A Multi–Resolution, Full–Colour Spatial Gamut Mapping Algorithm

*†Jan Morovic and †Yu Wang
*Hewlett Packard Company, Barcelona, Spain
†Colour & Imaging Institute, University of Derby, Derby, United Kingdom

Abstract

A multi–resolution, full–colour spatial gamut mapping algorithm (GMA) is proposed in this paper. Its aim is to maintain as much of an original image's overall, and in particular spatial, information as possible within the limits of a reproduction medium's gamut. First, the original image is decomposed into different spatial frequency bands. Second, lightness compression and initial gamut mapping are applied to the lowest frequency band image. Third, the next higher frequency band is added to the gamut mapped image and the result is processed by subsequent gamut mapping transformations. The third step is repeated until the highest frequency band is reached. The effect of this algorithm is that intra–image differences in the original image are well maintained in the gamut mapped reproduction. A psychophysical experiment is then described whose results show that this algorithm is in the pair of most accurate GMAs and can outperform all other algorithms tested here for images which are less accurately reproduced by all GMAs.

Introduction

Colour images have a wide variety of characteristics, ranging from the properties of their colours' distributions to what contents they represent. When attempting an accurate reproduction of colour images across media that have different colour gamuts it is therefore important to pay attention to the reproduction of all characteristics that an original image has. For example, original image features like having individual pixels of certain colours, having predominantly dark colours, having detail in certain parts of the image, or looking natural all ought to apply to the reproduction as well. Such attention to all image characteristics is particularly important when the reproduction medium has a colour gamut that is at least in some parts of colour space smaller than the original gamut.

Given the above challenge posed by accurate colour reproduction, it is worth looking at the properties of existing cross-media reproduction solutions to see whether they adequately address it. As most parts of cross-media reproduction workflows are descriptive (e.g. device characterization, colour appearance modelling), the work of preserving image characteristics beyond individual pixel colours falls to the gamut mapping algorithm (GMA). Looking at existing solutions to gamut mapping,[1] it can, however, be seen that the majority of them perform transformations that are determined only by factors derived from the original and reproduction media and a given original pixel's colour. As such these algorithms focus on colours of individual original image pixels and transform them without explicitly taking into account any other image characteristic, or at most taking into account the original image's colour gamut. A consequence of this is that when such algorithms are also intended for the reproduction of other image characteristics, and this is almost universally the intention of their authors, then their reproduction needs to be dealt with indirectly.

An important improvement as compared with such pixel–colour–only approaches is the GMA proposed by Braun and Fairchild,[2] which analyses an original image's lightness histogram and adjusts its behaviour accordingly. While such a method deals well with the distribution of original lightnesses, there is still significant room for improvement by addressing further important image characteristics.

The most obvious candidate for a next step is to improve the reproduction of an original's spatial properties. Over the years a handful of gamut mapping algorithms has already been proposed with the aim of explicitly dealing with this important characteristic.

In 1989 Meyer and Barth published the first spatial gamut mapping paper,[3] where the first step was lightness compression using low pass filtering in the Fourier spatial frequency domain. The dynamic range of the low pass filtered image was then compressed to that of the reproduction medium and the high pass filtered image detail information was added to it.

Later Nakauchi *et al.*[4] defined gamut mapping as an optimisation problem of finding such an image which is perceptually closest to a given original and has all pixels inside the reproduction gamut. As they defined perceptual difference by applying band–pass filters to Fourier–transformed CIELAB images and then weighting them according to human contrast sensitivity, they too performed spatial gamut mapping. McCann also proposed a solution to dealing with spatial detail in gamut mapping[5,6] and focuses on radiance ratios within images on the basis of the Retinex theory.[7]

Finally, Bala *et al.* describe a method,[8,9] which has similarities with the work of Meyer and Barth. Here the original image is first processed by a chroma–preserving gamut–clipping algorithm. Then, edges are extracted from the difference between the original and gamut–mapped luminance (or lightness) channels. A high-pass filtering method that takes the difference between a given pixel's value and the mean value of its neighbouring pixels is used. A window size of 3x3 pixels is proposed for business graphics and of 15x15 pixels for pictorials. This edge information is subsequently added to the gamut–mapped image and processed again using a luminance (or lightness) preserving gamut clipping method.
Against the background of the above proposals for spatial gamut mapping, we propose an alternative spatial gamut mapping algorithm operating in a multi–resolution and full–colour way. The aim of the algorithm is both to maintain colour appearance in the reproduction and to preserve spatial variance and its details will be discussed next.

Proposed Algorithm Framework

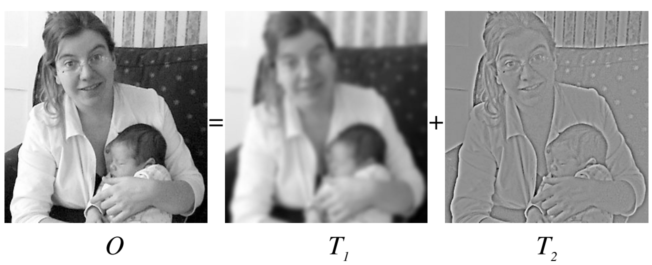

Figure 1. Two–resolution decomposition (r=3) (in T_2 mid–grey represents zero).

As has been noted previously it is more important in colour image reproduction to preserve the relationship of colours within images than to maintain their absolute values,[6-7] which is, in fact, impossible for some pixels. These relationships will here be expressed by the difference of a given pixel's colour from the mean of its neighbouring pixels in an $r \times r$ pixel region. This idea leads to the possibility of representing an original image (O) as a sum of two transformed images whereby the first transformed image (T_1) image contains at each pixel the mean value of the corresponding neighbourhood from O and the second transformed image (T_2) equals $O - T_1$ at each pixel (Fig. 1). Such a pair of transformed images then has spatially higher frequencies in T_2 and lower ones in T_1 and this allows for a different treatment of high versus low spatial frequency components.

Such a decomposition of an image can furthermore be repeated, whereby T_1 is taken as the original image and again represented by two transformed images (U_1 – lower frequency and U_2 – higher frequency image). As T1=U1+U2, O=U1+U2+T2 and the result is a three–resolution decomposition of the image (Fig. 2). Analogously greater numbers of bands can be used and this number can be meaningfully increased to the point where all pixels of the lowest frequency image will have the same value – i.e. the mean value of O.

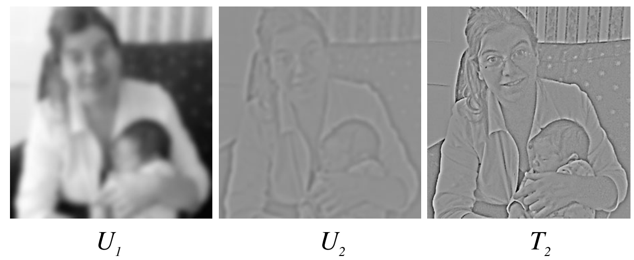

Figure 2. Three–resolution decomposition (r=3).

The point of using more than just two bands is that relationships can be considered not only between an original image's pixel and a neighbourhood of fixed area but also between neighbourhoods of different areas. Furthermore, differences between bands of this multi–resolution image representation can be computed in terms of all three dimensions of a colour space, rather than only in terms of lightness, as has been the case in some previous work.[3,8,9] This allows for dealing with local changes not only in lightness but also in chroma and hue.

Based on the above concepts, a multi–resolution and full–colour spatial gamut mapping algorithm (MSGM) is proposed in this paper. Its aim is to maintain an original image's overall colour appearance as well as spatial variation as much as is possible within the limits of a reproduction medium's gamut. This will be attempted by taking an original and first computing a multi–resolution decomposition of it. Then the lowest–resolution band will be gamut mapped and the difference between the lowest and next higher bands from the original decomposition will be added to it. The result will again be gamut mapped and the process will be repeated until all bands from the original decomposition are incorporated again into the gamut mapped image.

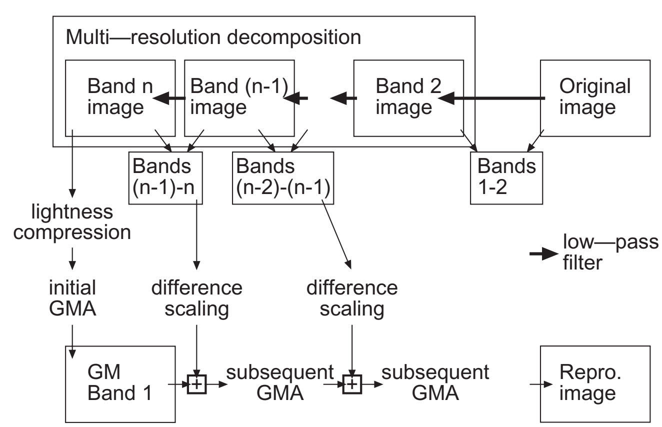

Figure 3. Framework of MSGM algorithm.

*In detail t*he basic framework of the proposed algorithm (Fig. 3) is the following whereby the effect of several options will be explored at each step:

1. **Multi-resolution decomposition** with various filter sizes r = {3, 7, 13, 19, 25, 31} and decomposition bands n = {2, 4, 6, 8, 10} and using mean filtering (MF) to obtain low-pass versions of an image. Note that image is filtered in all three channels of CMA97s2 *Jab*, as opposed to the processing being restricted to lightness, which is more typical of previously proposed spatial GMAs.
2. **Lightness compression** of the lowest spatial-frequency band either in a linear (LC) or sigmoidal[2] (SC) way or not at all.
3. **Initial gamut mapping** using the hue–preserving minimum ΔE clipping method (HPMINDE) for the lowest spatial frequency band (band n). This step establishes approximate colours for the reproduction pixels, which will be modified by subsequent steps. Other options for this gamut mapping are clipping towards the lightness of the cusp (CUSP) and clipping towards the centre of colour space *Jab*=[50,0,0] (SCLIP).
4. **Addition of difference** between the current (i) and next higher (i-1) band from the original image decomposition to the gamut mapped reproduction at band i. Here the absolute difference can be added or it can first be linearly compressed according the ratio of the lightness ranges of the reproduction and original media.
5. **Subsequent gamut mapping** using SCLIP clipping to preserve more spatial detail. Other options are HPMINDE and CUSP clipping.
6. If $n > 2$ and $i > 2$, go to step 4, otherwise terminate.

Given the above framework, the underlined option in each of the above steps is used to define a reference MSGM algorithm (rMSGM). The following sections will informally explore the effect of varying one of these options at a time as compared with the result of rMSGM.

Exploring Effects Of MSGM Options

Performance Of rMSGM

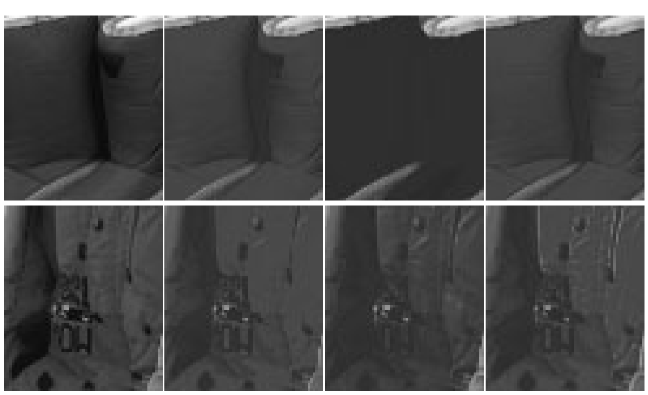

Figure 4. Effect of GMAs on image properties (left to right: original, HPMINDE, SGCK and rMSGM).

First, however, it is useful to see how rMSGM compares with existing non–spatial GMAs. Fig. 4 therefore shows how a number of parts from the SKI image are treated by rMSGM and the HPMINDE and SGCK algorithms. This image and the latter two GMAs are specified in the CIE TC8–03 *Guidelines for the Evaluation of Gamut Mapping Algorithms*.[10]

Using rMSGM shows that it performs better for these image parts than either HPMINDE or SGCK. It reduces colour differences as compared with SGCK especially for the blue boots and it also preserves more spatial detail than HPMINDE. However, for the red coat, rMSGM preserves more spatial detail than HPMINDE, but less than SGCK. The overall performance of the new algorithm can be seen to show a combination of the strengths of HPMINDE and SGCK as it can preserve both colour appearance and spatial information.

Filter Size And Number Of Decomposition Bands

The first pair of options – filter size and number of decomposition bands – relates to how the multi–resolution decomposition is computed. Here increasing filter size results in more spatial detail being maintained, however, with the drawback of clipping and halo (or blooming) artefacts becoming more serious. In Fig. 5 the effect of changing filter size is shown and it can be seen that while r=3 does not preserve detail as well as other options, going beyond r=7 only introduces artefacts.

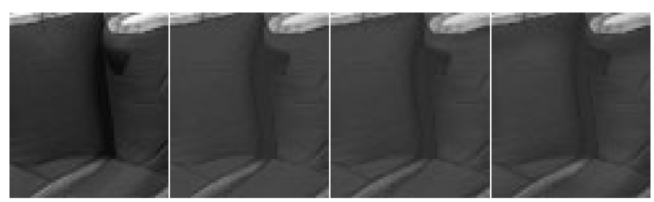

Figure 5. Effect of varying filter size. (left to right: original, r=3, r=7 and r=25).

In terms of reproduction quality increasing the number of decomposition bands has a similar effect to increasing filter size. However, the clipping and halo artefacts are not as strong as those for increased filter size. This is so partly because the combined effect of having multiple bands and a fixed filter size is in effect not a full mean filtering but a centrally weighted filtering. From among the numbers of decomposition bands considered here, four gave the best compromise between detail preservation and absence of artefacts.

Initial Lightness Compression

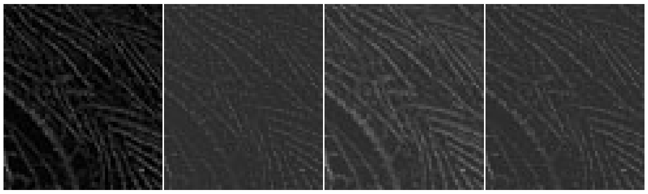

Figure 6. Initial lightness compression (left to right: original, none, linear, sigmoidal).

As the lowest frequency band contains information about the overall colour appearance of image regions, its reproduction could benefit from more advanced pixel–colour–only methods developed previously. Fig. 6 therefore shows the effect of having initial lightness compression, performed either linearly or in a sigmoidal way, as compared to not having any initial lightness processing, as is the case in rMSGM.

Figure 6 shows that dark neutral region can be reproduced better when lightness compression is used. However, linear compression can amplify detail while sigmoidal compression gives the closest match to the original.

Initial GMA

The choice of initial GMA is of great importance in the context of this algorithm as it serves two purposes. First, its aim is to give all parts of a reproduction a similar colour appearance to the corresponding parts of the original. Second, it needs to provide a starting point from which the subsequent addition and gamut mapping of higher frequency bands can preserve detail. While the first reason lead to the use of HPMINDE in the reference version of MSGM, Figure 7 shows that the SCLIP algorithm gives better results in terms of detail preservation. The reason for this is that while HPMINDE is, by definition, good for reproducing individual colours well, it suffers from mapping entire 2D regions of colour space onto single reproduction gamut boundary colours. SCLIP on the other hand always maps only 1D regions (i.e. lines) onto single colours and it therefore less likely to remove variation. While the CUSP clipping algorithm might in theory seems like a better choice than SCLIP, it can cause serious clipping artefact when the cusp is close to the extremes of the lightness range (second row of Fig. 7).

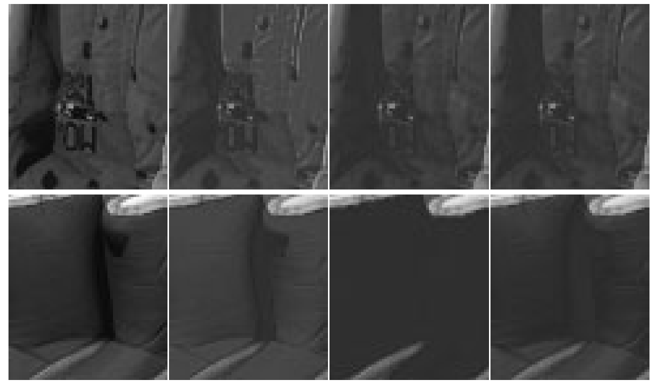

Figure 7. Effect of initial GMA (left to right: original, HPMINDE, CUSP, SCLIP).

Subsequent GMA

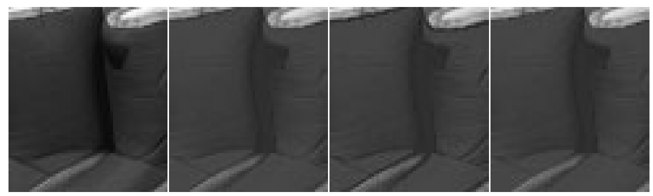

Figure 8. Effect of subsequent GMA (left to right: original, SCLIP, CUSP, HPMINDE).

The role of this algorithm is exclusively the preservation of spatial variation and the effect of using HPMINDE and the CUSP algorithm instead of SCLIP, which was chosen for rMSGM is shown in Fig. 8. Here it can be seen that the difference between using SCLIP and HPMINDE is negligible, while using CUSP results in some edge artefacts.

Proposed MSGM Options

Based on the above investigation of options affecting the proposed framework, the variant of the MSGM algorithm that will be psychophysically evaluated is the following:

First, use a low-pass mean filter with a size of 7×7 pixels and compute a 4–band decomposition of an original CAM97s2 *Jab*.[11] Second, use sigmoidal lightness compression on the lightness (*J*) channel of the lowest frequency band image, then apply initial gamut mapping using the SCLIP algorithm (clipping towards point on lightness axis with J=50) to the lightness–rescaled lowest frequency band image. Third, the difference between the current and the next higher frequency images is added to the current gamut mapped image after a linear compression using the ratio of the reproduction medium and original medium lightness ranges. Fourth, subsequent gamut mapping is performed using the SCLIP algorithm. The third and fourth steps are repeated until the highest frequency band is included in the reproduction.

Evaluation Of GMA Accuracy

To evaluate the performance of the proposed spatial gamut mapping algorithm, a psychophysical experiment was carried out following draft CIE TC8-03 *Guidelines*.[10] In this experiment the following six GMAs were compared: HPMINDE; SGCK; the basic spatial GMA proposed by Bala *et al.* (XSGM);[8-9] the spatial GMA proposed by Bala et al. that used inverse-gamma-inverse[12] lightness compression (XIGI-SGM); MSGM4 – the new spatial gamut mapping proposed in this paper; MSGM2 – a simplified version of MSGM4 using only two decomposition bands.

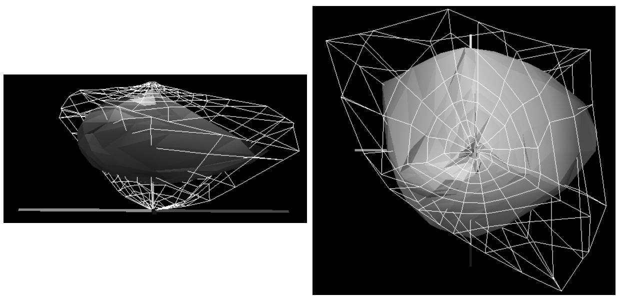

Figure 9. CAM97s2 colour gamuts of CRT (mesh) and printed medium (solid) used in experiment.

A total of 15 test images were used in this experiment whereby they were selected following the *Guidelines* and included a broad range of image types. For each original, the six gamut mapped images were printed on plain paper using a *Canon BJC-6100* bubble-jet printer and all images had a white border. The images sent to the printer had the same resolution as the originals displayed on an *Apple* 21-inch Studio Display CRT. Reproductions were viewed in a viewing booth with a light source simulating illuminant D65 and against a mid-grey background. The CRT's white point was set to be as close as possible to the plain paper's white in the viewing booth. Both monitor and viewing booth were set up in a dark room side–by–side and viewed by observers from approximately 75 cm distance. The colour gamuts of the original and reproduction media are shown in Fig. 9.

A category judgement technique[13] was then used and 15 observers judged the accuracy of an image's reproduction on an equi–interval accuracy scale from zero to six. Here zero represents the least accurate reproduction while six represents the most accurate reproduction. Observers were asked to judge into which category each particular stimulus belongs.

Experimental Results

To compare the accuracy of the chosen GMAs, the mean category values for each reproduction made by a given GMA were computed. From these values (Fig. 10), it can be seen that the XIGI-SGM and MSGM4 algorithms are significantly more accurate than the other methods, while not being significantly different from each other. The worst accuracy is had by HPMINDE and XSGM. Note that the scale on which these results are reported is very different from pair comparison scales, which are relative to a given set of images, whereas the present scale was defined for observers to span the entire range of accuracy levels they can imagine.

In order to get a better understanding of the experimental results additional methods of analysis are proposed here. In addition to considering the overall, mean performance of an algorithm for a set of test images, it is also of value to know how often each algorithm is in the group most accurate and statistically indistinguishable group of GMAs.

In addition to the above information, which again shows a similar result to what is seen in Fig. 10, it is also of interest to see what the same statistics are in a way that takes into account the mean accuracy of all reproductions of a given original. In other words this will show whether a given algorithm is in the most accurate group predominantly for images that are reproduced accurately or inaccurately using all GMAs. The frequency of a GMA being in the most accurate group as a function of original image mean accuracy is therefore shown in Fig. 12.

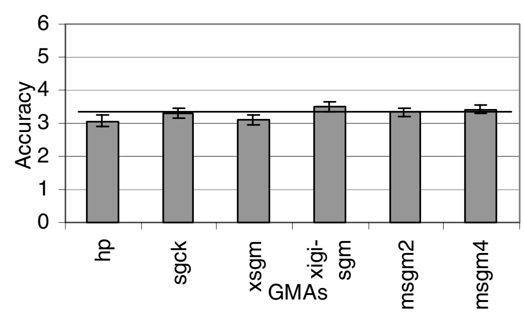

Figure 10. Overall accuracy of evaluated GMAs (error–bars represent 95 percent of judged category values).

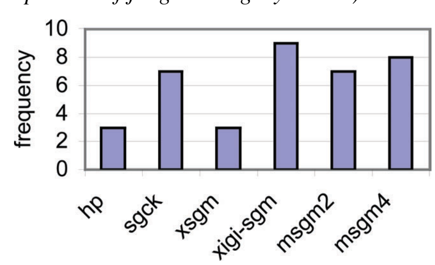

Figure 11. Frequency of GMA being in most accurate group.

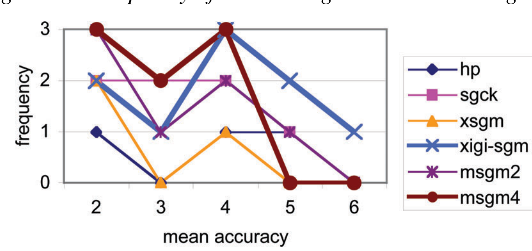

Figure 12. Frequency of being in most accurate group as function of mean accuracy.

Here it can be seen that at one end of the scale HPMINDE and XSGM are very unlikely to be the most accurate GMA for any kind of image, even though they do perform well for some specific images. Further it is interesting to compare the two GMAs that perform most accurately overall. For images having mean category values larger than 4, XIGI-SGM outperforms MSGM4 but for images having mean category values smaller than 4, MSGM4 outperforms XIGI-SGM. As in practice it is absolute performance that is important, MSGM4 has the advantage over all other methods tested here that it performs better especially for images that on average are reproduced inaccurately, while for images that are reproduced accurately by all algorithms it is not very different from the mean. Further details of the experimental results can be found elsewhere.[14]

Conclusions

The gamut mapping algorithm presented here performs full–colour gamut mapping of an image by transforming colour information in it sequentially for a number of spatial frequency bands. Its starts with an initial lightness compression followed by a gamut mapping transformation of the lowest frequency band. This transformation intends both to preserve the colour appearance of image regions and to provide a starting point for the processing of subsequent bands that allows for the reproduction of the original's spatial variation. In exploring the various options of such an approach a set of choices was made that informally resulted in the most faithful reproduction of original appearance.

The performance of new spatial GMA proposed here (MSGM4) was then evaluated psychophysically and it was shown that it overall performs equally with a previously published algorithm proposed by Bala *et al.*[8,9] The advantage of the present algorithm is, however, that it performs particularly well in relative terms for images that are not reproduced accurately by any of the GMAs evaluated here. As such this new GMA would be the best choice from among the tested algorithms since it works better than the others for images that are difficult to accurately reproduce and since it also works well for images that all of these algorithms reproduce accurately.

Acknowledgements

We would like to thank Tadashi Yoshida and Todd Newman (Canon Development Americas) for supporting this work. Many thanks go to Dr. Pei-Li Sun (Shih Hsin University) for his valuable help in conducting the experiment and to Dr. Raja Bala (Xerox) for advice on correctly implementing his GMAs. We would also like to thank all the observers at the Colour & Imaging Institute (UK) for their contribution to the psychophysical evaluation.

References

1. J. Morovic, Colour Gamut Mapping. In *Colour Engineering: Achieving Device Independent Colour*, Green P. and MacDonald L. W. (eds.), John Wiley & Sons, 297–318 (2002).
2. G. J. Braun and M. D. Fairchild, Image Lightness Rescaling Using Sigmoidal Contrast Enhancement Functions, *J. Electronic Imaging*, **8**(4): 380-393 (1999).
3. J. Meyer and B. Barth, Color gamut matching for hard copy, *SID Digest*, 86-89 (1989).
4. S. Nakauchi, S. Hatanaka, S. Usui, Color gamut mapping based on a perceptual image difference measure. *Color Research and Application*, **24**(4): 280–290 (1999).
5. J. McCann, Lessons learned from Mondrians applied to real images and color gamuts. *Proc. IS&T/SID 7th Color Imaging Conference*, 1-8 (1999).
6. J. McCann, Color gamut mapping using spatial comparisons. *Proc. of SPIE*, **4300**: 126-130 (2001).
7. E. Land, The Retinex, *American Scientist*, **52**:247–264 (1964).
8. R. Bala, R. DeQueiroz, R. Eschbach and W. Wu, Gamut mapping to preserve spatial luminance variations, *Proc. of IS&T/SID 8th Color Imaging Conference*, 122-128 (2000).
9. R. Bala, R. DeQueiroz, R. Eschbach and W. Wu, Gamut mapping to preserve spatial luminance variations, *J. Imaging Science and Technology*, **45**(5): 436-443 (2001).
10. CIE TC8–03, Guidelines for the Evaluation of Gamut Mapping Algorithms, *Technical Report, Draft No. 18M*, (2002).
11. C. J. Li, M. R. Luo and R. W. G. Hunt, The CAM97s2 Model, *Proc. of IS&T/SID 7th Color Imaging Conference*, 262-263 (1999).
12. K. Braun, R. Bala and R. Eschbach, Development and Evaluation of Six Gamut Mapping Algorithms for Pictorial Images, *Proc. IS&T/SID 7th Color Imaging Conference*, Springfield VA, 144-148 (1999).
13. W. S. Torgerson, A Law of Categorical Judgment, *Consumer Behaviour*, L. H. Clark (ed.), New York University Press, New York, 92–93 (1954).
14. J. Morovic and Y. Wang, Influence Of Test Image Choice On Experimental Results, *IS&T/SID 11th Color Imaging Conference, submitted for publication* (2003).

Biography

Ján Morovic received a Ph.D. in Colour Science at the Colour & Imaging Institute (CII) of the University of Derby (UK) in 1998, where the title of his thesis was *To Develop a Universal Gamut Mapping Algorithm*. He then worked as a lecturer in digital colour reproduction at the CII and is currently an image quality engineer at Hewlett Packard in Barcelona, Spain. He is also the chairman of the CIE Technical Committee 8-03 on Gamut Mapping.