

Performance Evaluation of Color Correction Approaches for Automatic Multi-view Image and Video Stitching

Wei Xu and Jane Mulligan
University of Colorado at Boulder
Boulder, CO 80302-0430 USA
{Wei.Xu, Jane.Mulligan}@Colorado.edu

Abstract

Many different automatic color correction approaches have been proposed by different research communities in the past decade. However, these approaches are seldom compared, so their relative performance and applicability are unclear. For multi-view image and video stitching applications, an ideal color correction approach should be effective at transferring the color palette of the source image to the target image, and meanwhile be able to extend the transferred color from the overlapped area to the full target image without creating visual artifacts. In this paper we evaluate the performance of color correction approaches for automatic multi-view image and video stitching. We consider nine color correction algorithms from the literature applied to 40 synthetic image pairs and 30 real mosaic image pairs selected from different applications. Experimental results show that both parametric and non-parametric approaches have members that are effective at transferring colors, while parametric approaches are generally better than non-parametric approaches in extendability.

1. Introduction and Motivation

Color correction or color balancing in automatic multi-view image and video stitching is the process of correcting the color differences between neighboring views which arise due to different exposure levels and view angles. Compared to the other major steps in image stitching of registration and blending, color correction has received less attention and relatively simpler treatment. Image blending, which has a similar end effect to color correction, has concealed the the role of the latter. Only recently, with the growing demand and popularity of high definition images and video, have people begun to recognize that image blending alone cannot remove all the color difference between different views under all situations (see Figure 1 for an example).

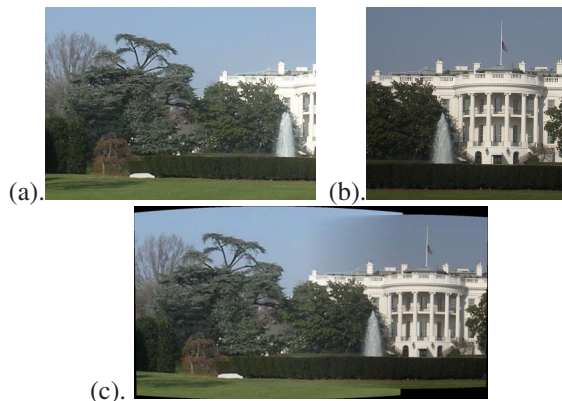


Figure 1. An example mosaic image pair that have dramatic color difference. (a) and (b) are the mosaic image pair, (c) is the stitching result of Autostitch [3, 4] with the multi-band image blending [6] functionality enabled. (Note that Autostitch projects the stitching result into the cylindrical plane). Color incoherence is still obvious in the scene even after multi-band image blending is performed.

In the computer vision and multi-view video processing communities, the initial efforts on solving the color balancing problem for multi-view stitching used exposure compensation (or gain compensation) [21, 3, 30]. This approach adjusts the intensity gain level of component images to compensate for appearance differences caused by different exposure levels. Although it works for some cases, it may fail to completely compensate for color difference between different views when the lighting conditions vary dramatically. Later work compensated for differences using all three color channels rather than via the single intensity channel [29, 12, 15, 16, 36, 34]. At the same time the image processing and computer graphics communities were developing similar color manipulation methods they called *color transfer techniques* (e.g., [25, 28, 33, 32, 24]). Technically speaking, there is no difference between color balancing and color transfer, exception the latter generally does not have to be restricted to the overlapped area — if we restrict color transfer techniques to operate using only

information from the overlapped area, then they can be easily used to solve the color balancing problem for multi-view image and video stitching. Actually, one of the current research topics in color transfer is how to exploit local spatial constraints/matches to guide the color transfer (e.g., the local color transfer technique [28]). From this perspective researches on color balancing and color transfer exhibit significant overlap.

In this paper, we perform a unified evaluation of color balancing and color transfer techniques in the context of automatic multi-view image and video stitching. In this context, we use the term “color correction approaches” to refer to the union of applicable color balancing and color transfer approaches. And we use the terms “color correction”, “color alteration”, “color transfer”, “color mapping”, and “color balancing” interchangeably in the following parts of this paper.

There is a trade-off of effectiveness and extendability for any color correction approach: *effectiveness* measures how genuinely the color mapping function (which is usually estimated from the overlapped area) transfers the color palette of the source image to the target image, and *extendability* measures how well this mapping extends to the non-overlapped areas of the target image without creating visual artifacts. We want to determine through evaluation how different approaches behave with respect to effectiveness and extendability given mosaic image pairs captured under different conditions. We focus on automatic approaches in our evaluation and exclude those approaches that need human invention or guidance (e.g. [22]) to complete the task. We also focus on techniques operating in the image domain for maximum generality, and thus those approaches that request pre-calibration information to operate in the radiance or irradiance domain (e.g. [8]) are not included. The inputs to our evaluation system are images of different appearance and unknown capture conditions, so those approaches for calibration of multi-view camera systems (e.g. [13, 35]) are also excluded. Our evaluation results should be of interest not only to the computer vision community, but also to other communities including computer graphics, image processing, and multi-view video processing.

This paper is organized as follows. In section 2 we present a set of state of the art color correction techniques. Section 3 gives the details of the evaluation experiment setup including the selection of approaches for comparison, test image datasets and parameter settings. In section 4 we present the experimental results. And in section 5 we summarize our results and conclusions.

2. Color Correction Approaches

The essence of all color correction algorithms is transferring the color (or intensity) palette of a source image to a target image. In the context of multi-view image and video

stitching, the source image corresponds to the view selected as the reference by the user, and the target image corresponds to the image whose color is to be corrected. Rather than giving an historic review of color balancing and color transfer respectively, here we will categorize the techniques used according to their basic approaches. At the highest level there are two classes of color correction approaches: parametric and non-parametric.

2.1. Model-based parametric approaches

2.1.1 Global color transfer

Model-based color correction techniques are parametric, and include global and local modeling approaches. Global modeling approaches assume the relation between the color of the target image and that of the source image can be described by a transform: $I_s = M * I_t$, where M is a 3x3 matrix representing the mapping of the three color channels. Here M can be a diagonal matrix, an affine matrix or an arbitrary 3x3 matrix, corresponding to the diagonal model, the affine model and the linear model respectively [29, 10]. Various approaches can be used to estimate M , depending on applications and inputs.

Exposure compensation (or gain compensation) is the technique initially employed to address the color balancing problem in panorama stitching where the inputs are partially overlapped images. Nanda and Cutler first incorporated gain adjustment as part of the “AutoBrihtness” function of their multi-head panoramic camera called RingCam [21]. Then, Brown and Lowe employed it in their well-known automatic panorama stitching software “Autostitch” [3, 4], and Uyttendaele *et al.* [30] applied it on a block-by-block basis followed by spline interpolation of piece-wise gain parameters. Since the gain compensation technique only operates in the intensity channel but not in full color space, it actually corresponds to a particular diagonal model where the values in the main diagonal of M have to be same. This particular diagonal model was also adopted in some later work that combines exposure compensation and vignetting correction [11, 19].

Other more general approaches in global modeling include Tian *et al.*’s work [29] using histogram mapping over the overlap area to estimate the transformation matrix M , and Zhang *et al.*’s work [36] using the principal regions mapping to estimate M where the highest peaks in the hue histogram are designated as principal regions.

Given two general images where there is no overlap, Reinhard *et al.* [25] proposed a linear transformation based on the simplest statistics of global color distributions of two images: $g(C_t) = \mu_s + \frac{\sigma_s}{\sigma_t}(C_t - \mu_t)$, where (μ_s, σ_s) and (μ_t, σ_t) are the mean and standard deviation of the global color distributions of the source and target images in the uncorrelated $l\alpha\beta$ color space and C_t is the color

of a pixel in the target image. This work was widely used as the baseline approach by other color correction approaches [36, 33, 23, 28]. Xiao *et al.* [33] proposed an ellipsoid mapping scheme which extends Reinhard *et al.*'s work to correlated RGB color space. An *et al.* [2] discussed the linear transformation in YUV space.

2.1.2 Local color transfer

Global modeling usually provides only a rough mapping between the color of two images. In practice, many applications require a more deliberate mapping, which suggests local color transfer approaches. Tai *et al.* [28] proposed a local color transfer scheme based on probabilistic image segmentation and region mapping using Gaussian mixture models (GMM) and the EM algorithm. Xiang *et al.* [32] improved this work in the case that multiple source images are available for selection. For both of these approaches, after the local regions of the two images are matched, region-wise color transfer is performed using a weighted version of Reinhard's method [25].

2.2. Modeless non-parametric approaches

Non-parametric methods assume no particular parametric format of the color mapping function and most of them use a look-up table to directly record the mapping of the full range of color/intensity levels. This look-up table is usually computed from 2D joint histogram of image feature correspondences or pixel pairs in the overlapped area of two images. Two points need to be kept in mind when one estimates a mapping function from the histogram: First, robust methods are usually needed because the data are prone to outliers and noise due to different lighting conditions, capturing angles and reflection properties of scene objects. Second, the monotonicity property of the color/intensity levels in the estimated mapping function has to be maintained. All of the existing non-parametric mapping approaches can be distinguished from each other in how they accomplish these two points.

Yamamoto *et al.* [34] proposed using the joint histogram of SIFT feature matches between two neighboring views in a multi-view camera network. An energy minimization scheme in the 2D histogram space was proposed to get a robust estimation of the color mapping function and meanwhile maintain its monotonicity.

Jia and Tang [15] proposed a two-stage approach to handle robustness and monotonicity separately: in the first stage 2D tensor voting was used to suppress the noise and fill in the data gaps (i.e. places where no correspondences are available for some color levels). This produced an initial estimate of the mapping function. In the second stage, a heuristic local adjustment scheme was proposed to adjust the initial estimate and make the mapping monotoni-

cally increasing. In other work by the same authors [14], a Bayesian framework was used to recover the mapping function between a poorly exposed image and a blurred image.

Similar to Jia's work, Kim and Pollefeys [17] proposed a likelihood maximization scheme for robust estimation of the Brightness Transfer Function (BTF) from the 2D joint intensity histogram of two overlapped images. The approach operated in each of the three color channels separately. Dynamic programming was used to find a robust estimate under the monotonicity constraint. The estimated BTF was further used to estimate and remove the exposure difference and vignetting effect in the images.

Fecker *et al.* [9] proposed the use of cumulative color histogram mapping for color correction. They used a closest neighbor mapping scheme to select the corresponding color level of the source image to each level of the target. Using cumulative histogram-based mapping automatically satisfies the monotonicity constraint. The authors also suggested some special adjustment to the mapping of the border bins (i.e. the first and last bin) to avoid possible visual artifacts.

Pitié *et al.* [23, 24] proposed a quite different approach for color correction, called iterative color distribution transfer. This approach does not employ any explicit mapping function of the global color distribution, but relies on applying a sequence of simple conversions with respect to randomly projected marginal color distributions. Specifically, it treats the colors of an image as a distribution in a high dimensional space (usually of order 3), and repeatedly projects this high dimensional distribution into a series of random 1D marginal distributions using the Radon Transform. The color distribution of the target image is converted to that of the source image by repeatedly mapping its 1D marginal distributions to those of the source image until convergence. In [24], a post-processing technique for reducing the grain artifacts over the converted image was also proposed.

2.3. Previous evaluation work

Although various color correction techniques have been proposed in the last decade, there does not exist an extensive evaluation comparing the performance of these approaches. Most authors either only demonstrated their systems on a few self-selected example images or compared with very simple baseline approaches.

3. Evaluation Setup

3.1. Selection of Approaches

Table 1 shows a list of nine color correction algorithms we selected for performance evaluation and comparison. These selected algorithms are either a widely used standard baseline in color alteration (e.g., [25]), or represent the

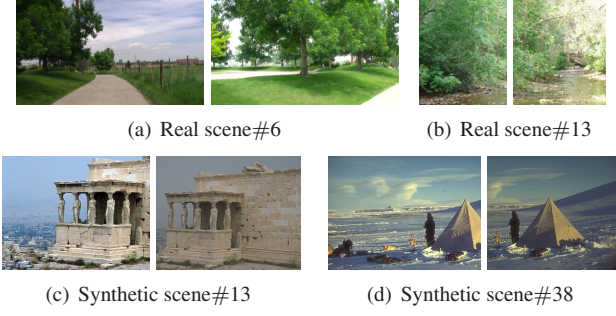


Figure 2. Example test image pairs. (a)-(b) example real mosaic image pairs, (c)-(d) example synthetic image pairs. For each pair, left is source image and right is target image.

most recent progress in color correction techniques. The selection includes both model-based parametric approaches and modeless non-parametric approaches, both approaches using global image statistics for estimating the color mapping function and approaches using local matches in the overlapped area, as well as approaches operating in various color spaces.

3.2. Test Image Sets

Both synthetic images and mosaic pairs selected from real image and video stitching tasks are used in our evaluation. The synthetic test image set includes 40 source/target image pairs. Each image pair is created in three steps: First, we selected images of poor exposure-levels from the Berkeley BSDS300 image segmentation dataset [20]. Then, for each of these selected images/frames an image processing tool [26] was used to auto adjust its color. This produces a new image of the same scene but differing in color properties (see Fig. 2 (c)-(d)). Finally, we visually compared the quality of the original image and the new image, and designated a clip from the image with better quality as the source image and another clip from the other image as the target image (If the two images are of similar quality then the assignment is random). When a color correction algorithm is executed with these synthetic image pairs, its ability to increase the quality of an image by altering its color distribution is thus tested.

The real test image set includes 30 example mosaic image pairs collected from various sources, including image frames from multi-view video applications, scenic or object photos taken with/without flash lighting or under different capture modes, and aerial image clips of the same place but at a different time (Fig. 2 (a)-(b)). For each of these real image pairs, the image which looks more natural is designated the source image and the other one the target image.

Each source/target image pair in the test image sets is partially overlapped and Autostitch is used to find the geo-

metric registration between them before color correction is performed. Figure 2 shows a few examples from the two test image sets.

3.3. Evaluation Criteria

A recently proposed theory on image quality evaluation is that from the perceptual point of view the goodness of a color altered target image should show both color coherence and structural coherence, since color correction may not only change the color of the target image but also the structure [31]. Based on this theory, we propose a criterion to evaluate the quality of transferring the color of a source image s to a target image t , which results in a converted image r . The proposed evaluation criterion includes two components: color similarity $CS(r, s)$ between the source image s and the transferred image r , and structure similarity $SS(r, t)$ between the target image t and the transferred image r . The color similarity $CS(r, s)$ is defined as:

$$CS(r, s) = PSNR(\hat{r}, \hat{s}) \quad (1)$$

where $PSNR = 20 * \log_{10}(L/RMS)$ is the peak signal-to noise ratio [7]. L is the largest possible value in the dynamic range of an image, and RMS is the root mean square difference between two images. \hat{r} and \hat{s} are the overlapped area of r and s respectively. The higher the value of $CS(r, s)$ the more similar the color between the two images r and s .

The structure similarity $CS(r, t)$ is defined as:

$$SS(r, t) = SSIM(r, t) \quad (2)$$

where $SSIM(r, t) = \frac{1}{N} \sum_{j=1}^N SSIM(a_j, b_j)$ is the Structural SIMilarity (SSIM) index [31]. N is the number of local windows for an image, a_j and b_j are the image contents at the j th local window of r and t respectively. SSIM itself is defined as a combination of luminance, contrast and structure components [31]:

$$SSIM(a, b) = [l(a, b)]^\alpha \cdot [c(a, b)]^\beta \cdot [s(a, b)]^\gamma \quad (3)$$

where $l(a, b) = \frac{2\mu_a\mu_b+A_1}{\mu_a^2+\mu_b^2+A_1}$, $c(a, b) = \frac{2\sigma_a\sigma_b+A_2}{\sigma_a^2+\sigma_b^2+A_2}$, $s(a, b) = \frac{\sigma_{ab}+A_3}{\sigma_a\sigma_b+A_3}$. μ_a and μ_b are the mean luminance values of windows a and b respectively; σ_a and σ_b are the standard variance of the of windows a and b respectively; σ_{ab} is the auto-covariance between a and b . Here A_1 , A_2 and A_3 are small constants to avoid divide-by-zero error, α , β and γ control the weighting between the three components. In our implementation we use the default settings recommended by [31]: $A_1 = (0.01 * L)^2$, $A_2 = (0.03 * L)^2$, $A_3 = A_2/2$, $L = 255$ for images of dynamic range $[0, 255]$ and $\alpha = \beta = \gamma = 1$. The higher the value of $SS(r, t)$ the less difference between the structure of r and t , and $SS(r, t) = 1$ if there is no structure difference.

Index#	Name of Approach	Reference	Parametric/Non-Parametric	Local/Global	color space
1	gain compensation	Brown 2007 [4]	parametric	local	intensity
2	principal regions mapping	Zhang 2004 [36]	parametric	local	CIECAM97
3	tensor voting in joint image space	Jia 2005 [16]	non-parametric	local	RGB
4	global color transfer	Reinhard 2001 [25]	parametric	global	$l\alpha\beta$
5	global color transfer in correlated color space	Xiao 2006 [33]	parametric	global	RGB
6	local color transfer	Tai 2005 [28]	parametric	local	$l\alpha\beta$
7	brightness transfer function	Kim 2008 [17]	non-parametric	local	RGB
8	iterative color distribution transfer	Pitie 2005 [23]	non-parametric	global	RGB
9	cumulative histogram mapping	Fecker 2008 [9]	non-parametric	local	YCbCr

Table 1. Color correction approaches selected for performance evaluation and comparison.

The color and structure similarities measure the effectiveness and extendability of a color correction approach respectively. To give the reader a perception of these measures, Figure 3 shows a real example.

3.4. Pre-processing

One problem that must be considered in our evaluation is that vignetting may exist in the test images. The traditional approach is to solve vignetting removal and color correction in a whole, but recent progress allows the two to be decoupled so vignetting can be removed separately for a single image [37, 38]. Since the focus of this study is color correction, we used Zheng’s approach [38] to remove vignetting effects that might exist in the test images before using the images to test selected color correction approaches.

3.5. Implementation details and parameter settings

We downloaded the source code of the iterative color distribution transfer approach [23] and SSIM [31] from the authors’ websites and used them directly in our evaluation. We also implemented the other eight color correction approaches, and the proposed “color+structure” evaluation criterion using MATLAB 7.7.0 (R2008b). We used the open source code OpenTVF [1] for the 2D tensor voting technique in our implementation of the tensor voting-based color correction approach [16].

In our implementation, most of the approaches and evaluation criteria use the same parameters as stated in the original papers. The only exception is for the principal regions mapping approach [36]: the original paper prefers to use three principle regions to construct order-2 polynomial mapping functions, while in our implementation we conservatively used only two principle regions which simplifies the mapping functions to an affine model. This is because in practice we found higher degree mapping functions are more prone to out-of-gamut errors (i.e. some transferred colors go out of the display range of the RGB color model) [18].

4. Evaluation Results

4.1. Experimental data and analysis

We tested all of the nine selected approaches on both the synthetic image set and the real image set by computing both CS and SS scores of the correction results. Figure 4 and Table 2 show statistics of these scores over the 40 synthetic image pairs and 30 real image pairs respectively. Here ‘alg#10’ is a “virtual baseline approach” added in for purely comparison purposes. It corresponds to “no correction performed”, that is, computing the CS/SS scores between the source image and the target image directly with no color adjustment.

First of all, the data tell us the synthetic image set and the real image set are pretty different. The CS score range of alg#10 (the baseline) is 20.5134 ± 5.2041 for the synthetic set, and is 18.4429 ± 5.4203 for the real image set. Since CS is built upon PSNR which is built upon RMS errors, this data shows there are more intra-pair and inter-pair differences in the real image set, and thus it is more challenging than the synthetic image set. Considering this factor that our test sets are different and our goal is to serve real applications, in the following we use the data from the real image set as main reference. From the data we find:

As a very simple and widely used approach in image and mosaic stitching, alg#1 (gain compensation) performs pretty well: this is reflected as relatively high (rank 4) mean SS scores with small variance, and good (rank 2) CS scores.

The biggest problem for alg#2 (principle regions mapping) is stability. It has pretty good performance on the synthetic image set, but very poor performance on the structural score on real image set. The possible explanation for this is that it simply designates peaks in hue histograms as principle regions, which might be too simple to work well for real scenes with complex contents.

Alg#5 (global color transfer in correlated color space) and alg#6 (local color transfer) are both variants to alg#4 (global color transfer). Compared to alg#4 which operates in uncorrelated $l\alpha\beta$ color space, alg#5 operates in correlated RGB color space, which leads to deteriorations in both color correction effectiveness and extendability. Alg#6 makes use of local spatial information to guide color transfer, which leads to a gain in color correction effectiveness

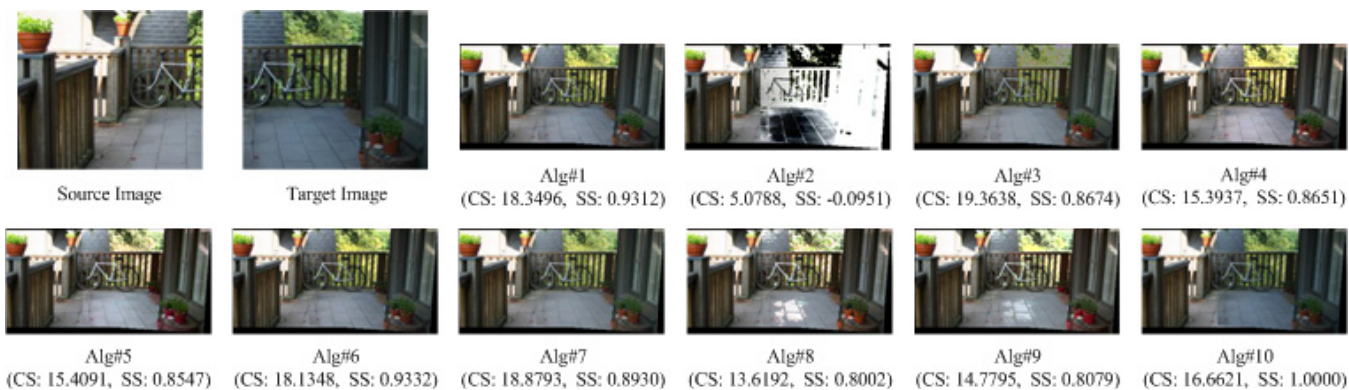


Figure 3. Color correction results on a real test image pair (real scene#16) and the corresponding CS and SS scores. The baseline approach alg#10 shows the basic CS and SS scores corresponding to directly stitching the source and target image together without taking any color correction measures. In this example, alg#8 and alg#9 produce out-of-gamut errors in the overlapped area, so they obtain lower CS scores than the baseline. alg#1 and alg#3 make the color in the overlapped area more balanced (seams become less obvious), so they obtain higher CS scores than the baseline. However, comparing to alg#1, the structure of the transferred image of alg#3 (over the tree area above the bike) is more distorted with respect to the original target image, so it gets a lower SS score than does alg#1. Note here and in our evaluation no advanced image blending techniques but simple averaging is applied over the overlapped area in order for color correction effect to be evaluated independently. This source/target image pair is originally from [5] and used in our evaluation with permission.

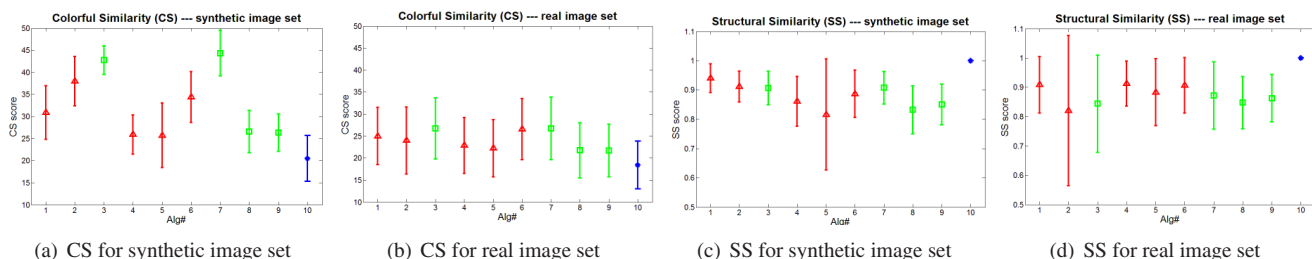


Figure 4. Errorbar statistics of CS and SS scores for all of the nine selected color correction algorithms. Red triangle: parametric approach; Green Square: non-parametric approach; Blue Star: the baseline approach.

and similar extendability.

Alg#3 (tensor voting in joint image space) and alg#7 (brightness transfer function) are representatives of non-parametric approaches that build the color transfer function upon exact mapping of the full range of color/intensity levels. Compared to alg#8 (iterative color distribution transfer) and alg#9 (cumulative histogram mapping) that use implicit or rough mapping, they show not only much better color correction effectiveness, but also similar (or slightly better) extendability.

From the perspective of class-level comparison, non-parametric approaches have better color transfer effectiveness but less extendability than parametric ones in general. But this is by no means absolute for individual approaches: some parametric approaches that make use of local information, such as alg#6 (local color transfer), have quite close performance in color transfer effectiveness as that of the most capable non-parametric approaches such as alg#3 (tensor voting in joint image space) and alg#7 (brightness transfer function).

4.2. Analysis of the worst cases

There are two questions of interest to us: 1) is there a common factor in practice that may affect the performance of all the approaches, and 2) what is the most challenging scene for all of the approaches. To answer these questions, we have found the five real scenes (i.e. image pairs) on which the nine selected approaches achieve the lowest average CS scores and SS scores (see Table 3).

Table 3 shows that scenes #6, #13 and #16 are in the worst case lists of both CS and SS scores. Especially, on scene #13 and #6 the nine selected approaches on average suffer deterioration of both CS and SS scores. Figures 2(a), 2(b) and 3 show these three scenes. It is easy to discover that all of them are affected by extreme lighting conditions: The target image of scene #6 contains saturated areas, which may cause problems in both calculating the color mapping function and in extending this mapping to the non-overlapped area. Scene #13 has saturated parts in the non-overlapped region of the target image, which may cause problems when extending the color mapping func-

		Alg#1	Alg#2	Alg#3	Alg#4	Alg#5	Alg#6	Alg#7	Alg#8	Alg#9	Alg#10
syn	μ_{CS}	30.8825	38.0145	42.7686	25.9149	25.6988	34.3836	44.3281	26.5597	26.3066	20.5134
	σ_{CS}	6.0941	5.5917	3.2322	4.4301	7.2960	5.7530	5.1852	4.7586	4.2376	5.2041
real	μ_{CS}	24.9928	23.9749	26.6989	22.8586	22.2535	26.5897	26.7329	21.7544	21.6938	18.4429
	σ_{CS}	6.5170	7.6458	6.9387	6.3581	6.5077	6.9391	7.1134	6.2852	5.9989	5.4203
syn	μ_{SS}	0.9404	0.9115	0.9064	0.8609	0.8159	0.8866	0.9075	0.8320	0.8507	1.0000
	σ_{SS}	0.0492	0.0523	0.0573	0.0847	0.1897	0.0808	0.0554	0.0818	0.0695	0
real	μ_{SS}	0.9085	0.8205	0.8440	0.9127	0.8830	0.9064	0.8716	0.8478	0.8625	1.0000
	σ_{SS}	0.0962	0.2561	0.1660	0.0771	0.1144	0.0947	0.1149	0.0893	0.0810	0

Table 2. Mean (μ) and standard deviation (σ) statistics of CS and SS scores for the nine selected color correction algorithms.

scene#	CS					SS				
	13	16	7	6	19	13	6	23	16	22
average score	13.0501	15.4465	15.9166	16.1193	17.5904	0.5909	0.6690	0.7155	0.7622	0.7627
baseline score	16.7105	11.9117	12.6996	16.6621	12.9240	1.0000	1.0000	1.0000	1.0000	1.0000
gain (percentage)	-28.05%	22.88%	20.21%	-3.37%	26.53%	-40.91%	-33.10%	-28.45%	-23.78%	-23.73%

Table 3. The real test scenes on which the nine selected approaches achieves the lowest average CS and SS scores.

tion to this area. The saturated area in scene#16 is in the non-overlapped part of the source image, which might not cause any problems in calculating and extending the color mapping function, but at least partly shows that the lighting conditions are very different between the source and target images.

5. Conclusions and Discussion

To the best of our knowledge, our work is the first work so far that performs an extensive, systematic and quantitative evaluation of the performance of color correction approaches in the context of automatic multi-view image and video stitching. Our evaluation and comparison of the approaches, has yielded a number of useful observations and conclusions.

From the perspective of color transfer effectiveness, both the non-parametric approaches of alg#3 (tensor voting in joint image space) and alg#7 (brightness transfer function) and the parametric approaches of alg#1 (gain compensation) and alg#6 (local color transfer) are superior according to the experimental data. From the perspective of extendability, parametric approaches (including alg#1 (gain compensation), alg#4 (global color transfer) and alg#6 (local color transfer)) are generally better and more stable than non-parametric ones. It is also worth mentioning that non-parametric approaches are much more complex than parametric ones. Alg#3 (tensor voting in joint image space) in particular is much slower than the other eight, according to our un-optimized implementation.

Considering all the above factors (effectiveness, extendability, stability and speed), we think alg#1 (gain compensation) and alg#6 (local color transfer) could be the first options to try for a general image and video stitching application in practice. Both of these two approaches are simple, fast, effective, and general. It is interesting to notice that to our best knowledge alg#1 may be one of the earliest color correction approaches developed for mo-

saic and panorama stitching [27], while the direct predecessor of alg#6, alg#4 (global color transfer), is widely used as the baseline approach in color transfer research. After alg#1 and alg#6, alg#3 (tensor voting in joint image space) and alg#7 (brightness transfer function) may also be good choices.

Based on our experience on studying various color correction approaches and implementing and evaluating nine of them, we think that future work on color correction approaches faces the following problems: 1) How to process extreme inputs like over-saturated parts of the input images that may affect the calculation of the color mapping function, 2) the handling of out-of-gamut errors, and 3) how to intelligently extend the color mapping function calculated from the overlapped area to non-overlapped regions. On the last problem, making use of image segmentation results to selectively extend the mapping might be a good exploration direction.

For the convenience of the reader to reproduce the experimental results shown in this paper, we make our implementation of the color correction approaches involved in our evaluation available for download on the project's website: http://ia.cs.colorado.edu/~wxu/color_correction.htm.

6. Acknowledgement

The authors would like to thank Dr. Matthew Brown for proving some of the test images. This work was funded by the NIDRR Rehabilitation Engineering Research Center on Recreational Technologies and Exercise Physiology Benefiting Persons with Disabilities (RERC RecTech), grant # H133E020715, by the National Science Foundation under Grant No. 0535269 and through a gift from the Coleman Institute for Cognitive Disabilities (<http://www.colemaninstitute.org/>).

References

- [1] Open tensor voting framework (OpenTVF). <http://sourceforge.net/projects/opentvf/>.
- [2] K. An, J. Sun, and L. Zhou. A linear color correction method for compressed images and videos. *IEICE Transactions on Information and Systems*, (10):2686–2689, 2006.
- [3] M. Brown and D. G. Lowe. Recognising panoramas. In *Proc. ICCV'03*, volume 2, pages 1218–1225, 2003.
- [4] M. Brown and D. G. Lowe. Automatic panoramic image stitching using invariant features. *IJCV*, 74(1):59–73, 2007.
- [5] M. A. Brown. *Multi-Image Matching using Invariant Features*. PhD thesis, University of British Columbia, 2005.
- [6] P. J. Burt and E. H. Adelson. A multiresolution spline with application to image mosaics. *ACM Transactions on Graphics*, 2:217–236, 1983.
- [7] Y. F. (ed.). *Fractal Image Compression: Theory and Application*. Springer Verlag, New York, 1995.
- [8] A. Eden, M. Uyttendaele, and R. Szeliski. Seamless image stitching of scenes with large motions and exposure differences. In *Proc. CVPR'06*, volume 2, pages 2498–2505, 2006.
- [9] U. Fecker, M. Barkowsky, and A. Kaup. Histogram-based prefiltering for luminance and chrominance compensation of multiview video. *IEEE Transactions on Circuits and Systems for Video Technology*, 18(9):1258–1267, 2008.
- [10] B. V. Funt and B. C. Lewis. Diagonal versus affine transformations for color correction. *Journal of the Optical Society of America A*, 17(11):2108–2112, 2000.
- [11] D. Goldman and J. Chen. Vignette and exposure calibration and compensation. In *Proc. ICCV'05*, pages 899–906, 2005.
- [12] D. Hasler and S. Süsstrunk. Mapping colour in image stitching applications. *Journal of Visual Communication and Image Representation*, 15:65–90, 2004.
- [13] A. Ilie and G. Welch. Ensuring color consistency across multiple cameras. In *Proc. ICCV'05*, pages 1268–1275, 2005.
- [14] J. Jia, J. Sun, C.-K. Tang, and H.-Y. Shum. Bayesian correction of image intensity with spatial consideration. In *Proc. ECCV'04*, volume 3, pages 342–354, 2004.
- [15] J. Jia and C.-K. Tang. Image registration with global and local luminance alignment. In *Proc. ICCV'03*, volume 1, pages 156–163, 2003.
- [16] J. Jia and C.-K. Tang. Tensor voting for image correction by global and local intensity alignment. *IEEE TPAMI*, 27(1):36–50, 2005.
- [17] S. J. Kim and M. Pollefeys. Robust radiometric calibration and vignetting correction. *IEEE TPAMI*, 30(4):562–576, 2008.
- [18] H.-C. Lee. *Introduction to Color Imaging Science*. Cambridge University Press, 2005.
- [19] A. Litvinov and Y. Schechner. Addressing radiometric non-idealities: A unified framework. In *Proc. CVPR'05*, volume 2, pages 52–59, 2005.
- [20] D. Martin, C. Fowlkes, D. Tal, and J. Malik. A database of human segmented natural images and its application to evaluating segmentation algorithms and measuring ecological statistics. In *Proc. ICCV'01*, volume 2, pages 416–423, 2001.
- [21] H. Nanda and R. Cutler. Practical calibrations for a real-time digital omnidirectional camera. In *Technical Sketches, CVPR'01*, 2001.
- [22] O. D.-C. no, M. Rivera, and P. P. Mayorga. Computing the α channel with probabilistic segmentation for image colorization. In *Proc. ICCV'07*, pages 1–7, 2007.
- [23] F. Pitić, A. C. Kokaram, and R. Dahyot. N-dimensional probability density function transfer and its application to color transfer. In *Proc. ICCV'05*, volume 2, pages 1434–1439, 2005.
- [24] F. Pitić, A. C. Kokaram, and R. Dahyot. Automated colour grading using colour distribution transfer. *CVIU*, 107(1):123–137, 2007.
- [25] E. Reinhard, M. Adhikhmin, B. Gooch, and P. Shirley. Color transfer between images. *IEEE Computer Graphics and Applications*, 21(5):34–41, 2001.
- [26] I. Skiljan. Irfanview (ver. 4.23). <http://www.irfanview.com>.
- [27] R. Szeliski. Image alignment and stitching: A tutorial. Technical Report MSR-TR-2004-92, One Microsoft Way, Redmond, WA, 2004.
- [28] Y.-W. Tai, J. Jia, and C.-K. Tang. Local color transfer via probabilistic segmentation by expectation-maximization. In *Proc. CVPR'05*, volume 1, pages 747–754, 2005.
- [29] G. Y. Tian, D. Gledhill, D. Taylor, and D. Clarke. Colour correction for panoramic imaging. In *Proc. 6th International Conference on Information Visualisation*, pages 483–488, 2002.
- [30] M. Uyttendaele, A. Eden, and R. Szeliski. Eliminating ghosting and exposure artifacts in image mosaics. In *Proc. CVPR'01*, volume 2, pages 509–516, 2001.
- [31] Z. Wang, A. C. Bovik, H. R. Sheikh, and E. P. Simoncelli. Image quality assessment: From error visibility to structural similarity. *IEEE Transactions on Image Processing*, 13(4):600–612, 2004.
- [32] Y. Xiang, B. Zou, and H. Li. Selective color transfer with multi-source images. *Pattern Recognition Letters*, 30(7):682–689, 2009.
- [33] X. Xiao and L. Ma. Color transfer in correlated color space. In *Proc. 2006 ACM international conference on Virtual reality continuum and its applications*, pages 305–309, 2006.
- [34] K. Yamamoto and R. Oi. Color correction for multi-view video using energy minimization of view networks. *International Journal of Automation and Computing*, 5(3):234–245, 2008.
- [35] K. Yamamoto, T. Yendo, T. Fujii, M. Tanimoto, and D. Suter. Color correction for multi-camera system by using correspondences. In *Research posters, SIGGRAPH'06*, 2006.
- [36] M. Zhang and N. D. Georganas. Fast color correction using principal regions mapping in different color spaces. *Real-Time Imaging*, 10(1):23–30, 2004.
- [37] Y. Zheng, S. Lin, and S. Kang. Single-image vignetting correction. In *Proc. CVPR'06*, volume 1, pages 461–468, 2006.
- [38] Y. Zheng, J. Yu, S. Kang, S. Lin, and C. Kambhampettu. Single-image vignetting correction using radial gradient symmetry. In *Proc. CVPR'08*, volume 2, pages 1–8, 2008.