Colour Correction for Stereoscopic Omnidirectional Images

[1]Simone Croci, [1]Mairead Grogan, [1,2]Sebastian Knorr, [1]Aljosa Smolic

[1]V-SENSE Project, School of Computer Science and Statistics, Trinity College Dublin;
[2]Communication Systems Group, Technical University of Berlin

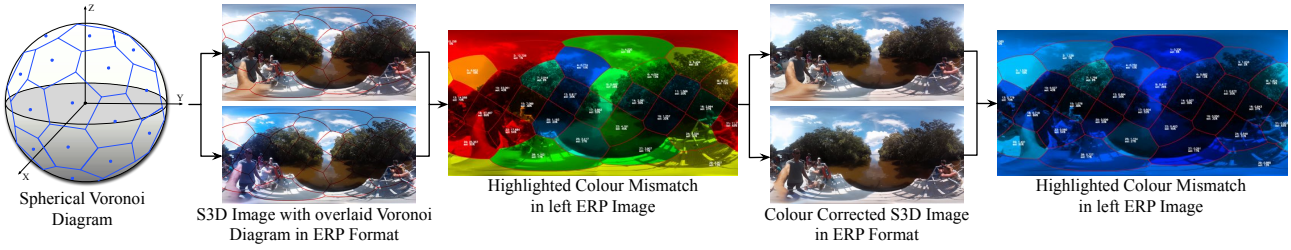

Figure 1: Colour mismatch correction and detection based on the spherical Voronoi diagram.

Abstract

Stereoscopic omnidirectional images (ODI) when viewed with a head-mounted display are a way to generate an immersive experience. Unfortunately, their creation is not an easy process, and different problems can be present in the ODI that can reduce the quality of experience. A common problem is colour mismatch, which occurs when the colours of the objects in the scene are different between the two stereoscopic views. In this paper we propose a novel method for the correction of colour mismatch based on the subdivision of ODIs into patches, where local colour correction transformations are fitted and then globally combined. The results presented in the paper show that the proposed method is able to reduce the colour mismatch in stereoscopic ODIs.

Keywords: Virtual reality, 360-video, colour matching, binocular rivalry, stereoscopic 3D

1 Introduction

One of the most popular formats to deliver virtual reality experiences is 360-video which is often referred to as VR-video or cinematic VR. Shooting 360-video is a technological challenge as there are many technical limitations which need to be overcome, especially for capturing and post-processing in stereoscopic 3D (S3D). 360-video is often captured with an omnidirectional multi-camera rig and stitched together in post-production [Zhang and Liu, 2015]. In general, the limitations inherent in 360 videos result in artifacts which cause visual discomfort when watching the content with a head-mounted display. The artifacts or issues can be divided into three categories: binocular rivalry issues (e.g. colour mismatch), conflicts of depth cues (e.g. vergence-accomodation conflicts), and artifacts which occur in both monocular and stereoscopic 360-degree content production (e.g. stitching artifacts) [Knorr et al., 2017].

In this paper, we focus on binocular rivalry issues, in particular on colour mismatch detection and correction in S3D omnidirectional content as illustrated in Figure 1. Colour mismatch in multi-camera systems is an inherent problem due to different camera and lens characteristics, different illumination and reflections resulting from different camera orientations, etc. Such colour mismatches occur during the stitching and blending process

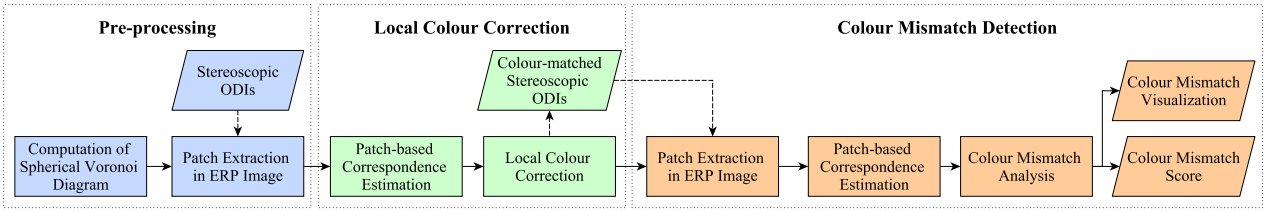

Figure 2: Overall system for colour mismatch correction and detection.

of multiple views into a single monocular panorama, but also between the left and right view of a stereoscopic panorama, which often result in visual discomfort [Knorr et al., 2012].

In this context, we introduce a novel approach and entire system for colour mismatch correction and detection in S3D omnidirectional content. The system consists of three main modules: pre-processing, local colour correction between the left and right stereoscopic ODIs, and colour mismatch detection as shown in Figure 2. During the pre-processing step, we compute a spherical Voronoi diagram and extract Voronoi patches in the equirectangular projection (ERP) format of the left and right ODI. Then, we estimate pixel correspondences between the corresponding patches in both views and apply a local colour transfer in order to match the colours of the corresponding patches, which is the main contribution of this paper. Finally, our patch-based colour mismatch detection module [Croci et al., 2017] measures and visualises colour mismatch, which might still be present, between the left and right stereoscopic ODIs. While the entire colour correction approach is applied in the RGB colour space and uses a colour transfer approach introduced in [Grogan and Dahyot, 2017], the colour mismatch detection module is applied in the Lab colour space and uses colour statistics analysis proposed in [Reinhard et al., 2001]. This allows a more objective and independent evaluation of still existing colour discrepancies between the views. Finally, we evaluate the entire system on 15 S3D ODIs with large colour mismatch and show that our system improves the quality of the ODIs significantly.

The remainder of the paper is organised as follows. In Section 2, related work in artifact detection and colour correction is reviewed. Then, in Section 3, we describe the proposed system for the subdivision of the ODI into patches, the colour correction step, and the colour mismatch detection. In Section 4, the proposed colour correction method is evaluated with 15 S3D ODIs. Finally, in Section 5, the paper concludes with a discussion and future work.

2 Related Work

Colour mismatch detection and colour correction in stereoscopic and multi-view applications has been an ongoing research topic for many years. In [Dong et al., 2013], a method for detecting stereo camera distortions based on statistical models was presented in order to evaluate vertical misalignment, camera rotation, unsynchronised zooming, and colour mismatch in S3D content.

A large variety of artifact detection methods, including a method for the detection of colour mismatch, was introduced in [Voronov et al., 2013]. More recently, in [Knorr et al., 2017, Croci et al., 2017], S3D quality assessment methods for stereoscopic ODIs were introduced which also focus on colour mismatch detection.

In the computer vision and multi-view video processing communities, the initial efforts on solving colour mismatches between multiple views used exposure compensation (or gain compensation) [Xu and Mulligan, 2010]. This approach adjusts the gain level of images to compensate for appearance differences caused by different exposure levels. However, this approach may fail in the case of local differences e.g. caused by lens flares or polarization.

[Wang et al., 2011] proposed a robust algorithm to correct the colour discrepancy between images, which neither requires a colour calibration chart/object, nor explicitly compensates for the image as a whole. Instead, they correct the image region by region using local feature correspondences. In [Zheng et al., 2017], a method is proposed that combines global and local colour information to correct colour discrepancies between stereo-

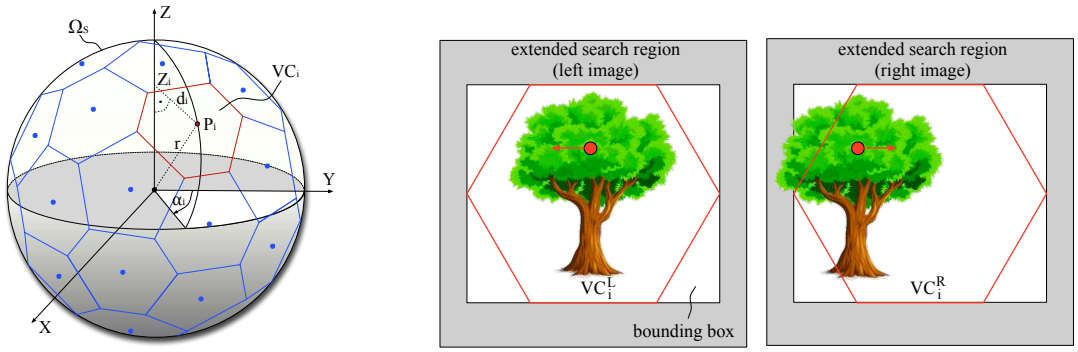

(a) Spherical Voronoi diagram.

(b) Disparity compensation.

Figure 3: Spherical Voronoi diagram and disparity compensation.

scopic image pairs. The algorithm uses dense stereo matching and global colour correction to initialise colour values, and then improves the local colour smoothness and global colour consistency of the resulting image.

For large baseline multi-view video, [Ye et al., 2017] introduced a robust colour correction method that enforces spatio-temporal colour consistencies and gradient preservation by solving a global optimization problem. The authors of [Xia et al., 2017] proposed an effective colour correction method for multi-view image stitching which first finds coherent content regions in inter-image overlaps, and then parameterise a colour remapping curve as transform model.

The image processing and computer graphics communities were developing similar colour manipulation methods, called *colour transfer* techniques. These methods transfer the colour feel from a palette image to a target image, and assume that the content of the images is different. The earliest work in this area was by [Reinhard et al., 2001], who proposed transforming the mean and standard deviation of each colour channel in the target image to match that of the palette image. Since then, more complex techniques have been used to model the colour distributions of the images more accurately, including histograms and Gaussian Mixture models [Pitie et al., 2005, Tai et al., 2005]. While global colour transfer functions are often used, including affine, radial basis and optimal transport functions [Pitié and Kokaram, 2007, Grogan et al., 2017, Bonneel et al., 2016], local techniques have also been proposed to allow for more flexibility in the recolouring [Wang et al., 2010, Shih et al., 2013]. Recently, Grogan and Dahyot [Grogan and Dahyot, 2017, Grogan et al., 2015] proposed a colour transfer technique that could also be enhanced to take into account colour correspondences between the target and palette images, ensuring the method could be used to colour correct images of the same scene. They showed that this method performed as well as other state of the art colour corrections techniques, with the advantage of being more robust to correspondence outliers. In this paper, we extend this method so that it can be used to successfully colour correct stereoscopic ODIs.

3 Proposed Method

The complete system for the correction and detection of colour mismatch is illustrated in Figure 2. Its three main components, that is, the pre-processing step, the local colour correction, and the colour mismatch detection are described in the next sections.

3.1 Pre-processing Step

The first step is the subdivision of the ODI into approximately equally sized patches based on the approach described in [Croci et al., 2017]. This ensures both a local colour correction of image patches and a localisation of colour mismatches in the detection module. First, a spherical Voronoi diagram [Aurenhammer, 1991] is computed from n evenly distributed points on the sphere as illustrated in Figure 3a. The evenly distributed points $\mathbf{P}_i = (X_i, Y_i, Z_i)$ with $i = 1 \ldots n$ on the sphere are defined by the following equations:

$\alpha_i = (i-1)\pi \cdot \left(3-\sqrt{5}\right)$, $Z_i = \left(1-\frac{1}{n}\right) \cdot \left(1-\frac{2(i-1)}{n-1}\right)$, $d_i = \sqrt{1-Z_i^2}$, $X_i = d_i \cdot cos(\alpha_i)$ and $Y_i = d_i \cdot sin(\alpha_i)$, where α_i is the azimuthal angle and d_i is the distance of the point from the z-axis. From the n evenly distributed points, the spherical Voronoi diagram basically partitions the surface of the sphere Ω_S into n cells VC_i, where each point in the cell VC_i is closer to the corresponding point $\mathbf{P}_i$ than to any of the other evenly distributed points $\mathbf{P}_j$. Formally, the cell VC_i is defined as follows:

$$VC_i = \{\mathbf{P} \in \Omega_S \mid d_S(\mathbf{P}, \mathbf{P}_i) \leq d_S(\mathbf{P}, \mathbf{P}_j) \ \forall j \neq i\}, \tag{1}$$

where $d_S(\mathbf{P}, \mathbf{P}_i)$ is the spherical distance between the current point $\mathbf{P}$ and the point $\mathbf{P}_i$, i.e., the length of the shortest path on the surface of the sphere connecting these two points.

Next, we map each cell VC_i of the spherical Voronoi diagram to a planar image patch Π_i. For each cell, a planar patch is positioned on the centroid of the cell, tangent to the sphere. The points on the sphere and the planar patch are related by central projection, and the pixel values of the patch are computed by sampling the ODI in ERP format using bilinear interpolation. The resolution of each patch is defined by the pixels per visual angle, a parameter that is kept constant for each patch. In the presence of disparity, it can occur that a region inside a Voronoi cell in one view is outside the same cell in the other view. In order to cope with the disparity, we add a border around the Voronoi cell when the patch is extracted, as shown in Figure 3b.

The number of patches and thus the size of each patch influences the reduction of the colour mismatch. If the colour mismatch is localised in a small region and the patch is large, then the proposed method could have difficulty in matching the colours between the two views. We have empirically found that 30 patches is a good number for most of the ODIs that we have processed.

3.2 Local Colour Correction

The local colour correction component of the system involves first estimating colour correspondences $\{c_L^{(k)}, c_R^{(k)}\}_{k=1..m}$ between corresponding patches of the left and right view. We investigated two methods for the estimation of correspondences: the Semi-Global Block Matching approach [Hirschmuller, 2008] and the Coarse-to-Fine PatchMatch approach [Hu et al., 2016], but we found no significant difference between the colour correction results generated using these approaches.

For each patch, we use the correspondences to estimate a colour transformation which recolours the patch of the right view so that it is more similar to the left, using the method proposed in [Grogan and Dahyot, 2017]. For each patch, we fit two Gaussian Mixture models GMM_L and GMM_R to the left and right colour correspondences respectively:

$$GMM_Y(x) = \frac{1}{\mathrm{m}} \sum_{k=1}^{\mathrm{m}} \mathcal{N}(x; c_Y^{(k)}, h\mathrm{I}), \quad \text{with } Y \in \{L, R\}, \tag{2}$$

where $x \in \mathbb{R}^3$ are colour values of the RGB colour space, and each Gaussian $\mathcal{N}$ is associated with an identical isotropic covariance matrix, hI. The goal is to align the two Gaussian Mixture models by warping the right one as follows:

$$GMM'_R(x|\theta) = \frac{1}{\mathrm{m}} \sum_{k=1}^{\mathrm{m}} \mathcal{N}(x; \phi(c_R^{(k)}, \theta), h\mathrm{I}), \tag{3}$$

where ϕ represents a parametric Thin Plate Spline (TPS) transformation controlled by the parameter θ. Technically, the alignment beteween GMM_L and the warped GMM'_R is obtained by minimising the $\mathscr{L}_2$ distance between them. This $\mathscr{L}_2$ technique has been shown to be robust to correspondence outliers, and the smooth TPS function ensures that similar colours in the patch remain similar after recolouring, eliminating artifacts in the gradient of the image which can appear when using other recolouring methods [Pitie et al., 2005].

Once the transformations ϕ_i have been estimated for each patch Π_i, they have to be combined to recolour the entire ODI of the right view. To ensure that there are no harsh colour changes between patches in the recoloured ODI, we use weight masks to blend the transformations. For each transformation ϕ_i, a corresponding weight mask G_i is computed. The masks G_i are in ERP format and their pixel values are used to weight the

contributions of each of the transformations when recolouring the ODI. To compute the value of a pixel in the weight mask G_i, the spherical distance between this pixel in the spherical ODI and the centroid of the Voronoi cell VC_i is computed, and a Gaussian function is applied to it. In this way, in G_i, pixels that lie close to the centroid in the ODI will have higher weights than those further away. Then, when recolouring the ODI of the right view I_R in ERP format to its corrected version $\hat{I}_R$, the colour of the pixel at location (j,k) is given by:

$$\hat{I}_R(j,k) = \frac{\sum_{i=1}^{n} G_i(j,k) \cdot \phi_i(I_R(j,k),\theta)}{\sum_{i=1}^{n} G_i(j,k)}. \tag{4}$$

In this manner, each local colour transformation has the most influence in the area from which it is estimated, and the colour transformations are smoothly blended without creating any artifacts at the patch borders.

3.3 Colour Mismatch Detection

The colour mismatch detection applied after the colour correction of the ODI, and useful for getting feedback on the remaining colour mismatch, is a simplification of the colour mismatch detection proposed in [Croci et al., 2017]. The simplification is obtained by discarding the saliency, since it is not available for all the 15 ODIs processed in this paper. As mentioned in Section 3.1, the detection module also uses Voronoi patches for colour mismatch detection and localization. First, the ODI is partitioned into patches. Then, correspondences are computed for corresponding patches between the two views using the Semi-Global Block Matching approach introduced in [Hirschmuller, 2008]. The correspondences are necessary in order to identify the regions that are present in both views of the patch. From the common regions, the colour means μ_L and μ_R, and the colour standard deviations σ_L and σ_R are estimated in the Lab colour space as proposed in [Reinhard et al., 2001]. Finally, for each patch Π_i the following colour mismatch score is computed:

$$CMS_i = \sqrt{\left\|\mu_L - \mu_R\right\|^2 + \lambda \left\|\sigma_L - \sigma_R\right\|^2}, \tag{5}$$

where λ is a tuning parameter that was set to one for the analysis of the ODIs. The patch scores can be visualised with the jet colourmap and overlaid with the ODI in ERP format as shown in the teaser in Figure 1. In order to get the global score CMS_{global} for the entire ODI, the patch scores are simply averaged:

$$CMS_{global} = \frac{\sum_{i=1}^{n} CMS_i}{n}. \tag{6}$$

4 Results

In order to evaluate the proposed method, we selected 14 ODIs with the highest colour mismatch scores from the dataset introduced in [Croci et al., 2018], and one ODI which was captured with a 360° mirror-rig. Figure 4a shows a bar chart with the global colour mismatch scores computed with Equation 6 before and after applying the proposed colour correction method. In this figure, one can see that the novel approach is able to reduce the colour mismatch in all ODIs by an average of 74%. The largest score reduction, equal to 89%, was observed for ODI 1.

Figure 4b shows for four ODIs the individual colour mismatch patch scores and their visualisation using the jet colourmap before and after colour correction, together with some close-ups. ODI 1 is a clear example of strong colour mismatch where our method is able to significantly reduce it, as can be clearly seen in the close-up. ODI 2 and ODI 3 also show a strong colour mismatch localised to a particular region. Even in these two cases, our method fixes the mismatch. For ODI 4, the close-up shows how even minor colour mismatches can be corrected.

Apart from the good results, we also observed some limitations of the proposed method. Since the method is based on pixel correspondences between corresponding patches in both views, colour mismatch correction and detection are dependent on the pixel correspondence accuracy. Another limitation occurs when the patch is too large compared to the region with colour mismatch, or when the patch contains regions with different types of colour mismatch. In this case the colour mismatch is reduced only partially.

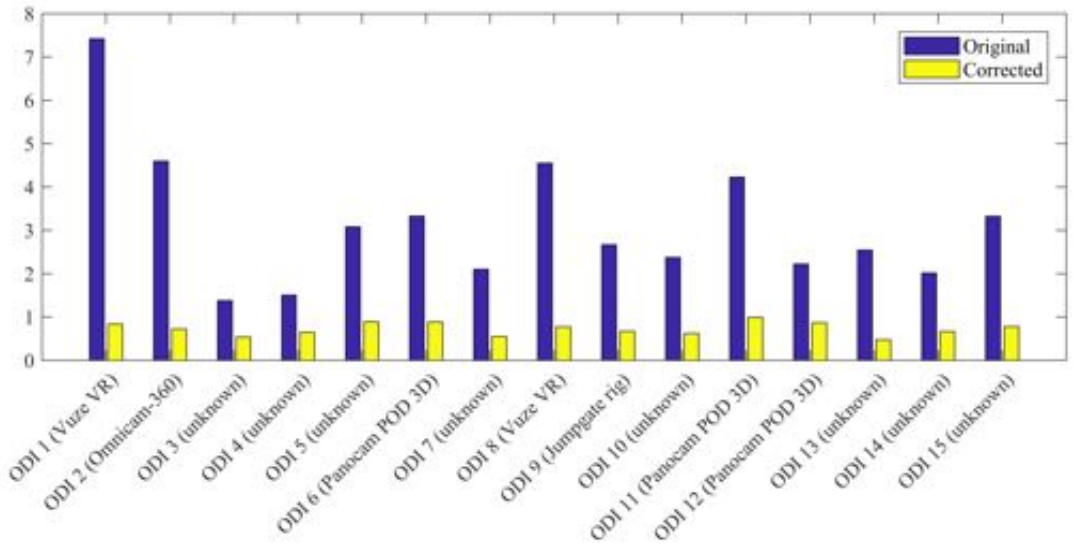

(a) Global colour mismatch scores before and after colour correction (used camera rigs are specified in brackets).

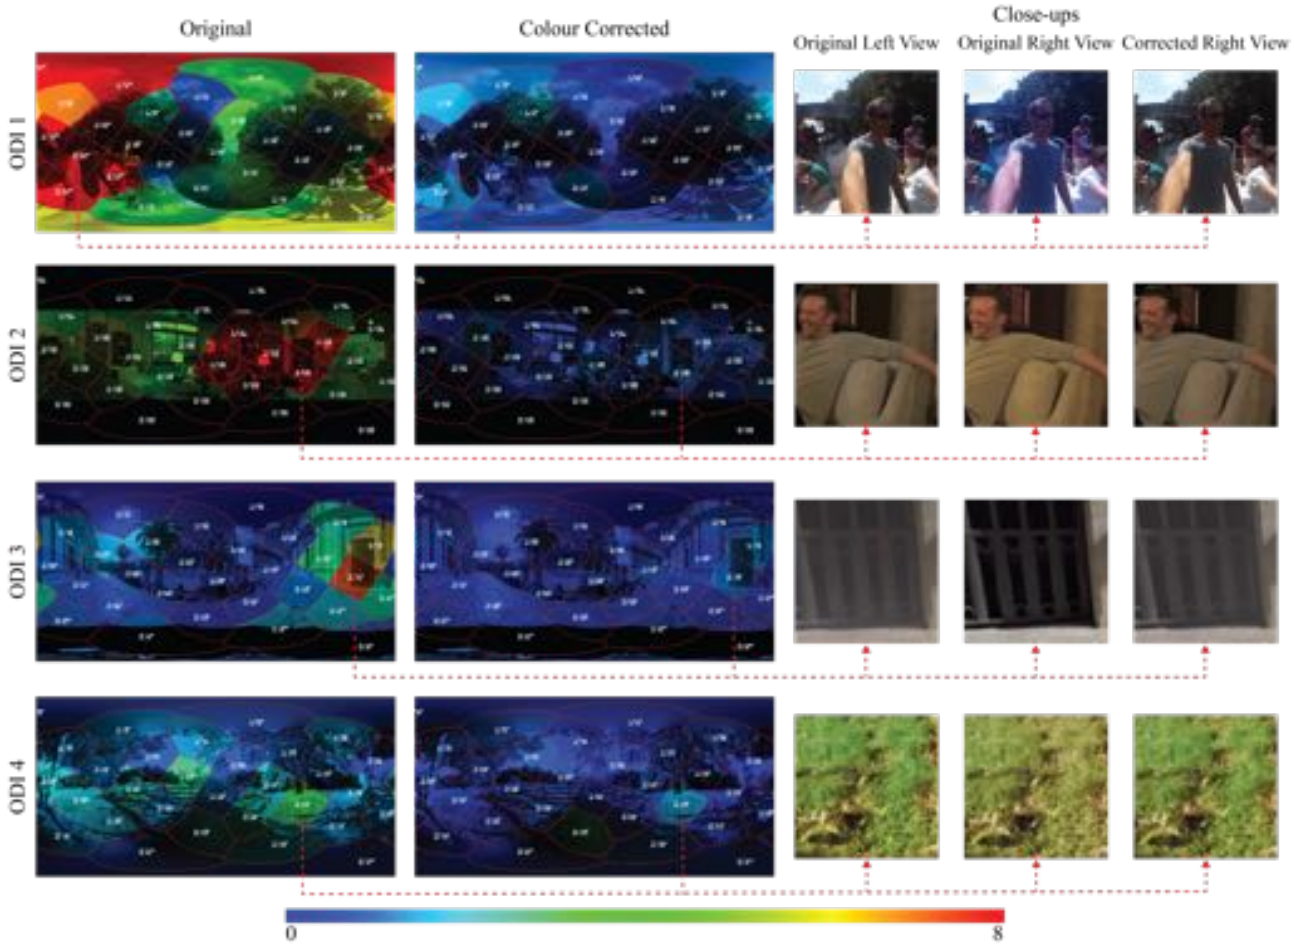

(b) Sample ODIs with colour mismatch visualisation (red: strong mismatch, blue: no mismatch) and close-ups before and after colour correction.

Figure 4: Results.

5 Conclusions

This paper presented a solution to the problem of colour mismatch in stereoscopic ODIs, which can cause visual discomfort. The proposed approach first divides the ODI into patches using the spherical Voronoi diagram from evenly distributed points on the sphere. In each patch a colour transformation is fitted from correspondences in the RGB colour space, and then the colour transformations are combined together using weight masks based on the spherical distance from the centroid of the Voronoi cells. In order to analyse colour mismatch objectively and independently, the processed ODI is analysed in the Lab colour space using an alternative correspondence estimation approach.

The results show that the proposed approach is able to reduce the colour mismatch significantly. This conclusion was obtained by computing the global and local patch colour mismatch scores before and after applying the colour correction approach on different ODIs. In particular, 89% is the largest global score reduction that was observed. However, the proposed method is not exempt from limitations related to the accuracy of the correspondence estimation, and the misalignment of the patches with the region characterized by colour mismatch.

In the future, we plan to improve the proposed approach by tackling some of the limitations, especially the misalignment of the patch with the region affected by colour mismatch. The plan is to investigate the possibility to have adaptable patches to the region with colour mismatch.

Acknowledgments

This publication has emanated from research conducted with the financial support of Science Foundation Ireland (SFI) under the Grant Number 15/RP/2776.

References

[Aurenhammer, 1991] Aurenhammer, F. (1991). Voronoi Diagrams - A Survey of a Fundamental Data Structure. *ACM Comput. Surv.*, 23(3):345–405.

[Bonneel et al., 2016] Bonneel, N., Peyré, G., and Cuturi, M. (2016). Wasserstein barycentric coordinates: Histogram regression using optimal transport. *ACM Trans. Graph.*, 35(4):71:1–71:10.

[Croci et al., 2017] Croci, S., Knorr, S., Goldmann, L., and Smolic, A. (2017). A framework for quality control in cinematic VR based on Voronoi patches and saliency. In *IEEE Int. Conf. 3D Immersion (IC3D)*, Brussels, Belgium.

[Croci et al., 2018] Croci, S., Knorr, S., and Smolic, A. (2018). Sharpness mismatch detection in stereoscopic content with 360-degree capability. *IEEE Int. Conf. Image Processing (ICIP)*.

[Dong et al., 2013] Dong, Q., Zhou, T., Guo, Z., and Xiao, J. (2013). A stereo camera distortion detecting method for 3DTV video quality assessment. In *IEEE Asia-Pacific Signal and Inform. Process. Assoc. Annu. Summit and Conf.*, pages 1–4.

[Grogan and Dahyot, 2017] Grogan, M. and Dahyot, R. (2017). Robust registration of Gaussian mixtures for colour transfer. *CoRR*, abs/1705.06091.

[Grogan et al., 2017] Grogan, M., Dahyot, R., and Smolic, A. (2017). User interaction for image recolouring using L2. In *Proc. of the 14th European Conf. on Visual Media Prod.*, CVMP 2017, pages 6:1–6:10, New York, NY, USA. ACM.

[Grogan et al., 2015] Grogan, M., Prasad, M., and Dahyot, R. (2015). L2 registration for colour transfer. In *23rd European Signal Process. Conf. (EUSIPCO)*, pages 1–5.

[Hirschmuller, 2008] Hirschmuller, H. (2008). Stereo processing by semiglobal matching and mutual information. *IEEE Trans. Pattern Anal. Mach. Intell.*, 30(2):328–341.

[Hu et al., 2016] Hu, Y., Song, R., and Li, Y. (2016). Efficient coarse-to-fine patch match for large displacement optical flow. In *IEEE Conf. on Comput. Vision and Pattern Recognition (CVPR)*, pages 5704–5712.

[Knorr et al., 2017] Knorr, S., Croci, S., and Smolic, A. (2017). A modular scheme for artifact detection in stereoscopic omni-directional images. In *Proc. of the Irish Mach. Vision and Image Process. Conf. (IMVIP)*.

[Knorr et al., 2012] Knorr, S., Ide, K., Kunter, M., and Sikora, T. (2012). The avoidance of visual discomfort and basic rules for producing "good 3D" pictures. *SMPTE Motion Imaging J.*, 121(7):72–79.

[Pitié and Kokaram, 2007] Pitié, F. and Kokaram, A. (2007). The linear monge-kantorovitch linear colour mapping for example-based colour transfer. In *Visual Media Prod., 2007. IETCVMP. 4th European Conf. on*, pages 1–9.

[Pitie et al., 2005] Pitie, F., Kokaram, A., and Dahyot, R. (2005). N-dimensional probability density function transfer and its application to color transfer. In *IEEE Int. Conf. Comput. Vision (ICCV)*, volume 2, pages 1434–1439 Vol. 2.

[Reinhard et al., 2001] Reinhard, E., Ashikhmin, M., Gooch, B., and Shirley, P. (2001). Color transfer between images. *IEEE Comput. Graph. Appl.*, 21(5):34–41.

[Shih et al., 2013] Shih, Y., Paris, S., Durand, F., and Freeman, W. T. (2013). Data-driven hallucination of different times of day from a single outdoor photo. *ACM Trans. Graph.*, 32(6):200:1–200:11.

[Tai et al., 2005] Tai, Y.-W., Jia, J., and Tang, C.-K. (2005). Local color transfer via probabilistic segmentation by expectation-maximization. In *IEEE Conf. on Comput. Vision and Pattern Recognition (CVPR)*, volume 1, pages 747–754 vol. 1.

[Voronov et al., 2013] Voronov, A., Vatolin, D., Sumin, D., Napadovsky, V., and Borisov, A. (2013). Methodology for stereoscopic motion-picture quality assessment. In *Proc. of the SPIE, Stereoscopic Displays and Appl. XXIV*, volume 8648.

[Wang et al., 2010] Wang, B., Yu, Y., Wong, T.-T., Chen, C., and Xu, Y.-Q. (2010). Data-driven image color theme enhancement. *ACM Trans. Graph.*, 29(6):146:1–146:10.

[Wang et al., 2011] Wang, Q., Yan, P., Yuan, Y., and Li, X. (2011). Robust color correction in stereo vision. In *Int. Conf. Image Processing (ICIP)*, number 2, pages 965–968.

[Xia et al., 2017] Xia, M., Yao, J., Xie, R., and Zhang, M. (2017). Color Consistency Correction Based on Remapping Optimization for Image Stitching. In *IEEE Int. Conf. Comput. Vision (ICCV)*.

[Xu and Mulligan, 2010] Xu, W. and Mulligan, J. (2010). Performance evaluation of color correction approaches for automatic multi-view image and video stitching. In *IEEE Conf. on Comput. Vision and Pattern Recognition (CVPR)*, pages 263–270.

[Ye et al., 2017] Ye, S., Lu, S. P., and Munteanu, A. (2017). Color correction for large-baseline multiview video. *Elsevier Signal Process.: Image Commun.*, 53(January):40–50.

[Zhang and Liu, 2015] Zhang, F. and Liu, F. (2015). Casual stereoscopic panorama stitching. In *IEEE Conf. on Comput. Vision and Pattern Recognition (CVPR)*, pages 2002–2010.

[Zheng et al., 2017] Zheng, X., Yuzhen, N., Chen, J., and Chen, Y. (2017). Color correction for stereoscopic image based on matching and optimization. In *IEEE Int. Conf. 3D Immersion (IC3D)*, Brussels.