the cameras and these are used to compensate for color differences between them. Existing methods include offset/gain optimization [JWV*05], linear cross-channel color correction either with respect to average colors [JWV*05] or they can be stated as a combined optimization problem [LDX10].

Although there are benefits to maintaining well calibrated cameras, in practice, calibration will have to be repeated with some regularity, for instance because camera settings changes, lenses are swapped out, or simply because the cameras are aging. As such, rather than prevent color mismatches, color transfer algorithms could be employed to correct for such mismatches without worrying about precise camera calibration.

The nature of stereoscopic photography is that there is significant overlap between the left and right frames. This makes geometry-based color transfer particularly amenable for this application [DN09] (see Section 2.1).

In image formation of stereoscopic images, color differences and disparity/depth are inherently linked. If spatially corresponding points in the left and right images have a significant color difference, then the likelihood that there is a depth discontinuity at that position is high. As a consequence, color differences and depth can be estimated in a jointly using an iterative scheme [HLL13]. To enhance depth estimation, color differences are then compensated iteratively using a linear model in a linearized chromaticity space.

The design of color transfer algorithms for stereo pair correction can be aided by the presence of ground-truth imagery. A common method to evaluate the performance of color mapping is to assess differences between a monoscopic presentation of left view, right view and mapped right view [WSW10]. Another way to evaluate color mapping in the framework of stereoscopic imaging is to measure the impact of color mapping on the quality disparity estimates. It was found that color transfer can indeed facilitate the creation of high quality disparity maps [WSW10, HLL13].

Commercially available tools for stereoscopic color mapping — mainly in post-production pipelines — include RE:Match, Matchgrade in NUKEX, MatchLight IMS.

3.5. Multi-View Video Coding

Multi-view video is concerned with capturing multiple video streams of the same scene. It can be seen as an extension of stereoscopic video, and therefore benefits from color transfer in similar ways.

Multi-view video allows in post-processing new virtual viewpoints to be derived [CSN07a]. Applications of multi-view video include free-viewpoint television [Tan04, TZMS04] and 3D TV [VMPX04]. The set of video captures is called a multi-view video sequence or a dynamic light field [FBK08]. Of course, if a large number of cameras is used, then the amount of data produced is enormous, and therefore efficient compression schemes are required.

Simulcast coding, whereby each video stream is encoded separately, would be sub-optimal as it does not leverage any redundancy that may exist between frames taken by different cameras. A more efficient approach would be to exploit the spatial correlations between frames that come from different cameras (by means of disparity estimation, e.g. [SJYH10]). However, as cameras are often calibrated differently, have (slightly) different lenses, and are operated in different illumination conditions, the efficiency of multi-view video coding can be improved by minimizing the differences between video streams captured by these cameras.

It is possible to apply color transfer prior to the encoding stage. This will minimize the color differences between images, and thus allow more efficient data compression. Techniques for this include modeling spatio-temporal variation [SJYH10], image segmentation [SJYC07], histogram-based techniques [FBK08, CMC10], view interpolation prediction [YYF*07], feature matching [YYF*07, YO08b] and global matching to the average of all corresponding frames [DN09].

3.6. Image Stitching

Creating panoramas from a set of individual photographs is known as image stitching. Typically, better panoramas can be obtained if the colors of the input images are made more consistent first [HS04, JT03, TGTC02, ADA*04, BL07]. Related to this application area is photo-tourism, whereby geotagged images are smoothly interpolated to visualize a landmark that was photographed by potentially many different cameras [SSS06]. This application also benefits from color consistency [SGSS08].

Figure 10 shows an example of image stitching performed with and without color correction. In the simplest case, in order to have a good overlap between input images, the viewpoint change between input images is not very important and, in order to avoid strong global or local color differences, the illumination conditions are quite similar between input images. In the worst case, many factors contribute to create global and local color differences. Among these factors we can mention vignetting, automatic camera settings (exposure, white balance, etc.), camera device (optics, sensor,




Figure 10: Color mapping (in this case the method by Faridul et al. [FST13]) can be used to remove color differences between overlapping images to improve the result of stitching operations.

nonlinear processing [KLL*12]), viewing conditions (viewpoint, illumination, scene content), etc. For example, input images may come from different camera devices and captured by different people at different times. In that case, most of stitching methods cannot compensate these complex color differences.

There are two ways to perform image stitching. One way is to color correct input images, next to stitch resulting images. For example, the method proposed by Hachen et al. [HSGL11] based on local set of dense correspondences first performs colors correction next stitching. This method is designed for input images depicting similar regions acquired by different cameras and lenses, under non-rigid transformations, under different lighting, and over different backgrounds. The two main drawbacks of this method is that it requires dense correspondences and that it is not robust to drastic color changes due to for example to illuminant changes. In a recent survey, Xu et al. pointed out the performance of color correction approaches for automatic multi-view image and video stitching [XM10].

Alternatively, it is possible to first stitch images and perform the color correction step as a post-process. Different color corrections methods can be used in this case, including color mapping, color transfer or color balance. For example, Wu et al. in [WDK*13] proposed a progressive color transfer method to preserve the visual coherence and the semantic correspondence of regions of input images. On the other hand, Qian et al. [QLZ13] proposed a manifold alignment method to perform color transfer by exploring manifold structures of partially overlapped input images.

Most of image stitching methods are based on a global luminance compensation followed by image blending. Other methods use gain compensation along with local image blending [BL07, Clo, NC01, UES01]. Others are based on brightness transfer [KP07], mapping based color constancy [TGTC02], or tensor voting to correct luminance [JT03]. The relevance of global gain compensation

methods depends of the overlap between input images. It can be performed either applied from the luminance channel [NC01, UES01] or from the three color channels [JT03, HS04, KP07, TGTC02].

In order to estimate and correct color differences between input images, correspondences can be detected using either a dense method or a sparse method, as discussed in Section 2.1. Sparse methods are based on the computation of geometric features and on features matching between input images. The relevance of sparse methods depends of the overlap between input images (i.e. the number of feature points common to input images) and of the relevance of the color sampling used to model color differences (i.e. the number of colors in input images for which there is no features matching). One of the main challenges of sparse methods is generating confident color correspondences from unreliable feature correspondences.

Another way to correct colors is to model color distributions of input images by statistics, followed by compensation of color differences between these distributions by statistic transfer methods (see Section 2.2). Although these methods can produce plausible results even when the input images are very different, a robust mapping suitable for stitching cannot be guaranteed.

To compensate color differences due to illuminant changes before stitching, Tian et al. proposed in [TGTC02] to correct colors by using computational color constancy. Nevertheless, this approach has three drawbacks. First, the scene illumination estimation requires several assumptions about the scene that may not hold in case of image stitching. Second, most of color constancy methods rely on simplifying assumptions, such that the input images are acquired under the same canonical illuminant. Third, color constancy methods cannot compensate color differences due to camera devices (sensors, optics, etc.), camera settings (white balance, exposures), or other nonlinear processing.

3.7. Image Fusion

Color transfer has found application in the general area of image fusion. Image fusion is concerned with capturing and representing images that were taken with camera systems that operate on different wavelength ranges, with applications in surveillance, face recognition, target tracking, concealed weapon detection, intelligence gathering and the detection of abandoned packages and buried explosives [TH12]. Civil applications include remote sensing, agriculture, medicine and art analysis [BL06]. For instance, night-time scenes can be captured with cameras that work in the normal range of visible wavelengths (400–700 nm), in the near-infrared range (700–900 nm) as well as the thermal middle wavelength band (3–5 μm) infrared images [Toe03].

If a same night-time scene is captured with multiple types of camera, then the results can be fused into a single false

color image that represents the disjunct information available in each wavelength regime. There are many ways in which this information can be fused. Traditionally, these different captures would be shown as grey tones, which is difficult to interpret.

This has given rise to color image fusion, which allows a significantly expanded range of different shades of color that can be distinguished by human vision. In addition, relative to monochrome images, fused color images can yield better and more accurate reaction times, scene understanding, scene recognition and object recognition [SH09, TH12].

It can be argued, however, that if the fusion process results in imagery that is not only in color, but additional appears in some sense natural, then it is possible to further increase situational awareness [TIWA97, Var98, GSFZ12] — an aspect of image perception that is important in security and military applications.

Visually plausible imagery can be created from a multi-modal set of image captures with the aid of color transfer algorithms [Toe03, TA05, HNZ07, LW07]. As noted, the purpose is to create a false color image that appears natural. The image will never be a true representation of the scene's content, as wavelengths outside the visible spectrum needs to be merged into an image that can be observed by humans. Nonetheless, color transfer can help create imagery with natural day-time qualities, which are therefore amenable to human interpretation.

A simple approach to achieve color transfer is to assign the output of the aforementioned three cameras to the red, green and blue channels of an RGB color space. This produces an unnatural looking false color image. However, this can be adjusted by selecting an appropriate natural day-time image and by applying a color transfer technique to the false color image. The result will then be a significantly more natural looking visualization of the input multi-modal night-time image. Matching means and standard deviations in $L\alpha\beta$ space has proven to be an appropriate approach [Toe03, SJLL05, YXM*12], a procedure previously applied to conventional photographs [RAGS01]. With lower computational complexity, this approach has also been proposed in the YUV color space [SWqLZ07].

This technique can be extended to a spatially varying algorithm with the aid of segmentation [ZHHE05, ZDB12], which is significantly more computationally expensive, but it has been argued that the visual quality of the results is superior [ZDB12].

Alternatively, this problem can be solved in real-time [HT08], for instance with the aid of look-up tables [HT10], and can even be applied in the context of augmented reality [THvSD12]. For further information on color transfer in the context of image fusion, see Toet and Hogervorst's recent paper [TH12].

3.8. Augmented Reality

Augmented reality is concerned with superimposing computer generated content onto live captured video and with presenting the result in real-time to the observer. In this application, the computer generated content will have been rendered with certain illumination and into a specific color space. At the same time, the camera will be capturing live footage with its own illumination, using camera-specific settings. To make the rendered imagery sit well with the captured footage, it will have to be adjusted.

Of course, it would be possible to simulate various aspects of the camera and apply the simulation to the rendered imagery. This may include features such as radial distortion, Gaussian blur, motion blur, Bayer sampling, as well as firmware processing including sharpening, quantization and various color transforms [KM10].

If only the color content is to be adjusted, then color transfer techniques come into play. Here, the differences in color spaces, illumination and possibly even a poor choice of materials in the rendered content, can be adjusted for. The simplest technique arguably involves matching means and standard deviations in a suitably chosen decorrelated color space [RAGS01], which was shown to produce plausible results in the context of augmented reality [RAC*04]. Polynomial regression, followed by leaky integration to ensure temporal stability is also applicable in this context [KTPW11].

Finally, for certain augmented reality applications, it may be appropriate to use tracking of objects so that a sparse set of color correspondences can be kept stable over time [OHSG12]. In the CIELab color space, a dense field of correspondences is then created by interpolating radial basis functions. The result can be used to map colors between images, or in this case, between rendered content and captured footage.

4. Color Mapping Evaluation

Most state of the art works on color mapping report only visual results. Interpretation of visual results is often not only subjective, but also difficult to compare. Therefore, quantitative evaluation of color mapping results remains a open and important research problem.

Ideally, the results of a color mapping method can be assessed through comparisons with a ground truth. In cases where the pair of input images are similar in structure, it is possible to produce image sets that provide a ground truth [FST12, FST13] (see Figure 11). Further, quality metrics can be employed to compare the result with the original input image [SDZL09, TD11].

In most color mapping scenarios, however, it is often impossible to construct or capture a ground truth. In such cases, it is possible to assess the mapped results using no-reference quality metrics [XM09, OZC12] that consider aspects such as the color distribution, hue or gradients of

the images. Alternatively, user studies may be employed to provide a quantitative assessment of the quality of the results [VdNJ*07, XADR13].

4.1. Ground Truth Evaluation

In the case of stereo or multiple views, Faridul et al. [FST12] provide a quantitative evaluation framework. Here, the idea is to capture *ground truth* color along with the test images. For example, let us consider a scene which can be acquired under two different illumination conditions and from two different viewpoints, e.g. simulating a stereo capture scenario. To create a ground truth image for this scenario, it is sufficient to capture an image with the same viewpoint as the input image (i.e. the one to be recolored) but with the same illumination as the reference. Now this ground truth image represents what the input image should look like after successful color mapping.

Instead of changing the illumination, the same principle can be extended to model other causes of variation in color between images. Faridul et al. [FST13] constructed a dataset suitable for evaluating color mapping methods by capturing ground truth sets for varying shutter speed and different white balance settings. Although this type of controlled input can be used to evaluate color mapping methods, ground truth images can only be acquired when the input and reference images represent the same scene under different viewpoints. It is therefore better suited for assessing methods where overlap in the scene is desirable, such as discussed in Section 2.1. Figure 11 shows some examples from the dataset described in [FST13].

4.2. Evaluation Metrics

Perceptual metrics defined in CIELab color space such as CIE $\Delta E00$ [LCR01] can be used to evaluate the quality of color mapping results when a *ground truth* is available. Additionally, a recent work [ZDB12] comparing the colorizing technique of multispectral images summarizes other CIELab based metrics. Such metrics include: objective evaluation index [ZDB12], phase congruency [Kov99], gray relational analysis [MTH05], gradient magnitude metric, image contrast metric.

On the other hand, PSNR is a popular performance evaluation metric notably in the literature of multiple view coding. Here, the idea is to report PSNR with and without color mapping. PSNR are often reported: separately [FBK08, DN09] for Y , C_b , C_r ; average PSNR of all color channels [YO08a]; Or, combination [XM10] of PSNR and SSIM [WBSS04].

4.3. Quality Assessment for Color Problems

Some solutions exist for the assessment of colored (impaired) pictures with respect to a reference (original) picture [SDZL09, TD11]. The reference picture has to match at

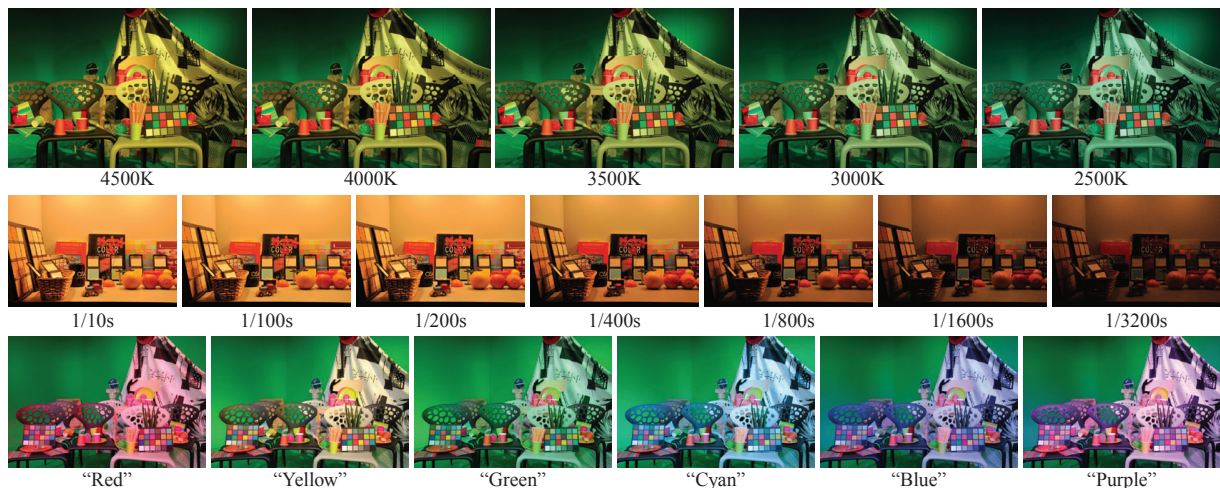


Figure 11: Color variability can be introduced in a controlled manner by systematically varying capture properties such as exposure time (top) and white point (middle), or scene aspects such as illumination (bottom), creating a dataset for assessing color mapping methods [FST13].

a pixel level the impaired picture limiting the application of such quality metric. On the other hand, no-reference quality metrics have a wider field of application, since they only compute statistics on the impaired picture. Their efficiency and accuracy is limited however in the context of assessing color mapping results since they rely on low-level features. Ouni et al. [OZC12] have proposed different color statistics analysis, relying on properties such as the distribution of hue histogram or the proportion and dispersion of the dominant color, to derive a quality score.

Although quality assessment is an active research area [AMMS08, HCA*12, MDMS05], the particular problem of no-reference color quality assessment has been largely unexplored. One solution proposed in the context of color mapping in particular measures mean squared error in the color and gradient domain, comparing the resulting image against the reference in terms of their color histogram, and against the input in terms of gradients [XM09]. Although this method offers a conceptually attractive solution, it cannot determine whether the color mapping is plausible or pleasing—it simply assesses whether it is close to the reference.

4.4. Evaluation through User Studies

One possibility for assessing the quality of a color mapping solution is simply to ask human observers. This is typically formalized as a user study or psychophysical experiment [CW11]. For instance, Viera et al. [VdNJ*07] proposed colorization of images as well as a user study to demonstrate the plausibility of their colorization algorithm. They conducted the user studies by asking participants whether “this

image is colored in a 1) totally plausible, 2) mostly plausible, 3) mostly implausible or 4) totally implausible way”. In a similar study evaluating an automatic method for applying color styles to image and video content, participants were asked to select the image better reflecting a particular style, such as “happy” or “sad”.

5. Conclusions

Some images have a more interesting look and feel to it than other images. Photographs may have more natural color distributions than some rendered content. This was the inspiration that led to the development of early color transfer algorithms. The idea was that an example-based post-process could make a mundane image appear more interesting by taking some color characteristics from an existing image.

This basic idea has since led to a proliferation of ever more sophisticated techniques as well as an explosion of applications that were found to benefit from example-based processing. Some of the more surprising application areas include interpretation of night vision imagery, correction of stereo and multi-view video, image stitching, artistic tools.

Whenever color transfer is applied between images that show overlap, including in stereo and image stitching applications, color transfer based on finding feature correspondences are applicable. When images come from entirely different domains, then statistical techniques may be more appropriate. In all cases, one of the main open problems is that there is often no good ground truth available. Especially in statistical transfers where the two input images are unrelated, it is difficult to assess whether a result is good or not.

Nonetheless, many color transfer algorithms often pro-

duce images that are intuitively plausible. As this is often good enough, they find increasingly often employ, including for movies and visual effects.

References

- [ADA*04] AGARWALA A., DONTCHEVA M., AGRAWALA M., DRUCKER S., COLBURN A., CURLLESS B., SALESIN D., COHEN M.: Interactive digital photomontage. *ACM Transactions on Graphics (TOG)* 23, 3 (2004), 294–302. 16
- [AK04] ABADPOUR A., KASAEI S.: A fast and efficient fuzzy color transfer method. In *Proceedings of the 4th IEEE International Symposium on Signal Processing and Information Technology* (2004), pp. 491–494. 8
- [AK07] ABADPOUR A., KASAEI S.: An efficient PCA-based color transfer method. *Journal of Visual Communication and Image Representation* 18 (2007), 15–34. 5, 6, 8
- [AMMS08] AYDIN T., MANTIUK R., MYSZKOWSKI K., SEIDEL H.: Dynamic range independent image quality assessment. *ACM Transactions on Graphics* 27 (2008), 69(1)–69(10). 19
- [AP08] AN X., PELLACINI F.: AppProp: all-pairs appearance-space edit propagation. *ACM Transactions on Graphics (TOG)* 27, 3 (2008), 40. 10
- [API0] AN X., PELLACINI F.: User controllable color transfer. *Computer Graphics Forum* 29 (2010), 263–271. 2, 10
- [Ber69] BERLIN B.: *Basic color terms: Their universality and evolution*. University of California Pr, 1969. 9
- [BL97] BEIS J. S., LOWE D. G.: Shape indexing using approximate nearest-neighbour search in high-dimensional spaces. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition* (1997), pp. 1000–1006. 3
- [BL06] BLUM R. S., LIU Z.: *Multi-sensor image fusion and its applications*. CRC Press, Boca Raton, FL, 2006. 17
- [BL07] BROWN M., LOWE D. G.: Automatic panoramic image stitching using invariant features. *International Journal of Computer Vision* 74, 1 (2007), 59–73. 2, 16
- [Bla01] BLAKE R.: A primer on binocular rivalry, including current controversies. *Brain and Mind* 2, 1 (2001), 5–38. 15
- [BNTS07] BOUSSEAU A., NEYRET F., THOLLOT J., SALESIN D.: Video watercolorization using bidirectional texture advection. *ACM Trans. Graph.* 26, 3 (July 2007). 14
- [BSFG09] BARNES C., SHECHTMAN E., FINKELSTEIN A., GOLDMAN D.: PatchMatch: a randomized correspondence algorithm for structural image editing. *ACM Transactions on Graphics* 28, 3 (2009), 24. 4
- [BSGF10] BARNES C., SHECHTMAN E., GOLDMAN D. B., FINKELSTEIN A.: The generalized PatchMatch correspondence algorithm. In *Proceedings of the European Conference on Computer Vision* (2010), pp. 29–43. 4
- [BSPP13] BONNEEL N., SUNKAVALLI K., PARIS S., PFISTER H.: Example-based video color grading. *ACM Transactions on Graphics (Proceedings of SIGGRAPH 2013)* 32, 4 (2013). 2, 14
- [BT12] BUGEAU A., TA V.-T.: Patch-based image colorization. In *Pattern Recognition (ICPR), 2012 21st International Conference on* (2012), IEEE, pp. 3058–3061. 11
- [BTP13] BUGEAU A., TA V.-T., PAPADAKIS N.: Variational exemplar-based image colorization. *IEEE Transactions on Image Processing* (2013). 11
- [BTVG06] BAY H., TUYTELAARS T., VAN GOOL L.: SURF: Speeded up robust features. In *Proceedings of the European Conference on Computer Vision* (2006), pp. 404–417. 3
- [BUC*13] BAVEYE Y., URBAN F., CHAMARET C., DEMOULIN V., HELLIER P.: Saliency-guided consistent color harmonization. In *Computational Color Imaging*, Tominaga S., Schettini R., Trémeau A., (Eds.), vol. 7786 of *Lecture Notes in Computer Science*. Springer Berlin Heidelberg, 2013, pp. 105–118. 13
- [BZS*07] BHAT P., ZITNICK C. L., SNAVELY N., AGARWALA A., AGRAWALA M., COHEN M., CURLLESS B., KANG S. B.: Using photographs to enhance videos of a static scene. In *Proceedings of the Eurographics Symposium on Rendering* (2007), pp. 327–338. 14
- [CHS08] CHARPIAT G., HOFMANN M., SCHÖLKOPF B.: Automatic image colorization via multimodal predictions. In *Proceedings of the 10th European Conference on Computer Vision: Part III* (Berlin, Heidelberg, 2008), ECCV '08, Springer-Verlag, pp. 126–139. 11
- [Clo] CLOUDBURST RESEARCH: AutoStitch: a new dimension in automatic image stitching. <http://cs.bath.ac.uk/brown/autostitch/autostitch.html>. 16
- [CM02] COMANICIU D., MEER P.: Mean shift: a robust approach toward feature space analysis. *IEEE Trans. on Pattern Analysis and Machine Intelligence* 24, 5 (2002), 603–619. 4, 9
- [CMC10] CHEN Y., MA K., CAI C.: Histogram-offset-based color correction for multi-view video coding. In *Proc. of ICIP* (2010), pp. 977–980. 2, 4, 16
- [COSG*06] COHEN-OR D., SORKINE O., GAL R., LEYVAND T., XU Y.: Color harmonization. *ACM Transactions on Graphics* 25, 3 (2006), 624–630. 10, 13
- [CSN07a] CHAN S. C., SHUM H. Y., NG K. T.: Image-based rendering and synthesis. *IEEE Signal Processing Magazine* 24, 6 (2007), 22–33. 16
- [CSN07b] CHANG Y., SAITO S., NAKAJIMA M.: Example-based color transformation of image and video using basic color categories. *IEEE Transactions on Image Processing* 16, 2 (feb. 2007), 329–336. 9, 15
- [CSUN06] CHANG Y., SAITO S., UCHIKAWA K., NAKAJIMA M.: Example-based color stylization of images. *ACM Transactions on Applied Perception* 2, 3 (2006), 322–345. 9
- [CUS04] CHANG Y., UCHIKAWA K., SAITO S.: Example-based color stylization based on categorical perception. In *Proceedings of the 1st Symposium on Applied perception in graphics and visualization (APGV 2004)* (New York, NY, USA, 2004), ACM, pp. 91–98. 8, 9
- [CW11] CUNNINGHAM D., WALLRAVEN C.: *Experimental Design: From User Studies to Psychophysics*. AK Peters, Natick, MA, 2011. 19
- [CZG*11] CHIA A. Y.-S., ZHUO S., GUPTA R. K., TAI Y.-W., CHO S.-Y., TAN P., LIN S.: Semantic colorization with internet images. *ACM Transactions on Graphics* 30, 6 (2011), 156. 12
- [DLR77] DEMPSTER A. P., LAIRD N. M., RUBIN D. B.: Maximum likelihood from incomplete data via the EM algorithm. *Journal of the Royal Statistical Society B* 39, 1 (1977), 1–38. 9
- [DN09] DOUTRE C., NASIOPOULOS P.: Color correction pre-processing for multiview video coding. *IEEE Transactions on Circuits and Systems for Video Technology* 19, 9 (2009), 1400–1406. 2, 5, 15, 16, 18
- [DS02] DECARLO D., SANTELLA A.: Stylization and abstraction of photographs. *ACM Transactions on Graphics* 21, 3 (July 2002), 769–776. 14
- [FB81] FISCHLER M. A., BOLLES R. C.: Random sample consensus: a paradigm for model fitting with applications to image analysis and automated cartography. *Communications of the ACM* 24, 6 (1981), 381–395. 3

- [FBK08] FECKER U., BARKOWSKY M., KAUP A.: Histogram-based prefiltering for luminance and chrominance compensation of multiview video. *IEEE Transactions on Circuits and Systems for Video Technology* 18, 9 (2008), 1258–1267. 16, 18
- [FLL10] FARBMAN Z., FATTAL R., LISCHINSKI D.: Diffusion maps for edge-aware image editing. *ACM Transactions on Graphics (TOG)* 29, 6 (2010), 145. 11
- [FL11] FARBMAN Z., LISCHINSKI D.: Tonal stabilization of video. *ACM Transactions on Graphics (Proceedings of ACM SIGGRAPH 2011)* 30, 4 (2011), 89:1–89:9. 15
- [FST11] FARIDUL H. S., STAUDER J., TRÉMEAU A.: Robust color correction for stereo. In *Proc. of Conference for Visual Media Production* (2011), pp. 101–108. 15
- [FST12] FARIDUL H. S., STAUDER J., TRÉMEAU A.: Optimization of sparse color correspondences for color mapping. In *Proc. of Color and Imaging Conference* (2012). 2, 3, 18
- [FST13] FARIDUL H. S., STAUDER J., TRÉMEAU A.: Approximate Cross Channel Color Mapping from Sparse Color Correspondences. In *Proceedings of International Conference on Computer Vision (ICCV)* (2013). 16, 18, 19
- [FWT*11] FENG J., WEI Y., TAO L., ZHANG C., SUN J.: Salient object detection by composition. In *Computer Vision (ICCV), 2011 IEEE International Conference on* (2011), IEEE, pp. 1028–1035. 10
- [GD05] GRUNDLAND M., DODGSON N.: Color histogram specification by histogram warping. *Color Imaging X: Processing, Hardcopy, and Applications* 5667, 1 (2005), 610–621. 5, 6
- [GH03] GREENFIELD G., HOUSE D.: Image recoloring induced by palette color associations. *Journal of the WSCG* 11, 1 (2003), 189–196. 8
- [GSFZ12] GU X., SUN S., FANG J., ZHOU P.: Kernel based color estimation for night vision imagery. *Optics Communications* 285, 7 (2012), 1697–1703. 17
- [Hae90] HAEBERLI P.: Paint by numbers: Abstract image representations. *ACM SIGGRAPH Computer Graphics* 24, 4 (1990), 207–214. 14
- [HČA*12] HERZOG R., ČADÍK M., AYDIN T. O., KIM K. I., MYSZKOWSKI K., SEIDEL H.-P.: Norm: No-reference image quality metric for realistic image synthesis. *Computer Graphics Forum* 31, 2 (2012), 545–554. 19
- [HE04] HAYS J., ESSA I. A.: Image and video based painterly animation. In *NPAR* (2004), pp. 113–120. 14
- [HEH05] HOIEM D., EFROS A. A., HEBERT M.: Geometric context from a single image. In *Computer Vision, 2005. ICCV 2005. Tenth IEEE International Conference on* (2005), vol. 1, IEEE, pp. 654–661. 10
- [HEH07] HOIEM D., EFROS A. A., HEBERT M.: Recovering surface layout from an image. *International Journal of Computer Vision* 75, 1 (2007), 151–172. 10
- [Hei72] HEIDER E. R.: Universals in color naming and memory. *Journal of experimental psychology* 93, 1 (1972), 10. 9
- [Her98] HERTZMANN A.: Painterly rendering with curved brush strokes of multiple sizes. In *Proceedings of the 25th Annual Conference on Computer Graphics and Interactive Techniques* (New York, NY, USA, 1998), SIGGRAPH '98, ACM, pp. 453–460. 14
- [Her10] HERTZMANN A.: Non-photorealistic rendering and the science of art. In *Proceedings of the 8th International Symposium on Non-Photorealistic Animation and Rendering* (New York, NY, USA, 2010), NPAR '10, ACM, pp. 147–157. 13, 14
- [HHH09] HYVÄRINEN A., HURRI J., HOYER P. O.: *Natural image statistics: A probabilistic approach to early computational vision*. Springer, Dordrecht, 2009. 7
- [HJO*01] HERTZMANN A., JACOBS C. E., OLIVER N., CURLESS B., SALESIN D.: Image analogies. In *Proceedings of ACM SIGGRAPH* (2001), pp. 327–340. 2
- [HLL13] HEO Y. S., LEE K. M., LEE S. U.: Joint depth map and color consistency estimation for stereo images with different illuminations and cameras. *Pattern Analysis and Machine Intelligence, IEEE Transactions on* 35, 5 (2013), 1094–1106. 15
- [HNZ07] HUANG G., NI G., ZHANG B.: Visual and infrared dual-band false color image fusion method motivated by land's experiment. *Optical Engineering* 46, 2 (2007), 027001. 17
- [HP00] HERTZMANN A., PERLIN K.: Painterly rendering for video and interaction. In *Proceedings of the 1st International Symposium on Non-photorealistic Animation and Rendering* (New York, NY, USA, 2000), NPAR '00, ACM, pp. 7–12. 14
- [HS04] HASLER D., SÜSSTRUNK S.: Mapping colour in image stitching applications. *Journal of Visual Communication and Image Representation* 15, 1 (2004), 65–90. 16, 17
- [HSGL11] HACOHEM Y., SHECHTMAN E., GOLDMAN D. B., LISCHINSKI D.: Non-rigid dense correspondence with applications for image enhancement. *ACM Transactions on Graphics* 30, 4 (2011), 70:1–70:9. 2, 4, 5, 16
- [HSGL13] HACOHEM Y., SHECHTMAN E., GOLDMAN D. B., LISCHINSKI D.: Optimizing color consistency in photo collections. *ACM Transactions on Graphics (TOG)* 32, 4 (2013), 38. 4
- [HT08] HOGERVORST M. A., TOET A.: Method for applying daytime colors to nighttime imagery in realtime. In *Proceedings of SPIE 6974, Multisensor, Multisource Information Fusion: Architectures, Algorithms, and Applications* (2008), Dasarathy B. V., (Ed.), vol. 697403. 14, 17
- [HT09] HUO X., TAN J.: An improved method for color harmonization. In *Image and Signal Processing, 2009. CISP '09. 2nd International Congress on* (Oct 2009), pp. 1–4. 13
- [HT10] HOGERVORST M. A., TOET A.: Fast natural color mapping for night-time imagery. *Image Fusion* 11, 2 (2010), 69–77. 14, 17
- [HTC*05] HUANG Y.-C., TUNG Y.-S., CHEN J.-C., WANG S.-W., WU J.-L.: An adaptive edge detection based colorization algorithm and its applications. In *Proceedings of the 13th annual ACM international conference on Multimedia* (2005), ACM, pp. 351–354. 12
- [HTF01] HASTIE T., TIBSHIRANI R., FRIEDMAN J.: *The elements of statistical learning*. Springer, 2001. 4
- [ICOL05] IRONY R., COHEN-OR D., LISCHINSKI D.: Colorization by example. In *Rendering Techniques* (2005), Deussen O., Keller A., Bala K., Dutré P., Fellner D. W., Spencer S. N., (Eds.), Eurographics Association, pp. 201–210. 2, 11
- [Itt61] ITTEN J.: *Kunst der Farbe*. Reinhold Pub. Corp., 1961. 13
- [IW05] ILIE A., WELCH G.: Ensuring color consistency across multiple cameras. In *IEEE International Conference on Computer Vision* (2005), vol. 2, pp. 1268–1275. 15
- [JT03] JIA J., TANG C.-K.: Image registration with global and local luminance alignment. *IEEE International Conference on Computer Vision* 1 (2003), 156. 16, 17
- [JWV*05] JOSHI N., WILBURN B., VAISH V., LEVOY M. L., HOROWITZ M.: *Automatic color calibration for large camera arrays*. [Department of Computer Science and Engineering], University of California, San Diego, 2005. 15

- [KKD09] KYPRIANIDIS J. E., KANG H., DÖLLNER J.: Image and video abstraction by anisotropic kuwahara filtering. *Computer Graphics Forum* 28, 7 (2009), 1955–1963. 14
- [KLC09] KANG H., LEE S., CHUI C.: Flow-based image abstraction. *Visualization and Computer Graphics, IEEE Transactions on* 15, 1 (Jan 2009), 62–76. 14
- [KLL*12] KIM S. J., LIN H. T., LU Z., SABINE S., ET AL.: A new in-camera imaging model for color computer vision and its application. *IEEE Transactions on Pattern Analysis and Machine Intelligence* (2012), 2289–2302. 15, 16
- [KM07] KUMAR R., MITRA S. K.: Motion estimation based color transfer and its application to color video compression. *Pattern Analysis and Applications* 11, 2 (2007), 131–139. 8
- [KM10] KLEIN G., MURRAY D. W.: Simulating low-cost cameras for augmented reality compositing. *IEEE Transaction on Visualization and Computer Graphics* 16, 3 (2010), 369–380. 18
- [KMHO09] KAGARLITSKY S., MOSES Y., HEL-OR Y.: Piecewise-consistent color mappings of images acquired under various conditions. In *Proc. of ICCV* (2009), pp. 2311–2318. 4
- [KMS98] KOTERA H., MORIMOTO T., SAITO R.: Object-oriented color matching by image clustering. In *Proceedings of the 6th IS&T Color Imaging Conference* (1998), pp. 154–158. 8
- [Kot05] KOTERA H.: A scene-referred color transfer for pleasant imaging on display. In *Proc. of ICIIP* (2005), vol. 2, IEEE, pp. II–5. 8
- [Kov99] KOVESI P.: Image features from phase congruency. *Journal of Computer Vision Research* 1, 3 (1999), 1–26. 18
- [KP07] KIM S., POLLEFEYS M.: Robust radiometric calibration and vignetting correction. *IEEE Trans. on Pattern Analysis and Machine Intelligence* (2007), 562–576. 15, 16, 17
- [KT04] KOOI F. L., TOET A.: Visual comfort of binocular and 3D display. *Displays* 25, 2 (2004), 99–108. 15
- [KTPW11] KNECHT M., TRAXLER C., PURGATHOFER W., WIMMER M.: Adaptive camera-based color mapping for mixed-reality applications. In *Proceedings of the 10th IEEE International Symposium on Mixed and Augmented Reality* (2011), pp. 165–168. 18
- [LCR01] LUO M., CUI G., RIGG B.: The development of the CIE 2000 colour-difference formula: CIEDE2000. *Color Research & Application* 26, 5 (2001), 340–350. 18
- [LDX10] LI K., DAI Q., XU W.: High quality color calibration for multi-camera systems with an omnidirectional color checker. In *Acoustics Speech and Signal Processing (ICASSP), 2010 IEEE International Conference on* (2010), IEEE, pp. 1026–1029. 15
- [LFUS06] LISCHINSKI D., FARBMAN Z., UYTENDAELE M., SZELISKI R.: Interactive local adjustment of tonal values. *ACM Transactions on Graphics* 25, 3 (2006), 646–653. 10
- [Lin98] LINDBERG T.: Feature detection with automatic scale selection. *International Journal of Computer Vision* 30, 2 (1998), 79–116. 3
- [Lit97] LITWINOWICZ P.: Processing images and video for an impressionist effect. In *Proceedings of the 24th Annual Conference on Computer Graphics and Interactive Techniques* (New York, NY, USA, 1997), SIGGRAPH '97, ACM Press/Addison-Wesley Publishing Co., pp. 407–414. 14
- [LJH10] LI Y., JU T., HU S.-M.: Instant propagation of sparse edits on images and videos. *Computer Graphics Forum* 29, 7 (2010), 2049–2054. 11
- [LJYS05] LI Z., JING Z., YANG X., SUN S.: Color transfer based remote sensing image fusion using non-separable wavelet frame transform. *Pattern Recognition Letters* 26 (2005), 2006–2014. 8
- [LLD07] LI Y., LIZHUANG M., DI W.: Fast colorization using edge and gradient constrains. *Proceedings of WSCG'07* (2007), 309–315. 12
- [LLW04] LEVIN A., LISCHINSKI D., WEISS Y.: Colorization using optimization. In *ACM SIGGRAPH 2004 Papers* (New York, NY, USA, 2004), SIGGRAPH '04, ACM, pp. 689–694. 2, 8, 10, 11, 12
- [Low04] LOWE D.: Distinctive image features from scale-invariant keypoints. *Int'l J. Computer Vision* 60, 2 (2004), 91–110. 3
- [LRFH13] LIN S., RITCHIE D., FISHER M., HANRAHAN P.: Probabilistic color-by-numbers: suggesting pattern colorizations using factor graphs. *ACM Transactions on Graphics* 32, 4 (2013), 37. 11, 13
- [LW07] LI G., WANG K.: Applying daytime colors to nighttime imagery with an efficient color transfer method. In *Proceedings of SPIE 6559, Enhanced and Synthetic Vision* (2007). 17
- [LWCO*07] LUAN Q., WEN F., COHEN-OR D., LIANG L., XU Y.-Q., SHUM H.-Y.: Natural image colorization. In *Rendering Techniques* (2007), Kautz J., Pattanaik S. N., (Eds.), Eurographics Association, pp. 309–320. 8, 12
- [LWQ*08] LIU X., WAN L., QU Y., WONG T.-T., LIN S., LEUNG C.-S., HENG P.-A.: Intrinsic colorization. *ACM Transactions on Graphics (TOG)* 27, 5 (2008), 152. 12
- [Mar84] MARKLE W.: The development and application of colorization®. *SMPTE journal* 93, 7 (1984), 632–635. 11
- [Mat95] MATSUDA Y.: *Color Design*. Asakura Shoten, 1995. 13
- [MDMS05] MANTIUK R., DALY S., MYSZKOWSKI K., SEIDEL H.-P.: Predicting visible differences in high dynamic range images - model and its calibration. In *Human Vision and Electronic Imaging X, IST SPIE's 17th Annual Symposium on Electronic Imaging* (2005), Rogowitz B. E., Pappas T. N., Daly S. J., (Eds.), vol. 5666, pp. 204–214. 19
- [MS03] MOROVIC J., SUN P.-L.: Accurate 3D image colour histogram transformation. *Pattern Recognition Letters* 24 (2003), 1725–1735. 8
- [MTH05] MA M., TIAN H., HAO C.: New method to quality evaluation for image fusion using gray relational analysis. *Optical Engineering* 44, 8 (2005), 087010–087010. 18
- [NC01] NANDA H., CUTLER R.: Practical calibrations for a real-time digital omnidirectional camera. *CVPR'01 Technical Sketch* (2001). 16, 17
- [NN05] NEUMANN L., NEUMANN A.: Color style transfer techniques using hue, lightness and saturation histogram matching. In *Computational Aesthetics in Graphics, Visualization and Imaging* (2005), pp. 111–122. 8
- [OAH11] O'DONOVAN P., AGARWALA A., HERTZMANN A.: Color compatibility from large datasets. *ACM Transactions on Graphics* 30, 4 (2011), 63. 13
- [OHSG12] OSKAM T., HORNING A., SUMNER R. W., GROSS M.: Fast and stable color balancing for images and augmented reality. In *Second International IEEE Conference on 3D Imaging, Modeling, Processing, Visualization and Transmission (3DIM-PVT)* (2012), pp. 49–56. 3, 4, 10, 11, 18
- [OZC12] OUNI S., ZAGROUBA E., CHAMBAH M.: Article: A new no-reference method for color image quality assessment. *International Journal of Computer Applications* 40, 17 (February 2012), 24–31. Published by Foundation of Computer Science, New York, USA. 18, 19

- [PCR10] POULI T., CUNNINGHAM D. W., REINHARD E.: Image statistics and their applications in computer graphics. In *Eurographics State-of-the-Art Reports* (2010). 7
- [PKD05] PITIÉ F., KOKARAM A., DAHYOT R.: 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* (Washington, DC, USA, 2005), vol. 2, IEEE Computer Society, pp. 1434–1439. 8, 9
- [PKD07] PITIÉ F., KOKARAM A., DAHYOT R.: Automated colour grading using colour distribution transfer. *Computer Vision and Image Understanding* 107, 1-2 (2007), 123–137. 2, 8
- [PM12] PRZEMYSŁAW S., MICHAL K.: Investigating human color harmony preferences using unsupervised machine learning. In *European Conference on Colour in Graphics, Imaging, and Vision* (2012), pp. 59–64. 13
- [PPC11] PAPADAKIS N., PROVENZI E., CASELLES V.: A variational model for histogram transfer of color images. *Image Processing, IEEE Transactions on* 20, 6 (2011), 1682–1695. 7
- [PR10] POULI T., REINHARD E.: Progressive histogram reshaping for creative color transfer and tone reproduction. In *Proc. of International Symposium on Non-Photorealistic Animation and Rendering* (2010), pp. 81–90. 5, 7, 8
- [PR11] POULI T., REINHARD E.: Progressive color transfer for images of arbitrary dynamic range. *Computers & Graphics* 35, 1 (2011), 67–80. 2, 5, 7, 8, 10
- [PRC13] POULI T., REINHARD E., CUNNINGHAM D. W.: *Image Statistics in Visual Computing*. CRC Press, 2013. 5, 7
- [QLZ13] QIAN Y., LIAO D., ZHOU J.: Manifold alignment based color transfer for multiview image stitching. In *Proceedings of the IEEE International Conference on Image Processing (ICIP'13)* (2013). 16
- [QWH06] QU Y., WONG T.-T., HENG P.-A.: Manga colorization. *ACM Transactions on Graphics (TOG)* 25, 3 (2006), 1214–1220. 12
- [RAC*04] REINHARD E., AKYÜZ A. O., COLBERT M., HUGHES C. E., O'CONNOR M.: Real-time color blending of rendered and captured video. In *Interservice/Industry Training, Simulation and Education Conference* (2004). 18
- [RAGS01] REINHARD E., ADHIKMIN M., GOOCH B., SHIRLEY P.: Color transfer between images. *IEEE Computer Graphics and Applications* 21, 5 (2001), 34–41. 2, 5, 8, 9, 11, 14, 17, 18
- [RCC98] RUDERMAN D. L., CRONIN T. W., CHIAO C. C.: Statistics of cone responses to natural images: implications for visual coding. *Journal of the Optical Society of America A* 15, 8 (1998), 2036–2045. 5, 7
- [Rey09] REYNOLDS D. A.: Gaussian mixture models., 2009. 9
- [RKAJ08] REINHARD E., KHAN E. A., AKYÜZ A. O., JOHNSON G. M.: *Color Imaging: Fundamentals and Applications*. A K Peters, Wellesley, 2008. 5
- [RPer] REINHARD E., POULI T.: Colour spaces for colour transfer. In *IAPR Computational Color Imaging Workshop*, Schettini R., Tominaga S., Trémeau A., (Eds.), vol. 6626 of *Lecture Notes in Computer Science*. Springer Berlin Heidelberg, 2011 (invited paper), pp. 1–15. 5, 7, 8
- [SA07] SENANAYAKE C., ALEXANDER D.: Colour transfer by feature based histogram registration. In *British Machine Vision Conference* (2007), British Machine Vision Association. 5, 7
- [SBŽ04] SYKORA D., BURIÁNEK J., ŽÁRA J.: Unsupervised colorization of black-and-white cartoons. In *Proceedings of the 3rd international symposium on Non-photorealistic animation and rendering* (2004), ACM, pp. 121–127. 12
- [SDZL09] SHI Y., DING Y., ZHANG R., LI J.: Structure and hue similarity for color image quality assessment. In *International Conference on Electronic Computer Technology* (2009), pp. 329–333. 18
- [SGSS08] SNAVELY N., GARG R., SEITZ S. M., SZELISKI R.: Finding paths through the world's photos. *ACM Transactions on Graphics* 27, 3 (2008), 15. 16
- [SH09] STUART G. W., HUGHES P. K.: *Towards understanding the role of colour information in scene perception using night vision devices*. Tech. Rep. DSTO-RR-0345, Defense Science and Technology Organisation, Fishermans Bend, Australia, 2009. 17
- [SJLL05] SUN S., JING Z., LIU G., LI Z.: Transfer color to night vision images. *Chinese Optics Letters* 3, 8 (2005), 448–450. 17
- [SJYC07] SHAO F., JIANG G., YU M., CHEN K.: A content-adaptive multi-view video color correction algorithm. In *IEEE International Conference on Acoustics, Speech and Signal Processing* (2007), vol. 1, pp. I-969. 4, 16
- [SJYH10] SHAO F., JIANG G.-Y., YU M., HO Y.-S.: Fast color correction for multi-view video by modeling spatial-temporal variation. *Journal of Visual Communication and Image Representation* 21 (2010), 392–403. 16
- [SM08] SAWANT N., MITRA N. J.: Color harmonization for videos. In *Proceedings of the 2008 Sixth Indian Conference on Computer Vision, Graphics & Image Processing* (Washington, DC, USA, 2008), ICVGIP '08, IEEE Computer Society, pp. 576–582. 13
- [SS02] SCHARSTEIN D., SZELISKI R.: A taxonomy and evaluation of dense two-frame stereo correspondence algorithms. *International Journal of Computer Vision* 47, 1 (2002), 7–42. 15
- [SSS06] SNAVELY N., SEITZ S. M., SZELISKI R.: Photo tourism: exploring photo collections in 3D. *ACM Transactions on Graphics* 25, 3 (2006), 835–846. 16
- [SVS*10] SAUVAGET C., VITTAUT J.-N., SUAREZ J., BOYER V., MANUEL S.: Automated colorization of segmented images based on color harmony. *JMPT* 1, 4 (2010), 228–244. 10, 11, 12
- [SWqLZ07] SHI S., WANG L., QI LIN W., ZHAO Y.: Color night vision based on color transfer in YUV color space. In *Proceedings of SPIE 6623, International Symposium on Photoelectronic Detection and Imaging* (2007). 17
- [TA05] TSAGARIS V., ANASTASSOPOULOS V.: Fusion of visible and infrared imagery for night color vision. *Displays* 26, 4–5 (2005), 191–196. 17
- [Tan04] TANIMOTO M.: Free viewpoint television — FTV. In *Proceedings of the Picture Coding Symposium* (2004). 16
- [TD11] THAKUR N., DEVI S.: Article: A new method for color image quality assessment. *International Journal of Computer Applications* 15, 2 (February 2011), 10–17. Published by Foundation of Computer Science. 18
- [TGTC02] TIAN G. Y., GLEDHILL D., TAYLOR D., CLARKE D.: Colour correction for panoramic imaging. In *Proceedings of Information Visualisation* (2002), pp. 483–488. 2, 16, 17
- [TH12] TOET A., HOGERVORST M. A.: Progress in color night vision. *Optical Engineering* 51, 1 (2012). 17
- [THvSD12] TOET A., HOGERVORST M. A., VAN SON R., DIJK J.: Augmenting full colour-fused multi-band night vision imagery with synthetic imagery in real-time. *International Journal of Image and Data Fusion* 2, 4 (2012), 287–308. 17
- [TISK10] TEHRANI M. P., ISHIKAWA A., SAKAZAWA S., KOIKE A.: Iterative colour correction of multicamera systems using corresponding feature points. *Journal of Visual Communication and Image Representation* 21, 5-6 (2010), 377 – 391. 3

- [TIWA97] TOET A., IJSPEERT J. K., WAXMAN A. M., AGUILAR M.: Fusion of visible and thermal imagery improves situational awareness. *Displays* 18, 2 (1997), 85–95. 17
- [TJT05] TAI Y., JIA J., TANG C.: Local color transfer via probabilistic segmentation by expectation-maximization. In *Proc. of CVPR* (2005), vol. 1, pp. 747–754. 9
- [TJT07] TAI Y.-W., JIA J., TANG C.-K.: Soft color segmentation and its applications. *Pattern Analysis and Machine Intelligence, IEEE Transactions on* 29, 9 (2007), 1520–1537. 9
- [TMI02] TOKUMARU M., MURANAKA N., IMANISHI S.: Color design support system considering color harmony. In *Fuzzy Systems, 2002. FUZZ-IEEE'02. Proceedings of the 2002 IEEE International Conference on* (2002), vol. 1, IEEE, pp. 378–383. 13
- [TMW10] TANG Z., MIAO Z., WAN Y.: Image composition with color harmonization. In *Image and Vision Computing New Zealand (IVCNZ), 2010 25th International Conference of* (Nov 2010), pp. 1–8. 13
- [Toe03] TOET A.: Natural colour mapping for multiband night-vision imagery. *Information Fusion* 4, 3 (2003), 155–166. 8, 17
- [TZMS04] THEOBALT C., ZIEGLER G., MAGNOR M., SEIDEL H.-P.: Model-based free-viewpoint video: Acquisition, rendering, and encoding. In *Proceedings of the Picture Coding Symposium* (2004). 16
- [UB87] UCHIKAWA K., BOYNTON R. M.: Categorical color perception of japanese observers: Comparison with that of americans. *Vision Research* 27, 10 (1987), 1825–1833. 9
- [UES01] UYTENDAELE M., EDEN A., SKELISKI R.: Eliminating ghosting and exposure artifacts in image mosaics. In *Computer Vision and Pattern Recognition, 2001. CVPR 2001. Proceedings of the 2001 IEEE Computer Society Conference on* (2001), vol. 2, IEEE, pp. II-509. 16, 17
- [Var98] VARGA J. T.: *Evaluation of operator performance using true color and artificial color in natural scene perception*. Tech. Rep. AD-A363036, Naval Postgraduate School, Monterey, CA, 1998. 17
- [VdNJ*07] VIEIRA L. F. M., DO NASCIMENTO E. R., JR F. A. F., CARCERONI R. L., VILELA R. D., DE A. ARAÚJO A.: Fully automatic coloring of grayscale images. *Image and Vision Computing* 25, 1 (2007), 50–60. 12, 18, 19
- [vGE40] VON GOETHE J. W., EASTLAKE C. L.: *Goethe's Theory of Colours: Translated from the German*. Murray, 1840. 13
- [VMPX04] VETRO A., MATUSIK W., PFISTER H., XIN J.: Coding approaches for end-to-end 3-D TV systems. In *Proceedings of the Picture Coding Symposium* (2004). 16
- [WAM02] WELSH T., ASHIKHMIN M., MUELLER K.: Transferring color to greyscale images. *ACM Transactions on Graphics* 21, 3 (2002), 277–280. 10, 11
- [WBSS04] WANG Z., BOVIK A., SHEIKH H., SIMONCELLI E.: Image quality assessment: From error visibility to structural similarity. *IEEE Trans. on Image Processing* 13, 4 (2004), 600–612. 18
- [WDK*13] WU F., DONG W., KONG Y., MEI X., PAUL J.-C., ZHANG X.: Content-based colour transfer. *Computer Graphics Forum* 32, 1 (2013), 190–203. 10, 16
- [Wei01] WEISS Y.: Deriving intrinsic images from image sequences. In *Computer Vision, 2001. ICCV 2001. Proceedings. Eighth IEEE International Conference on* (2001), vol. 2, IEEE, pp. 68–75. 12
- [WHCM08] WEN C. L., HSIEH C. H., CHEN B. Y., MING O.: Example-based multiple local color transfer by strokes. *Computer Graphics Forum* 27, 7 (2008), 1765–1762. 2, 8, 10
- [WHH06] WANG C. M., HUANG Y. H., HUANG M. L.: An effective algorithm for image sequence color transfer. *Mathematical and Computer Modeling* 44 (2006), 608–627. 8, 14, 15
- [WS00] WYSZECKI G., STILES W. S.: *Color science: Concepts and methods, quantitative data and formulae*, 2nd ed. John Wiley and Sons, New York, 2000. 5
- [WSW10] WANG Q., SUN X., WANG Z.: A robust algorithm for color correction between two stereo images. In *Proc. of ACCV* (2010), vol. 5995, pp. 405–416. 4, 15
- [WYW*10] WANG B., YU Y., WONG T.-T., CHEN C., XU Y.-Q.: Data-driven image color theme enhancement. *ACM Transactions on Graphics* 29, 6 (2010), 146. 2, 10
- [WZJ*07] WANG L., ZHAO Y., JIN W., SHI S., WANG S.: Real-time color transfer system for low-light level visible and infrared images in YUV color space. In *Proceedings of the SPIE* (2007), vol. 6567, p. 65671G. 8
- [XADR13] XUE S., AGARWALA A., DORSEY J., RUSHMEIER H.: Learning and applying color styles from feature films. *Computer Graphics Forum* 32 (2013). 2, 14, 18
- [XM06] XIAO X., MA L.: Color transfer in correlated color space. In *Proceedings of the 2006 ACM international conference on Virtual reality continuum and its applications* (2006), ACM, pp. 305–309. 2, 5, 6, 8
- [XM09] XIAO X., MA L.: Gradient-preserving color transfer. *Computer Graphics Forum* 28, 7 (2009), 1879–1886. 2, 5, 7, 8, 18, 19
- [XM10] XU W., MULLIGAN J.: Performance evaluation of color correction approaches for automatic multi-view image and video stitching. In *Proc. of CVPR* (2010), pp. 263–270. 16, 18
- [XYJ13] XU L., YAN Q., JIA J.: A sparse control model for image and video editing. *ACM Trans. Graph.* 32, 6 (Nov. 2013), 197:1–197:10. 11
- [XZL09] XIANG Y., ZOU B., LI H.: Selective color transfer with multi-source images. *Pattern Recognition Letters* 30, 7 (2009), 682–689. 8, 9
- [XZZY05] XU S., ZHANG Y., ZHANG S., YE X.: Uniform color transfer. In *IEEE International Conference on Image Processing* (2005), pp. III – 940–943. 8, 9
- [YC04] YIN J., COOPERSTOCK J. R.: Color correction methods with applications to digital projection environments. *Journal of the WSCG* 12, 1 (2004). 8
- [YKK*07] YAMAMOTO K., KITAHARA M., KIMATI H., YENDO T., FUJII T., TANIMOTO M., SHIMIZU S., KAMIKURA K., YASHIMA Y.: Multiview video coding using view interpolation and color correction. *IEEE Transactions on Circuits and Systems for Video Technology* 17, 11 (2007), 1436–1449. 3
- [YO08a] YAMAMOTO K., OI R.: Color correction for multi-view video using energy minimization of view networks. *International Journal of Automation and Computing* 5, 3 (2008), 234–245. 3, 18
- [YO08b] YAMAMOTO K., OI R.: Color correction for multi-view video using energy minimization of view networks. *International Journal of Automation and Computing* 5, 3 (2008), 234–245. 16
- [YPCL13] YOO J.-D., PARK M.-K., CHO J.-H., LEE K. H.: Local color transfer between images using dominant colors. *Journal of Electronic Imaging* 22, 3 (2013), 033003–033003. 9
- [YS06] YATZIV L., SAPIRO G.: Fast image and video colorization using chrominance blending. *Image Processing, IEEE Transactions on* 15, 5 (2006), 1120–1129. 12
- [YXM*12] YUAN Y., XU H., MIAO Z., LIU F., ZHANG J.,

- CHANG B.: Real-time infrared and visible image fusion system and fusion image evaluation. In *IEEE Symposium on Photonics and Optoelectronics* (2012), pp. 1–4. 17
- [YYF*07] YAMAMOTO K., YENDO T., FUJII T., TANIMOTO M., SUTER D.: Colour correction for multiple-camera system by using correspondences. *Journal of the Institute of Image Information and Television Engineers* 61, 3 (2007), 213–222. 3, 16
- [ZDB12] ZHENG Y., DONG W., BLASCH E. P.: Qualitative and quantitative comparisons of multispectral night vision colorization techniques. *Optical Engineering* 51, 8 (2012). 17, 18
- [ZHHE05] ZHENG Y., HANSEN B. C., HAUN A. M., ESSOCK E. A.: Coloring night-vision imagery with statistical properties of natural colors by using image segmentation and histogram matching. In *Proceedings of SPIE 5667, Color Imaging X: Processing, Hardcopy, and Applications* (2005), vol. 107. 17
- [ZW94] ZAHIB R., WOODFILL J.: Non-parametric local transforms for computing visual correspondence. In *Proceedings of the European Conference on Computer Vision* (1994), pp. 151–158. 4
- [ZWJS07] ZHAO Y., WANG L., JIN W., SHI S.: Colorizing biomedical images based on color transfer. In *IEEE/ICME International Conference on Complex Medical Engineering* (2007), pp. 820–823. 8

Author Biographies

Hasan Sheikh Faridul is a postdoctoral researcher in the Security & Content Protection Labs of Technicolor Research & Innovation, France. He received his PhD in 2014 from the University of Saint Etienne, France. Before joining Technicolor, he completed the Erasmus Mundus masters program titled ‘Color in Informatics & Media Technology (CIMET)’.

Tania Pouli is a Researcher at Technicolor Research & Innovation in Rennes. After completing her PhD in 2011 at the University of Bristol, she moved to the Max Planck Institute for Informatics as a postdoc, focusing on high dynamic range and color imaging until 2013. She has organized and co-presented a number of successful courses and tutorials at SIGGRAPH, SIGGRAPH Asia and Eurographics and she is the lead author of the book ‘Image Statistics in Visual Computing’, published by CRC Press.

Christel Chamaret is currently working as a Senior Scientist at Technicolor Research & Innovation in Rennes. Her research focuses on color image processing and more particularly on color harmonization and color grading. She has previously worked on visual attention models, quality metrics as well as aesthetic models. She has also designed and conducted a number of user studies to validate computational models and image processing algorithms. She was previously involved in video compression projects at Philips Research, The Netherlands and INRIA Rennes, France. She holds two master’s degree defended in 2003 from the University of Nantes and Ecole Centrale de Nantes, France.

Jurgen Stauder received his PhD degree in 1999 from the University of Hannover (Germany) in the field of computer vision. He then stayed with INRIA (France) for 18 months before joining Technicolor research labs in France (formerly Thomson). He is guest lecturer at the University of Paris 13 for applied color science. His research interests are computer vision, color science, and computer graphics with application to the media industry.

Alain Trémeau is Professor in Computational Color Imaging at the Université Jean Monnet – Saint-Etienne, France. He was the head of LIGIV (<http://www.ligiv.org>), a research laboratory working on computer Graphics, Vision Engineering and Color Imaging Science till 2008. Since 2009, he is the head of the Color Content Aware Image Processing Group of the Laboratoire Hubert Curien CNRS UMR 5516, Saint-Etienne, France. He has written numerous papers, book chapters on Computational Color Imaging and Processing.

Erik Reinhard is a Principal Scientist at Technicolor R&I since July 2013 prior to holding various academic positions at universities and research institutes in Europe and North America. He was founder and editor-in-chief of ACM Transactions on Applied Perception, and authored books on high dynamic range imaging, color imaging, computer graphics and natural image statistics. He has published more than 100 papers in these areas and was member of more than 50 program committees. He was program co-chair of 6 conferences and workshops, including the Eurographics Symposium on Rendering 2011 and delivered keynotes at 4 conferences including Eurographics 2010.