Colour Spaces for Colour Transfer

Erik Reinhard and Tania Pouli

University of Bristol, Dept. of Computer Science
{reinhard,pouli}@cs.bris.ac.uk

Abstract. Colour transfer algorithms aim to apply a colour palette, mood or style from one image to another, operating either in a three-dimensional colour space, or splitting the problem into three simpler one-dimensional problems. The latter class of algorithms simply treats each of the three dimensions independently, whether justified or not. Although they rarely introduce spatial artefacts, the quality of the results depends on how the problem was split into three sub-problems, i.e. which colour space was chosen. Generally, the assumption is made that a decorrelated colour space would perform best, as decorrelation makes the three colour channels semi-independent (decorrelation is a weaker property than independence). However, such spaces are only decorrelated for well-chosen image ensembles. For individual images, this property may not hold. In this work, the connection between the natural statistics of colour images and the ability of existing colour transfer algorithms to produce plausible results is investigated. This work aims to provide a better understanding of the performance of different colour spaces in the context of colour transfer.

Keywords: Colour Transfer, Colour Spaces, Correlation.

1 Introduction

Colour is one of the main image attributes used in art, photography and visualisation for relaying information, or for conveying a specific mood. Tools available to alter and manipulate the colour content of images tend to require extensive user input or simply do not offer sufficient control over the result. This leaves many such tools out of reach of untrained users. Considering such users, colour transfer forms a class of techniques that allow the colour palette of an image to be altered using a second image as a reference. The task is then to select a reference image whose colours are preferred. Subsequently the algorithm will modify the original image such that it acquires the palette of that reference.

In essence this operation can be seen as a function that, given two images, produces a third that has maintained the semantic content of the one while acquiring the colours of the second. This of course can be achieved in many different ways, using different descriptors for the image content, or operating in different colour spaces, leading to an extensive collection of colour transfer techniques.

This set of techniques can be broadly categorised into two classes, namely those that treat colour transfer as a 3D problem, and those that decompose the problem into a set of three 1D problems. The latter class is the focus of this paper, and relies on an important insight from the area of natural image statistics.

Natural image statistics is the study of statistical regularities in images, and usually aims to help understand how the human visual system operates. The argument is that human vision has evolved in the presence of a natural environment, and is therefore in some sense specialised to observe and interpret natural images. Such images are normally taken to be those that occur in nature, and for instance do not contain man-made structures.

In the context of human vision, Ruderman et al tested this hypothesis. They converted a set of spectral natural images to LMS cones space, a colour space that resembles the responses of the L, M and S cones in the human retina (the letters stand for Long, Medium and Short wavelength, respectively, indicating the peak sensitivity of each cone type). This image ensemble was then subjected to Principal Component Analysis (PCA), showing that the three principal components obtained closely correspond to the opponent colour channels found in the human visual system, in particular in the retina [23]. As PCA yields components that are maximally decorrelated, we can infer that the human retina aims to decorrelate its input. While formally PCA only decorrelates, it was found that transforming natural images according to this colour space, termed $L\alpha\beta$, yields channels that are close to independent.

By producing a decorrelated image signal, and for natural images largely independent signal, the efficiency of transmitting information through the optic nerve is increased, especially considering that the optic nerve constitutes a bottleneck in terms of transmission bandwidth. On the other hand, it is now possible to use this knowledge for computational tasks. The observation is that if the three colour channels of an image can be made independent, then image processing can take place in each of the three channels independently [19].

This has led to an early colour transfer algorithm which essentially matches means and standard deviations between a pair of images [19]. It can be shown that the choice of colour space is important, i.e. that decorrelation matters. An example is given in Figure 1, where the RGB colour space does not lead to an effective transfer of the colour palette, while the $L\alpha\beta$ space produces a much more plausible result.

Although $L\alpha\beta$ has been adopted as the space of choice for several colour transfer algorithms [19,5,25,34,11,33,26,12,35,32,30], related spaces such as CIELAB are also used [4,28,18]. Other colour spaces implicated in colour transfer are CIE LCh* [14], YC_bC_r [9], Yuv [10,27], as well as colour appearance models [13]. On the other hand, several authors suggest to compute a dedicated colour space for each transfer, based on PCA [8,1,7,31,2].

We argue that the choice of colour space is important, especially for the class of algorithms that decompose the problem into a set of three 1D problems. To test this hypothesis, we present a pilot study for which we collected a small set of images, and divided them into four distinct image sets, containing natural
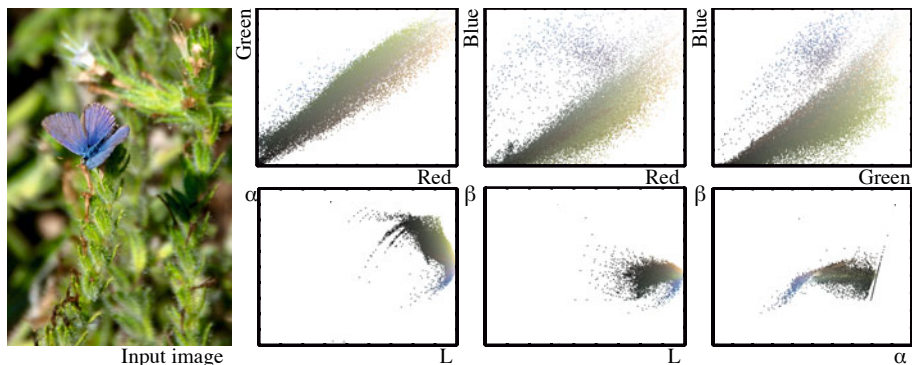
Fig. 1. Decorrelation properties of colour spaces. The top row shows pixel values for pairs of channels in an RGB colour space, taken from the image at the top left. The bottom row shows the same pixels plotted in pairs of channels from the $L\alpha\beta$ space instead. Values in channel in RGB spaces tend to be good predictors of other channels, resulting in an almost diagonal distribution.

scenes, manmade scenes, indoor scenes and night-time scenes. We first measure how well different colour spaces decorrelate each category, predicting how well we expect these spaces to perform for each image category. We then apply the colour transfer algorithm to pairs of images within each image category in each colour space to assess our predictions. We hope the insights presented in this paper provide a useful background for future developments in colour transfer.

2   Previous Work

One of the earliest colour transfer algorithms matches means and standard deviations in the $L\alpha\beta$ colour space that is decorrelated for natural images [19,20]. To achieve a more complete transfer, histogram matching can be used to transfer the distributions of images [14,32]. Some argue that a post-processing step to preserve gradient magnitudes is desirable [32]. A more refined approach sits in-between Reinhard's original technique and full histogram matching, allowing histograms to be partially matched in a controllable manner [17,18].

Recognising that no single colour space may suit all occasions, some authors propose to apply Principal Components Analysis to individual images, creating a dedicated colour space for each transfer [8,1,7,31,2]. To a greater degree this can also be achieved by means of Independent Components Analysis [6].

Most of these methods treat the three colour channels independently, meaning that the choice of colour space as well as the choice of source and target images determine the visual quality of the transfer. Alternatively, it is possible to carry out the colour transfer in 3D colour space, for instance by transferring probability distributions from one image to another [15,16], or by employing Gaussian Mixture Models [24,33,30]. A further technique aims to minimise artefacts by limiting transfers to remain within colour naming categories [4]. Some

techniques allow stroke-based user input to define target colours [10] or to define corresponding regions [28,3].

3   Method

An important aspect both in colour processing in the human visual system and in many colour transfer algorithms is the decorrelation of the input. This can only be assessed relative to a set of images. In this section we discuss the data sets that we captured, the colour spaces that we test, as well as the colour transfer algorithm employed. Sections 4 and 5 then present our results, followed by a discussion in Section 6.

3.1   Image Ensembles

To analyse the ability of existing colour spaces to decorrelate different scene categories as well as to evaluate the different colour spaces resulting from these scene categories we analyse several image databases. For this work, we focus both on different scene categories:

- **ND - Natural Day:** contains 95 day-time images shot outdoors in natural locations, without manmade structures. A variety of landscapes and weather conditions are included.
- **MD - Manmade Day:** contains 100 day-time images shot outdoors in manmade locations. These include urban landscapes as well as constructions and monuments from different eras.
- **IN - Indoors:** contains 100 images shot in indoors locations, such as rooms, offices or shops. In many scenes, artificial and natural light sources may be visible, including windows.
- **MN - Night:** contains 52 images shot at dusk or night, in urban locations. Many contain light sources visible within the scene such as street lights.

For the purpose of creating PCA-based colour spaces for each category, we use all the images in each category. When assessing the plausibility of colour transfer, we use a random subset of 6 images from each category (11 from the natural day ensemble).

3.2   Colour Spaces

A set of candidate spaces was selected based on the colour spaces seen in use in colour transfer research. The chosen colour spaces are $L\alpha\beta$, CIELAB (using both D65 and E illuminants), Yxy, Yuv, Yu'v', XYZ, RGB and HSV [22]. We have also computed a dedicated colour space for each of the image ensembles using PCA. Although most colour spaces are well known, the $L\alpha\beta$ space is perhaps least known, and is therefore described in more detail here.

To transform the images to $L\alpha\beta$ space, the input sRGB values are first converted to the device independent XYZ space:

$$\begin{bmatrix} X \\ Y \\ Z \end{bmatrix} = \begin{bmatrix} 0.4124 & 0.3576 & 0.1805 \\ 0.2126 & 0.7152 & 0.0722 \\ 0.0193 & 0.1192 & 0.9505 \end{bmatrix} \begin{bmatrix} R \\ G \\ B \end{bmatrix} \tag{1}$$

After the RGB to XYZ conversion, values are further converted to the LMS cone space to prepare them for the final transform into the $L\alpha\beta$ space:

$$\begin{bmatrix} L \\ M \\ S \end{bmatrix} = \begin{bmatrix} 0.3897 & 0.6890 & -0.0787 \\ -0.2298 & 1.1834 & 0.0464 \\ 0.0000 & 0.0000 & 1.0000 \end{bmatrix} \begin{bmatrix} X \\ Y \\ Z \end{bmatrix} \tag{2}$$

Following Ruderman et al. [23], the LMS values are logarithmically compressed to reduce skew in the data before applying the linear transform to $L\alpha\beta$:

$$\begin{bmatrix} L \\ \alpha \\ \beta \end{bmatrix} = \mathrm{diag}\left(\frac{1}{\sqrt{3}} \ \frac{1}{\sqrt{6}} \ \frac{1}{\sqrt{2}} \right) \begin{bmatrix} 1 & 1 & 1 \\ 1 & 1 & -2 \\ 1 & -1 & 0 \end{bmatrix} \begin{bmatrix} \log L \\ \log M \\ \log S \end{bmatrix} \tag{3}$$

This colour space is characterised by a luminance channel L and two opponent colour channels α and β, which encode yellow-blue and red-green chromaticities. The space is derived from Principal Components Analysis of natural images, and is therefore expected to be a good candidate for colour transfer algorithms, a hypothesis tested in this paper.

3.3   Ensemble-Based Colour Spaces

For each of our image ensembles we compute dedicated colour spaces following a process similar to the one described in [23]. More specifically, the following steps result in a $3x3$ matrix conversion matrix, converting between LMS cone space and a space which is decorrelated with respect to a specific image ensemble. Each of the rows of the matrix correspond to a principal component and therefore an axis of the ensemble-specific colour space.

A square patch of 128×128 pixels is selected from each of the images in an ensemble. These patches are then converted to LMS cone space. The patches are logarithmically compressed, and the data in each channel is mean centred, simulating a simplified Von Kries adaptation step [29]. Finally Principal Component Analysis (PCA) is performed on the N $\times$ 3 matrix which results from concatenating the scanlines of all image patches. This produces a set of 3 components (eigenvectors) sorted according to their respective eigenvalues. For each of our four image ensembles, this leads to a colour space which in the PCA sense would be more decorrelated than any of the standard colour spaces.

3.4   Colour Transfer

In its simplest form, the algorithm proposed by Reinhard et al. [19,21] shifts and scales the pixel values of the source image to match the mean and standard deviation of the target:

$$L_o = \frac{\sigma_{t,l}}{\sigma_{s,l}} (L_s - \mu_s) + \mu_t \tag{4}$$
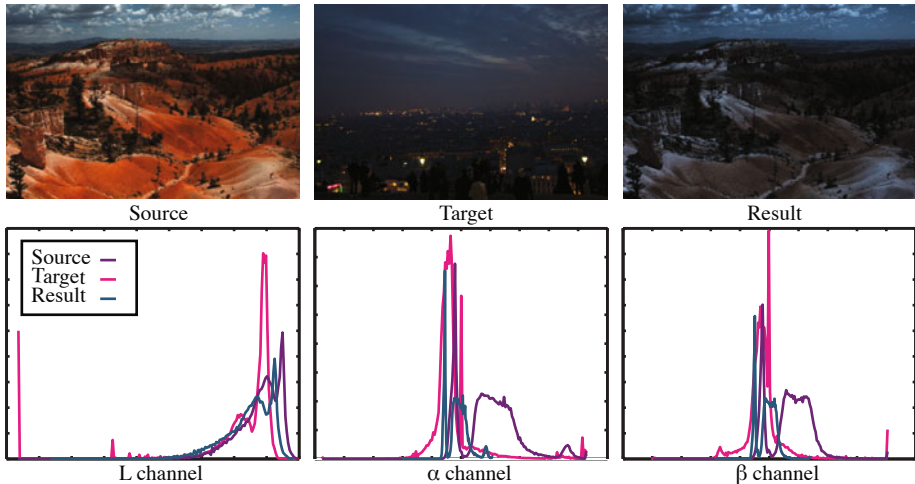
Fig. 2. An example pair of source and target images as well as their resulting output after using the colour transfer technique by Reinhard et al. [19] are shown in the top row. The corresponding histograms are shown at the bottom for the three channels of the $L\alpha\beta$ colour space.

Here, the subscripts s, t and o correspond to the source, target and output images and μ and σ are their respective means and standard deviations. Equation 4 describes the distribution transfer between source and target for the luminance channel only. The same process is repeated for the two opponent channels of the $L\alpha\beta$ space to complete the colour transfer. Figure 2 shows an example pair of source and target images as well as their corresponding histograms before and after transferring the colours between them. Rather than employ the $L\alpha\beta$ space, in the following experiments we operate this procedure in a variety of different spaces to assess their viability.

4   Correlation Analysis

To analyse the extent to which a colour space decorrelates its colour channels for a specific image ensemble, we use a procedure based on measuring the covariance between channels. First, each of our image sets is converted to the different spaces that we wish to compare. Following these conversions, the centre patch of 128 × 128 pixels is selected from each image, arranged to form an N × 3 vector. From this vector a 3-by-3 covariance matrix can be computed, containing the covariance between pairs of channels in the off-diagonal elements and the channel variance on the diagonal elements.

To make the covariance values comparable between colour spaces, we normalise the data used to compute the covariances. This is not straightforward as each colour space has different ranges and possibly different ratios between its

Table 1. The volumes of different colour spaces, approximated by a cuboid

Colour Space	Volume	Colour Space	Volume
$L\alpha\beta$	6.4480	**YUV'**	0.1318
CIELab (E)	6.0447e+03	**HSV**	0.8333
CIELab (D65)	8.8841e+03	**XYZ**	1.0351
Yxy	0.2646	**RGB**	1.0000
YUV	0.0879		

channels. Simply normalising the data by its maximum – whether per channel or per set – is not sufficient as this would potentially change the relations between the three channels of each space and would not necessarily ensure that the results are comparable.

For our analysis, we choose to normalise the different colour spaces according to their volume. Each space is defined by 3 axes, with different ranges and ratios between them. To approximate the volume of each colour space, the maximum and minimum achievable values for each channel are determined. In some colour spaces, such as RGB, these are well defined. In cases such as CIELAB however, values may not be explicitly constrained. Instead, we constrain the input to our analysis such that the initial RGB values are normalised between 0 and 1. Using a simple 3-channel image containing all permutations of 0 and 1, we can determine the minimum and maximum values that can be achieved in each channel for each of the colour spaces that we are examining. These values are then used to compute the volumes of the spaces, shown in Table 1. Since the RGB input is normalised, the RGB space already has unit volume. The normalisation procedure then consists of scaling the image values by the cube root of the colour space's volume, which is sufficient to allow comparisons between them.

4.1   Standard Colour Spaces

Tables 2 and 3 show the resulting covariance matrices for each colour space and for each of our image ensembles.

From these covariance matrices, we take the off-axis elements and average their absolute values, leading to an average covariance for each colour space. We then use these numbers to rank each colour space, leading to the orderings shown in Table 4. The first colour space is thus the most effective at decomposing the specified category of images to decorrelated components.

The first observation is that for all image ensembles CIELAB with illuminant E decorrelates the data more effectively than all other colour spaces. Especially for day-time images, both natural and artificial, this space decorrelates significantly better than the second best colour spaces. Surprisingly, CIELAB with illuminant D65 does not decorrelate well. On the other hand, we see that RGB and XYZ are significantly worse than all other colour spaces. Finally, we note that the ranking is mostly consistent across image ensembles. We expect that this measure of correlation is indicative for the performance of these colour spaces in colour transfer tasks.

Table 2. Covariance matrices for the Natural and Manmade Day ensembles

Natural Day				Manmade Day			
$L\alpha\beta$	0.0462 -0.0059 -0.0002			$L\alpha\beta$	0.0611 -0.0051 -0.0004		
	-0.0059 0.0025 0.0001				-0.0051 0.0027 0.0002		
	-0.0002 0.0001 0.0000				-0.0004 0.0002 0.0000		
CIELab (E)	0.0062 -0.0016 -0.0002			CIELab (E)	0.0057 -0.0016 0.0003		
	-0.0016 0.0009 -0.0002				-0.0016 0.0009 0.0002		
	-0.0002 -0.0002 0.0013				0.0003 0.0002 0.0018		
CIELab (D65)	0.0048 -0.0030 -0.0107			CIELab (D65)	0.0044 -0.0028 -0.0099		
	-0.0030 0.0022 0.0061				-0.0028 0.0021 0.0065		
	-0.0107 0.0061 0.0287				-0.0099 0.0065 0.0278		
Yxy	0.0616 -0.0027 -0.0041			Yxy	0.0567 -0.0031 -0.0025		
	-0.0027 0.0010 0.0008				-0.0031 0.0015 0.0010		
	-0.0041 0.0008 0.0016				-0.0025 0.0010 0.0013		
YUV	0.1286 -0.0007 -0.0030			YUV	0.1183 -0.0024 -0.0021		
	-0.0007 0.0005 0.0001				-0.0024 0.0007 0.0002		
	-0.0030 0.0001 0.0004				-0.0021 0.0002 0.0004		
YUV'	0.0981 -0.0005 -0.0035			YUV'	0.0903 -0.0018 -0.0024		
	-0.0005 0.0004 0.0001				-0.0018 0.0006 0.0002		
	-0.0035 0.0001 0.0006				-0.0024 0.0002 0.0007		
HSV	0.0075 -0.0019 0.0019			HSV	0.0142 -0.0029 0.0018		
	-0.0019 0.0264 -0.0146				-0.0029 0.0262 -0.0077		
	0.0019 -0.0146 0.0303				0.0018 -0.0077 0.0311		
XYZ	0.0221 0.0234 0.0240			XYZ	0.0201 0.0214 0.0207		
	0.0234 0.0248 0.0253				0.0214 0.0228 0.0223		
	0.0240 0.0253 0.0276				0.0207 0.0223 0.0241		
RGB	0.0269 0.0254 0.0232			RGB	0.0260 0.0234 0.0190		
	0.0254 0.0258 0.0236				0.0234 0.0239 0.0209		
	0.0232 0.0236 0.0238				0.0190 0.0209 0.0208		

4.2   Ensemble-Specific Colour Spaces

We compute dedicated colour spaces for each of the four image ensembles. For comparison, we note that the $L\alpha\beta$ colour space, derived from 12 spectral images of natural environments, takes the form of 3 approximately orthonormal principal axes [23]:

$$L = (\log L + \log M + \log S)/\sqrt{3} \tag{5a}$$

$$\alpha = (\log L + \log M - 2\log S)/\sqrt{6} \tag{5b}$$

$$\beta = (\log L + \log M)/\sqrt{2} \tag{5c}$$

For comparison, the precise components for this colour space are replicated in Table 5. We find that these axes closely approximate the actual components resulting from our principal component analysis, which are shown with their corresponding eigenvalues λ in Table 6.

Table 3. Covariance matrices for the Indoors and Night ensembles

	Indoors				Manmade Night		
$L\alpha\beta$	0.1153	-0.0081	-0.0005	$L\alpha\beta$	0.0688	-0.0045	-0.0008
	-0.0081	0.0052	0.0003		-0.0045	0.0054	0.0004
	-0.0005	0.0003	0.0001		-0.0008	0.0004	0.0001
CIELab (E)	0.0088	-0.0024	0.0018	**CIELab (E)**	0.0058	-0.0016	0.0019
	-0.0024	0.0016	-0.0005		-0.0016	0.0018	-0.0005
	0.0018	-0.0005	0.0028		0.0019	-0.0005	0.0045
CIELab (D65)	0.0068	-0.0043	-0.0126	**CIELab (D65)**	0.0045	-0.0029	-0.0067
	-0.0043	0.0033	0.0080		-0.0029	0.0027	0.0043
	-0.0126	0.0080	0.0323		-0.0067	0.0043	0.0223
Yxy	0.0865	-0.0040	-0.0036	**Yxy**	0.0575	-0.0036	0.0001
	-0.0040	0.0025	0.0016		-0.0036	0.0029	0.0010
	-0.0036	0.0016	0.0025		0.0001	0.0010	0.0018
YUV	0.1803	-0.0028	-0.0028	**YUV**	0.1200	-0.0052	-0.0006
	-0.0028	0.0014	0.0003		-0.0052	0.0024	0.0003
	-0.0028	0.0003	0.0008		-0.0006	0.0003	0.0005
YUV'	0.1376	-0.0022	-0.0032	**YUV'**	0.0915	-0.0039	-0.0007
	-0.0022	0.0010	0.0004		-0.0039	0.0018	0.0003
	-0.0032	0.0004	0.0013		-0.0007	0.0003	0.0009
HSV	0.0145	0.0039	0.0023	**HSV**	0.0263	0.0009	0.0051
	0.0039	0.0422	-0.0095		0.0009	0.0282	-0.0021
	0.0023	-0.0095	0.0494		0.0051	-0.0021	0.0388
XYZ	0.0306	0.0326	0.0287	**XYZ**	0.0205	0.0216	0.0175
	0.0326	0.0348	0.0307		0.0216	0.0232	0.0188
	0.0287	0.0307	0.0303		0.0175	0.0188	0.0204
RGB	0.0424	0.0371	0.0275	**RGB**	0.0324	0.0248	0.0153
	0.0371	0.0363	0.0281		0.0248	0.0242	0.0170
	0.0275	0.0281	0.0254		0.0153	0.0170	0.0177

Table 4. Colour space rankings according to average covariance ($\times\ 10^{-3}$). E and *D65* indicate the CIELAB colour space using illuminants E and D65 respectively.

Rank	Natural Day		Manmade Day		Indoors		Night	
	Space	Corr.	Space	Corr.	Space	Corr.	Space	Corr.
1	E	0.7	E	0.7	E	1.6	E	1.4
2	Yuv	1.3	Yu'v'	1.5	Yu'v'	1.9	Yxy	1.6
3	Yu'v'	1.4	Yuv	1.6	Yuv	2.0	Yu'v'	1.7
4	$L\alpha\beta$	2.1	$L\alpha\beta$	1.9	$L\alpha\beta$	3.0	$L\alpha\beta$	1.9
5	Yxy	2.6	Yxy	2.2	Yxy	3.1	Yuv	2.0
6	HSV	6.1	HSV	4.1	HSV	5.2	HSV	2.7
7	D65	6.6	D65	6.4	D65	8.3	D65	4.6
8	RGB	24.1	RGB	21.1	XYZ	30.6	RGB	19.0
9	XYZ	24.2	XYZ	21.4	RGB	30.9	XYZ	19.3

Table 5. The three component axes of the $L\alpha\beta$ colour space and their corresponding standard deviations [23]

| Vector	L	M	S	σ
$\frac{1}{\sqrt{3}}L$	1.004	1.005	0.991	0.353
$\frac{1}{\sqrt{6}}\alpha$	1.014	0.968	-2.009	0.0732
$\frac{1}{\sqrt{2}}\beta$	0.993	-1.007	0.016	0.00745

Table 6. The three axes and their corresponding standard deviations for each of the image ensembles

Manmade Day					Natural Day				
Vector	L	M	S	σ	Vector	L	M	S	σ
$\frac{1}{\sqrt{3}}L^\star_{MD}$	0.9619	0.9749	1.0603	0.9174	$\frac{1}{\sqrt{3}}L^\star_{ND}$	0.9449	0.9510	1.0966	0.8354
$\frac{1}{\sqrt{6}}\alpha^\star_{MD}$	1.2488	0.8600	-1.9238	0.1241	$\frac{1}{\sqrt{6}}\alpha^\star_{ND}$	1.2116	0.9771	-1.8914	0.1032
$\frac{1}{\sqrt{2}}\beta^\star_{MD}$	0.9291	-1.0582	0.1301	0.0178	$\frac{1}{\sqrt{2}}\beta^\star_{ND}$	0.9568	-1.0386	0.0763	0.0159
Manmade Night					**Manmade Indoors**				
Vector	L	M	S	σ	Vector	L	M	S	σ
$\frac{1}{\sqrt{3}}L^\star_{MN}$	0.9871	1.0012	1.0115	0.9928	$\frac{1}{\sqrt{3}}L^\star_{IN}$	0.9796	0.9854	1.0341	1.2813
$\frac{1}{\sqrt{6}}\alpha^\star_{MN}$	1.2466	0.7585	-1.9674	0.1432	$\frac{1}{\sqrt{6}}\alpha^\star_{IN}$	1.2310	0.8236	-1.9510	0.1567
$\frac{1}{\sqrt{2}}\beta^\star_{MN}$	0.9123	-1.0677	0.1664	0.0230	$\frac{1}{\sqrt{2}}\beta^\star_{IN}$	0.9247	-1.0614	0.1354	0.0220

To simplify comparisons between the dedicated colour spaces computed through our analysis and $L\alpha\beta$, we will refer to the resulting axes as $L^\star_{\text{ens}}$, $\alpha^\star_{\text{ens}}$ and $\beta^\star_{\text{ens}}$ in decreasing eigenvalue order, with $ens \in \{ND, MD, IN, MN\}$. The component with the largest eigenvalue for the Manmade Day set for instance will be referred to as $L^\star_{MD}$. The resulting spaces are listed in Table 6.

The relations between the colour channels resulting from this analysis are similar to the ones in the $L\alpha\beta$ colour space. Specifically, the first channel in all cases, $L^\star$, is approximately equal to $L + M + S$, effectively encoding luminance information. The second channel, $\alpha^\star$, is close to $L+M-2S$, encoding information along a blueish-yellowish axis, while finally the third channel, $\beta^\star$, is close to $L - M$, encoding information along a reddish-greenish axis. Unsurprisingly, each of these colour spaces is therefore approximately an opponent space.

Nonetheless, there are individual differences between the various scene categories that merit further exploration. First, the colour spaces resulting from the natural ensemble (ND) have a larger L contribution in the $\alpha^\star$ channel compared to the α channel in $L\alpha\beta$. Second, the L contribution in the same channel is even larger for the MD ensemble, with values for M being smaller to their equivalent in $L\alpha\beta$. In fact, this trend continues in the Indoors and Night datasets. As we explore less 'natural' image classes, the $\alpha^\star$ channel becomes increasingly skewed, containing a larger contribution from L and a smaller contribution from M. Effectively, the yellow-blue axis in $L\alpha\beta$ space becomes more akin to an orange-blue axis.
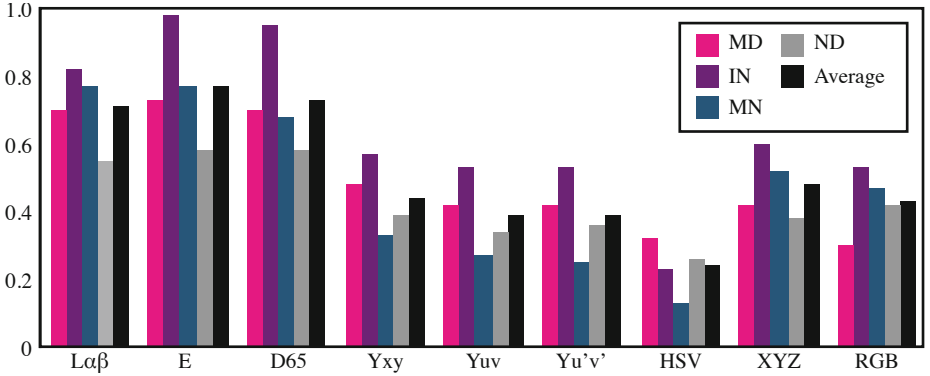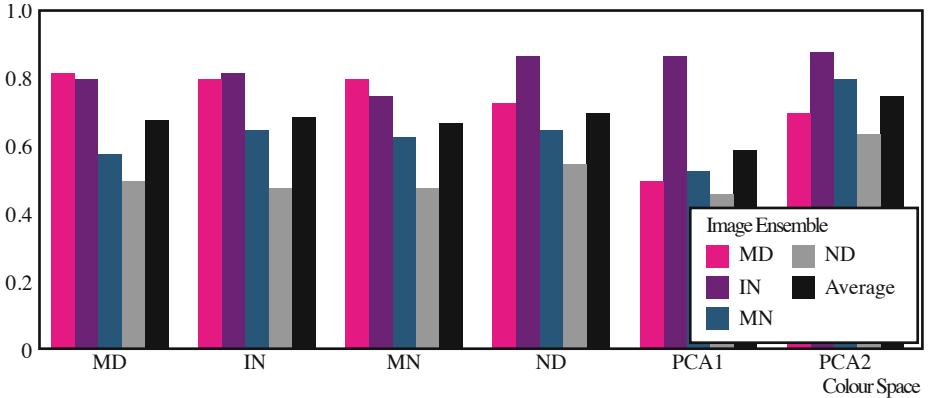
Fig. 3. Percentage of successful colour transfers for each colour space for each of the ensembles, as well as averaged over all four image ensembles. *E* and *D65* refer to CIELAB with illuminants E and D65 respectively.

Further, as we move to less natural sets, a second trend appears. In the third channel, $\beta^\star$, the S component contributes more than in typical opponent spaces, essentially replacing the red-green axis in $L\alpha\beta$ with an axis resembling purple-green. Although still significantly smaller compared to the L and M contributions, this cannot be discounted.

5   Colour Transfer

To understand whether these results can be used to inform the utility of colour spaces in the area of colour transfer, we have taken a subset of six images from each of the four image ensembles. Each combination of pairs was used to transfer colours from one image to the other. In this pilot experiment, we have counted the number of images for each colour space and image ensemble that we deem to be plausible results. This approach will only allow us to draw informal conclusions, as a definite result would require both more images as well as a full psychophysical experiment. However, we think that this informal experiment is useful as it may show a trend, and may steer a future validation study.

Figure 3 shows the percentage of successful colour transfers for each of the image ensembles, as well as averaged for all images in all ensembles. These percentages were computed by first running the colour transfer algorithm for all combinations of images within each ensemble (36 for the three manmade ensembles and 121 for the natural day ensemble). Both authors have independently observed each combination, assessing whether the colour transfer could be deemed successful. The criteria were deliberately kept informal, rejecting all cases where the palette was not transferred sufficiently, and all cases which have created obvious gamut problems. The results shown in the figure are averaged over the two sets generated by each author.

Fig. 4. Percentage of successful colour transfers for each ensemble-specific colour space for each of the ensembles, as well as averaged over all four image ensembles. We also show results for per-image PCA, carried out in the colour space derived from the source image (PCA1) and the target image (PCA2).

We see that both CIELAB variants as well as the $L\alpha\beta$ space tend to perform well, in particular for the indoors ensemble. A peculiar observation is that colour transfer for natural images appears to be a more difficult problem than for any of the other ensembles. This, however, may be attributed to the relatively small number of images involved in this experiment.

Further, the three spaces that have luminance-chrominance channel (Yxy, Yuv and Yu'v') as well as HSV do not perform well. This is predominantly due to the fact that the colour transfer algorithm produces out-of-gamut values that lead to visible artefacts. The RGB and XYZ spaces do not typically produce artefacts, but the transfer of colour tends to be too subtle.

Next, we analyse the ensemble-specific colour spaces that were created by computing PCA on each of our ensembles. By design, this means that colour transfer should be most effective when applied to images with their corresponding ensemble colour space. However, Figure 4 shows that this is not the case. Instead, it appears that the manmade day and indoor datasets have given rise to colour spaces that are effective for transfers in all four image ensembles. A further surprising result is that the natural day colour space is less effective than any of the other three colour spaces.

We have also tested whether ensemble statistics are fine-grained enough, or whether it would be advantageous to compute a separate colour space for each colour transfer. To this end, we have run PCA on the source images only, converting both source and target image to the resulting colour space, and applied the colour transfer algorithm in that space (PCA1). Conversely, we have done the same, instead running PCA on the target image (PCA2). The results are shown in Figure 4. We note that deriving a colour space from the target image is significantly more successful than doing the same on the source image. In fact,

this leads to a higher success rate than for any of the ensemble-specific colour spaces.

Comparing Figures 3 and 4 we note that the four most successful colour spaces are CIELAB (E) with a 77% overall success rate, followed by PCA2 (75% success), CIELAB (D65; 73% success) and $L\alpha\beta$ (71%) success. After that the ensemble-specific colour spaces perform reasonably well (67% - 70%). All other colour spaces do not perform acceptably well on our datasets.

Although one may argue that decorrelating individual images would be the most fine-grained approach to colour transfer, we find that similarly excellent results can be obtained by simply choosing the CIELAB space. Thus, it does not appear to be necessary to expend computational resources to deriving either category-based colour spaces or image-specific colour spaces for the purpose of colour transfer.

Finally, we ask whether these results are predicted by the covariance measures that we computed in the preceding section. Comparing Table 4 with Figure 3, it is clear that the average absolute covariance is a reasonable indicator of whether a colour space is suitable for colour transfer tasks. However, this measure does not take into account gamut issues that may arrise. Such problems are more prevalent in some colour spaces than others. In particular, the Yxy, Yuv, Yu'v' and HSV spaces lead to many out-of-gamut pixels, and therefore perform worse than may be expected on the grounds of their covariance statistics.

6   Conclusions

On the basis of four image ensembles we present an assessment of the viability of different colour spaces in the context of colour transfer. First, we find that a space's covariance with respect to each image ensemble is a good measure of how well such a space is likely to perform for colour transfer.

Second, we ask if one space could be found that is universally suitable. To this end we compare the performance of colour spaces derived for single images, those derived for our four image ensembles, and a set of general colour spaces. Surprisingly, we find that CIELAB, if used with illuminant E as the white point leads on average to the best performance, yielding a plausible colour transfer in 77% of all cases tested. This is slightly more than could be obtained with per-image PCA. CIELAB (E) also outperforms all other generic as well as ensemble-specific colour spaces. We therefore conclude that one space indeed fits all.

References

1. 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, pp. 491–494 (2004)
2. Abadpour, A., Kasaei, S.: An efficient PCA-based color transfer method. Journal of Visual Communication and Image Representation 18, 15–34 (2007)
3. An, X., Pellacini, F.: User controllable color transfer. Computer Graphics Forum 29, 263–271 (2010)

4. 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), pp. 91–98 (2004)
5. Greenfield, G., House, D.: Image recoloring induced by palette color associations. Journal of WSCG 11(1), 189–196 (2003)
6. Grundland, M., Dodgson, N.: The decolorize algorithm for contrast enhancing, color to grayscale conversion. Tech. Rep. UCAM-CL-TR-649, University of Cambridge (2005)
7. Kotera, H.: A scene-referred color transfer for pleasant imaging on display. In: IEEE International Conference on Image Processing, vol. 2, pp. 5–8 (2005)
8. Kotera, H., Morimoto, T., Saito, R.: Object-oriented color matching by image clustering. In: Proceedings of the 6th Color Imaging Conference, pp. 154–158 (1998)
9. Kumar, R., Mitra, S.K.: Motion estimation based color transfer and its application to color video compression. Pattern Analysis and Applications 11(2), 131–139 (2007)
10. Levin, A., Lischinski, D., Weiss, Y.: Colorization using optimization. ACM Transactions on Graphics 24(3), 689–694 (2004)
11. 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, 2006–2014 (2005)
12. Luan, Q., Wen, F., Xu, Y.Q.: Color transfer brush. In: Proceedings of Pacific Graphics, pp. 465–468 (2007)
13. Morovic, J., Sun, P.L.: Accurate 3D image colour histogram transformation. Pattern Recognition Letters 24, 1725–1735 (2003)
14. Neumann, L., Neumann, A.: Color style transfer techniques using hue, lightness and saturation histogram matching. In: Computational Aesthetics in Graphics, Visualization and Imaging, pp. 111–122 (2005)
15. Pitié, F., Kokaram, A., Dahyot, R.: N-dimensional probability density function transfer and its application to colour transfer. In: ICCV 2005: Proceedings of the 2005 IEEE International Conference on Computer Vision, vol. 2, pp. 1434–1439 (2005)
16. Pitié, F., Kokaram, A., Dahyot, R.: Automated colour grading using colour distribution transfer. Computer Vision and Image Understanding 107(2), 1434–1439 (2007)
17. Pouli, T., Reinhard, E.: Progressive histogram reshaping for creative color transfer and tone reproduction. In: NPAR 2010: Proceedings of the 8th International Symposium on Non-Photorealistic Animation and Rendering, pp. 81–90 (2010)
18. Pouli, T., Reinhard, E.: Progressive color transfer for images of arbitrary dynamic range. Computers and Graphics 35(1), 67–80 (2011)
19. Reinhard, E., Ashikhmin, M., Gooch, B., Shirley, P.: Color transfer between images. IEEE Computer Graphics and Applications 21, 34–41 (2001)
20. Reinhard, E., Shirley, P., Ashikhmin, M., Troscianko, T.: Second order image statistics in computer graphics. In: Proceedings of the 1st Symposium on Applied Perception in Graphics and Visualization, pp. 99–106 (2004)
21. 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)
22. Reinhard, E., Khan, E.A., Akyüz, A.O., Johnson, G.M.: Color Imaging: Fundamentals and Applications. A K Peters, Wellesley (2008)

23. Ruderman, D., Cronin, T., Chiao, C.: Statistics of cone responses to natural images: implications for visual coding. Journal of the Optical Society of America A 15(8), 2036–2045 (1998)
24. Tai, Y., Jia, J., Tang, C.: Local color transfer via probabilistic segmentation by expectation-maximization. In: IEEE Computer Society Conference on Computer Vision and Pattern Recognition (CVPR 2005), vol. 1, pp. 747–754 (2005)
25. Toet, A.: Natural colour mapping for multiband nightvision imagery. Information Fusion 4, 155–166 (2003)
26. Wang, C.M., Huang, Y.H., Huang, M.L.: An effective algorithm for image sequence color transfer. Mathematical and Computer Modelling 44, 608–627 (2006)
27. 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, vol. 6567, p. 65671G (2007)
28. Wen, C.L., Hsieh, C.H., Chen, B.Y., Ming, O.: Example-based multiple local color transfer by strokes. Computer Graphics Forum 27(7), 1762–1765 (2008)
29. Wyszecki, G., Stiles, W.S.: Color science: Concepts and methods, quantitative data and formulae, 2nd edn. John Wiley and Sons, New York (2000)
30. Xiang, Y., Zou, B., Li, H.: Selective color transfer with multi-source images. Pattern Recognition Letters 30, 682–689 (2009)
31. Xiao, X., Ma, L.: Color transfer in correlated color space. In: VRCIA 2006: Proceedings of the 2006 ACM International Conference on Virtual Reality Continuum and its Applications, pp. 305–309 (2006)
32. Xiao, X., Ma, L.: Gradient-preserving color transfer. Computer Graphics Forum 28, 1879–1886 (2009)
33. Xu, S., Zhang, Y., Zhang, S., Ye, X.: Uniform color transfer. In: IEEE International Conference on Image Processing, pp. III-940–III-943 (2005)
34. Yin, J., Cooperstock, J.R.: Color correction methods with applications to digital projection environments. Journal of the WSCG 12(1-3) (2004)
35. Zhao, Y., Wang, L., Jin, W., Shi, S.: Colorizing biomedical images based on color transfer. In: IEEE/ICME International Conference on Complex Medical Engineering, pp. 820–823 (2007)