

Image and Video Decolorization by Fusion

Non Peer-reviewed author version

ANCUTI, Codruta; ANCUTI, Cosmin; HERMANS, Chris & BEKAERT, Philippe
(2010) Image and Video Decolorization by Fusion. In: Proceedings of the Asian
Conference on Computer Vision (ACCV 2010)..

Handle: <http://hdl.handle.net/1942/11618>

Image and Video Decolorization by Fusion

Codruta O. Ancuti, Cosmin Ancuti, Chris Hermans, Philippe Bekaert

Hasselt University - tUL -IBBT,
Expertise Center for Digital Media, Wetenschapspark 2, Diepenbeek, 3590, Belgium

Abstract. In this paper we present a novel decolorization strategy, based on image fusion principles. We show that by defining proper inputs and weight maps, our fusion-based strategy can yield accurate decolorized images, in which the original discriminability and appearance of the color images are well preserved. Aside from the independent R, G, B channels, we also employ an additional input channel that conserves color contrast, based on the Helmholtz-Kohlrausch effect. We use three different weight maps in order to control saliency, exposure and saturation. In order to prevent potential artifacts that could be introduced by applying the weight maps in a per pixel fashion, our algorithm is designed as a multi-scale approach. The potential of the new operator has been tested on a large dataset of both natural and synthetic images. We demonstrate the effectiveness of our technique, based on an extensive evaluation against the state-of-the-art grayscale methods, and its ability to decolorize videos in a consistent manner.

1 Introduction

Although color plays an important role in images, applications such as compression, visualization of medical imaging, aesthetical stylization, and printings require reliable decolorized image versions. The widely-used standard color-to-grayscale conversion employs the luminance channel only, disregarding the important loss of color information. In many cases, a decolorized image obtained in this way will not fulfill our expectations, as the global appearance is not well preserved (illustrated in Figure 1). This limitation of the standard transformation is due to the fact that isoluminant regions are mapped onto the same output intensity.

In this paper we present a novel decolorization method, built on the principle of image fusion. This well-studied topic of computational imaging has found many useful applications, such as single image dehazing [1], interactive photomontage [2], image editing [3], image compositing [4, 5] and HDR imaging [6, 7]. The main idea is to combine several images into a single one, retaining only the most significant features.

The main difference between fusion methods that makes them application-specific, is the choice of inputs and weights. Our algorithm employs the three independent RGB channels and an additional image that conserves the color




Fig. 1. In contrast with the standard color-grayscale approach, our method seeks to preserve the global appearance.

contrast, based on Helmholtz-Kohlrausch effect, as image inputs. This fourth input better preserves the global appearance of the image, as it enforces a more consistent gray-shades ordering. The weights used by our algorithm are based on three different forms of local contrast: (a) a saliency map which helps us preserve the saliency of the original color image; (b) a second weight map that advantages well-exposed regions; and (c) a chromatic weight map which enhances the color contrast in addition to the effect of $H - K$ input. In order to minimize artifacts introduced by the weight maps, our approach is designed in a multi-scale fashion, using a Laplacian pyramid representation of the inputs combined with Gaussian pyramids of normalized weights.

To the best of our knowledge we are the first that introduce a fusion-based decolorization technique. Our method performs faster than existing color-to-gray methods since it does not employ color quantization [8], that tends to introduce artifacts, or cost function optimization, which commonly is computationally expensive (e.g. Gooch et al. [9] approach) and risks not converging to a global extremum. Our new operator has been tested on a large dataset of both natural and synthetic images. In addition, we demonstrate that our operator is able to decolorize videos. Our multi-scale fusion approach demonstrates consistency over varying palettes, and is able to maintain temporal coherence within videos. Furthermore, we have performed a comparative evaluation of the contrast enhancement qualities of the recent state-of-the-art color-to-grayscale techniques.

2 Related Work

Recently, grayscale image conversion has received an increasing amount of attention in the computer vision and graphics communities. There have been attempts to solve dimensionality reduction problem by both local [10, 8, 9, 11] and




Fig. 2. An overview of our fusion-based approach. Based on the original input image, we derive four input images (R , G , B and $H - K$ lightness) and three weight maps that blended by a multi-scale image fusion strategy yields the decolorized output.

global [12–14] mapping strategies, using different linear but also non-linear mapping techniques.

The gamut-mapping method of Bala et al. [10] combines the luminance channel with a high-pass filtered chroma channel. Gooch et al. [9] attempt to preserve the sensitivity of the human visual system by iteratively comparing each color pixel value with the average of its surrounding region. Since, in their experiments, small neighborhoods appear to introduce artifacts, the entire image has been used as neighbor region by default. Similarly, the method of Rasche et al. [8] computes the distribution of all the image colors previously quantized in a number of *landmark* points. The main drawback of these methods [9, 8] is that they are computationally expensive, as their computation time depends heavily on the number of colors within the image.

Queiroz et al. [15] generate grayscale images that encode color information in the output image. The colors are mapped onto low-visibility high-frequency texture, which can be identified by a decoder, and finally recovered. Their work has proven to be a practical solution for office documents. Neumann et al. [12] compute the gradient field in both the $CIEL^*a^*b$ and *Coloroid* [16] color space. Based on an extensive user-study, the indices of relative luminance are determined. The global mapping technique of Grundland and Dodgson [13] performs a dimensionality reduction using predominant component analysis, which is similar to principal component analysis (PCA). This technique reduces the processing cost substantially, compared to previous approaches. The multi-spectral method of Alsam and Drew [17] improves the approach of Socolinsky and Wolff [18] in terms of computation efficiency. Both approaches aim to preserve the *maximum local contrast*.

Smith et al. [11] developed a decolorization algorithm that exploits the Helmholtz-Kohlrausch effect, by applying the lightness measure of Fairchild and Pirrotta [19]. The algorithm uses a multi-scale chromatic filter to enhance the discriminability over the salient color features. Additionally, the authors have proven the applicability of their method to decolorizing videos. The recent technique of Kim et al. [14] optimizes a nonlinear global mapping function. The




Fig. 3. a. The four image fusion inputs (R , G , B and $H - K$ *Lightness*) and the corresponding Gaussian pyramids of the image fusion weights (b. saliency; c. exposedness ; d. chromatic).

method is built on the Gooch et al. [9] approach, but is computationally more effective.

In contrast with the existing techniques, our decolorization algorithm employs a multi-scale fusion strategy. The method is straightforward to implement, and only uses classical concepts. By selecting the appropriate weights and inputs, our approach has demonstrated robustness and consistency in decolorizing both images and videos.

3 Fusion-based Decolorization Approach

The standard grayscale transformation tends to reduce the amount of variations and sharpness within an image. Qualitatively, the dull appearance is due to the loss of contrast that is more visually noticeable on dimmed highlights and shadows. In order to obtain pleasing decolorized images, photographers might compensate these limitations by tedious work in the darkroom, applying elaborate lighting techniques or using photo-editor programs to manually adjust the contrast, luminance or histogram distribution.

We argue that the image appearance in black-and-white is tightly connected with models of color appearance, and that measurable values like salient features and color contrast are difficult to integrate by simple per pixel blending, without

introducing artifacts into the image structure. For this reason, we have opted for the multi-scale approach of image fusion, combining the Helmholtz-Kohlrausch lightness predictor [19] with a set of pixel weights depending on important image qualities. This will ensure that regions with superior gain are well depicted in the decolorized image. Practically, the resulted grayscale image is obtained by fusing four input images (a lightness image that incorporates the $H - K$ effect, and the R , G , B color channels), weighted by normalized coefficients maps determined by saliency, pixel exposure, and chromatic weights. An overview of our approach is given in Figure 2.

3.1 $H - K$ Chromatic Adapted Lightness

Our algorithm requires four input images to be used in the fusion process. Besides the color channels R , G , B , we define an additional input that preserves the global contrast, based on Helmholtz-Kohlrausch effect. As observed by Smith et al. [11], the $H - K$ effect can be used to resolve potential ambiguities regarding the difference between the isoluminant colors. Therefore, given two isoluminant patches, the most colorful one will be mapped onto a brighter output intensity. For this fusion input channel, we used Fairchild’s chromatic lightness metric [19], which predicts the $H - K$ effect, defined in the $CIE L^*c^*h^*$ color space by the expression:

$$L_{H-K} = L^* + (2.5 - 0.025L^*)(0.116 \left| \sin \left(\frac{h^* - 90}{2} \right) \right| + 0.085)c^* \quad (1)$$

This L_{H-K} predictor has also been used in the work of Smith et al. [11], in which it was demonstrated to be more appropriate for the task of image decolorization than the chromatic lightness metric of Nayatani [20]. Nayatani [20] predictor often tends to map bright colors to white, which makes it harder to discriminate between images that contain bright isoluminant colors. However, as can be observed in the comparative results (see the image but as well video results of Smith et al. [11]), relying only to the Helmholtz-Kohlrausch effect the decolorized outputs might not preserve accurately the original saliency (this feature is mostly ensured in our operator by integrating the saliency weight map).

3.2 Weight Maps Assignment

In the following section, we present how the weight maps are defined in our fusion-based decolorization algorithm. Our approach is based on the principle that the output decolorized image needs be both visually pleasing and meet the application requirements. In the case of grayscale conversion, aside from the luminance which is the main contributor to the perceived lightness, there are also several other image qualities that guide our visual system during its analysis of the incoming light. Practically, the attention of an observer tends to be focused on the salient regions that stand-out within their neighborhood. In order to maintain this focus, it is desirable that these predominant regions are well preserved by the grayscale version. Therefore, in order to meet this




Fig. 4. From color to gray: from the original color image (a), we obtain our decolorized result (b) by applying an image fusion approach, using the four inputs (c), weighted by the corresponding normalized weight maps (d).

requirement, we first introduce a *saliency weight map*. Furthermore, as commonly the over- and underexposed regions are advantaged by the saliency map, we also define an *exposedness weight map* that overcomes perception degradation in these regions. Finally, we assign a third weight map, the *chromatic weight map*, which has the main goal of balancing the influence of chromatic stimuli into the perception of lightness. Practically, by smoothly fusing the input channels weighted by these weight maps, the original consistency of the image is well preserved, while ghosting and haloing artifacts are reduced. Moreover, we believe that these weight maps are intuitive concepts for the users.

Saliency weight map ($\mathcal{W}_S$) reveals the degree of conspicuousness with respect to the neighborhood regions. For this measurement, our algorithm employs the recent saliency algorithm of Achanta et al. [21]. Their strategy is inspired by the biological concept of center-surround contrast. The saliency weight at pixel position (x, y) of input I^k is defined as:

$$\mathcal{W}_S(x, y) = \|I_\mu^k - I_{\omega_{hc}}^k\| \quad (2)$$

where I_μ^k represents the arithmetic mean pixel value of the input I^k while $I_{\omega_{hc}}^k$ is the blurred version of the same input that aims to remove high frequency noise and textures. $I_{\omega_{hc}}^k$ is obtained by employing a small 5×5 ($\frac{1}{16}[1, 4, 6, 4, 1]$) separable binomial kernel with the high frequency cut-off value $\omega_{hc} = \pi/2.75$. For small kernels the binomial kernel is a good approximation of its Gaussian counterpart, but it can be computed more effectively. The approach of Achanta et al. [21] is very fast, and has the additional advantage of the extracted maps being characterized by well-defined boundaries and uniformly highlighted salient regions, even at high resolution scales. Based on extensive experiments, we found that this saliency map tends to favor highlighted areas. In order to increase the accuracy of results, we introduce the exposedness map to protect the mid tones that might be altered in some specific cases.

Exposedness weight map ($\mathcal{W}_E$) estimates the degree to which a pixel is exposed. The function of this weight map is to maintain a constant appearance of

the local contrast, neither exaggerated nor understated. Practically, this weight avoids an over- or underexposed look by constraining the result to match the average luminance. Pixels are commonly better exposed when they have normalized values close to the average value of 0.5. Inspired by the approach of Mertens et al. [7], who employ a similar weight in the context of tone mapping, the exposedness weight map is expressed as a Gaussian-modeled distance to the average normalized range value (0.5):

$$\mathcal{W}_E(x, y) = \exp\left(-\frac{(I^k(x, y) - 0.5)^2}{2\sigma^2}\right) \quad (3)$$

where $I^k(x, y)$ represents the value of the pixel location (x, y) of the input image I^k , while the standard deviation is set to $\sigma = 0.25$. This mapping conserves those tones that are characterized by distance close to zero, while larger distance values are related with the over- and underexposed regions. As a result, the impact of over- and underexposed regions filtered by the saliency map is tempered, keeping the original image appearance well preserved.

Chromatic weight map ($\mathcal{W}_C$) controls the saturation contribution of the inputs in the decolorized image. This is expressed as the *standard deviation* between every input and the saturation S (in *HSL* color space) of the original image. Due to the fact that in general humans prefer increased saturation, it is desirable that more saturated areas are mapped onto brighter tones. This balances the chromatic contrast loss with the desired amount of enhancement. We have observed that the impact of this gain is reduced for the $H-K$ chromatic adapted lightness input.

In our framework these weight maps (saliency, exposedness and chroma) have the same contribution to the resulted decolorized images. As an example, Figure 3 shows the computed weights for the considered inputs.

3.3 Multi-scale Fusion of the Inputs

Having defined the inputs (R, G, B color channels and $H-K$ chromatic adapted lightness) and the weight maps, in the following section we present how this information is blended by our fusion strategy. As previously mentioned, during the fusion process the inputs are weighted by specific maps in order to conserve the most significant features, and finally combined into a single output image:

$$\mathcal{F}(x, y) = \sum_{k=1}^K \bar{\mathcal{W}}^k(x, y) \mathcal{I}^k(x, y) \quad (4)$$

where the value of every pixel location (x, y) of the fused result $\mathcal{F}$ is obtained by taking the sum of the corresponding locations of the inputs $\mathcal{I}^k$ (k is the input index), weighted by the normalized weight maps $\bar{\mathcal{W}}^k$. The number of the inputs is counted by the index k (in our case $K = 4$). The normalized weights $\bar{\mathcal{W}}$ are obtained by normalizing over the M weight maps $\mathcal{W}$ ($M = 3$) in order that the

value of each pixel (x, y) weights to sum up to unity ($\sum \mathcal{W}^k = 1$ for each pixel location) (see Figure 4).

Unfortunately, applying eq. 4 directly sometimes introduces haloing artifacts, mainly in locations close to strong transitions between weight maps. In order to solve this problem, a more effective strategy needs to be devised. Generally, this task is solved by multi-scale decomposition strategies that use linear [22, 23] or non-linear filters [24–26]. While the class of non-linear filters has shown to be better at preserving edges, the linear filters are computationally more effective. Even though more refined multi-scale solutions might be applied as well, we have opted for the classical multi-scale Laplacian pyramid decomposition [22]. In this linear decomposition, every input image is represented as a sum of patterns computed at different scales based on the Laplacian operator. The inputs are convolved by a Gaussian kernel, yielding a low pass filtered versions of the original. In order to control the cut-off frequency, the standard deviation is increased monotonically. To obtain the different levels of the pyramid, initially we need to compute the difference between the original image and the low pass filtered image. From then on, the process is iterated by computing the difference between two adjacent levels of the Gaussian pyramid. The resulting representation, the Laplacian pyramid, is a set of quasi-bandpass versions of the image.

In our case, each input is decomposed into a pyramid by applying Laplacian operator to different scales. Similarly, for each normalized weight map $\bar{\mathcal{W}}$ a Gaussian pyramid is computed. Considering that both the Gaussian and Laplacian pyramids have the same number of levels, the mixing between the Laplacian inputs and Gaussian normalized weights is performed at each level independently yielding finally the fused pyramid:

$$\mathcal{F}^l(x, y) = \sum_{k=1}^K G^l \{ \bar{\mathcal{W}}^k(x, y) \} L^l \{ \mathcal{I}^k(x, y) \} \quad (5)$$

where l represents the number of the pyramid levels (determined by the image dimensions), $L \{ \mathcal{I} \}$ is the Laplacian version of the input $\mathcal{I}$, and $G \{ \bar{\mathcal{W}} \}$ represents the Gaussian version of the normalized weight map $\bar{\mathcal{W}}$. This step is performed successively for each pyramid layer, in a bottom-up manner. Basically, this approach solves the *cut and paste* problem among the inputs with respect to the normalized masks. A similar adjustment of the Laplacian pyramid, but in the context of exposure fusion has been applied in other contexts [6]. The final decolorized image is obtained by summing the fused contribution of all inputs.

This linear multi-scale strategy performs relatively fast (takes approximately 1.4 seconds per image in our unoptimized MATLAB implementation) representing a good trade off between speed and accuracy. By employing independently a fusion process at every scale level the potential artifacts due to the sharp transitions of the weight maps are minimized. Multi-scale techniques are broadly used due to their efficiency in image compression, analysis and manipulation. This operation has the advantage that it respects the perceptual system of the human eye, which is known to be more sensitive to modifications into high frequencies than changes in low frequencies.




Fig. 5. Coherence of color-to-gray methods. Note how differently the methods map the background (e.g. leaves) and the flower. Compared with the methods of [8, 13] our operator maps into the same grayscale level the leaves while the flower is converted into different grayscale levels.

4 Results and Discussion

Our fusion-based approach addresses the preservation of several important image features: saliency, well-exposedness and chromatic contrast. One major benefit of fusing the inputs guided by weight maps is that this principle allows for a direct transfer of the important characteristics of the color image to the decolorized version. We believe that strong perceptual similarity between colorized and decolorized images can be obtained by algorithms that consider both global and local impressions. In our approach, the global appearance of the image is preserved by imposing a gray-shades order that respect the $H - K$ color appearance model. The weight maps contribute to the local preservation of the original relations between neighbor patches. A similar idea has been experimented with, using Poisson solvers [27] in a related approach of Gooch et al. [9]. However, their approach performs poorly for images with extended disconnected regions that represent isoluminant features. The main reason for this is that the Poisson solver ignores differences in gradients over distances larger than one pixels. Our fusion technique proves that by employing well defined quality measures and inputs, consistent results can be produced even for these difficult cases.

The new operator has been tested extensively for a large set of images. Figure 7 presents several comparative results against recent grayscale operators (for additional results the reader is referred to the supplementary material).

4.1 Evaluation of Grayscale Operators

In order to measure the quality of the conversions, we performed a contrast-based evaluation of the recent state-of-the-art operators. For this task we adapted the recent technique of Aydin et al. [28], which is used to compare a pair of images with significantly different dynamic ranges. Instead of detecting only contrast changes, this metric is sensitive to three types of structural changes: *loss of visible contrast* (**green**) - a contrast that was visible in the reference image becomes invisible in the transformed version, *amplification of invisible contrast* (**blue**) - a contrast that was invisible in the reference image becomes visible in the transformed version and *reversal of visible contrast* (**red**) - a contrast is visible in both images, but has different polarity. They observed that the contrast loss (green) is related with blurring, while contrast amplification (blue) and reversal (red) with sharpening. An online implementation¹ of this metric is made available by the authors.

We tested several grayscale operators for a set of 24 images that have also been used in the perceptual evaluation of Cadik [29]. Besides the *CIEY*, Bala and Eschbach [10], Gooch et al. [9], Rasche et al. [8], Grundland and Dodgson [13], Coloroid [12], Smith et al. [11] methods, we also reviewed the recent technique of Kim et al. [14] and our fusion-based decolorization operator. The measure of Aydin et al. [28] is applied using the default parameter set of the authors, and considering the original color image as a reference. The results of applying the *IQA* measure are shown in Figure 8. The graphics in the figure display the average ratio over the 24 *IQA* images of the pixels with the contrast changed after applying the corresponding transformation. Only the pixels with a probability higher than 70% have been counted.

Based on these graphics, we can observe that our operator, together with Smith et al. [11] and Kim et al. [14], shows the minimal amount of produced blurring artifacts (rendered with the *green* after applying *IQA*). Regarding the sharpening effects (*blue* and *red* pixels), in general all methods, except the one of Smith et al. [11] and *CIEY*, perform in a relative similar range of values.

4.2 Video Decolorization

Video decolorization adds an other dimension to the problem of image decolorization, as temporal coherence needs to be guaranteed for the entire video sequence. In order to speak of consistency, an algorithm has to map similar regions from the color input onto similar areas in the decolorized output. Recently Smith et al. [11] have shown that local approaches are suitable for this task. As our strategy retains both global and local characteristics, it is able to maintain consistency over varying palettes (see Figure 5), yielding temporal coherence for videos (see as well Figure 6). Figure 5 shows several versions of the same image, in which the flower is colored differently at each instance. Global pallet mapping techniques like the one of Grundland and Dodgson [13] generate dissimilar gray

¹ http://www.mpi-inf.mpg.de/resources/hdr/vis_metric/




Fig. 6. Isoluminant video test. For a well consistency the same color patch needs to be converted into a similar gray level in all images. Notice the artifacts introduced by Smith et al. [11] approach but also the different grayscale mapping of the same colored patch yielded by Grundland and Dodgson method [13] (please refer to the supplementary material for the entire sequence).

levels for the same region on different instances (note the leaves and the background mapping). Our operator and the method of Smith et al. [11] yield more consistent outputs.

A similar limitation can be observed by analyzing Figure 6 that displays the frames of a synthetically generated footage with isoluminant color patches. By a close inspection it might be seen that even though the technique of Grundland and Dodgson [13] decolorizes each frame perceptually accurate, this technique is not able to preserve the same grayscale level corresponding to the same color patch along the entire sequence of the frames. On the other hand, the method of Smith et al. [11] introduces some non-homogeneity artifacts along edges (please refer to the supplementary material for the comparative videos).

In our extensive experiments, the fusion-based operator performed generally well. However, we observed that in some situation due to the inexact selection of the saliency map, our technique is unable to improve substantially the results yielded by standard conversion. In general, the chosen weights can yield proper results. As well, the algorithm shares a weight map related problem in common with other fusion algorithms. As noticed by previous methods [7], the exposedness map may generate an artificial appearance of the image when its gain is exaggerated.

Our algorithm is computationally effective (our unoptimized implementation takes approximately 2.5 seconds for a 800x600 image), having a processing time comparable to recent CPU approaches (e.g. Smith et al. [11] method takes 6.7 - 10.8 seconds for 570x593 image, Decolorize [13] -unoptimized code - 3.5 seconds for a 800x600 image and the (extremely) optimized code of Kim et al. [14]




Fig. 7. Comparative results. From left to right the grayscale results obtained by applying *CIE Y*, Bala and Eschbach [10], Gooch et al. [9], Rasche et al. [8], Grundland and Dodgson [13], Coloroid [12], Smith et al. [11], Kim et al. [14] methods and our fusion-based operator.

decolorizes a 800x600 image in 1-2 seconds). However, even relatively fast since it employs an effective nonlinear global mapping optimization, the method of Kim et al. [14] did not solve the rendering limitations of the related technique of Gooch et al [9], tending to diminish the global contrast and to loose the original saliency (please refer to their image results and the rose video in the additional material). In addition, we believe that an optimized CPU implementation would make our operator suitable for real-time applications.

5 Conclusion

In this paper we have introduced a new color-to-grayscale conversion strategy, in which we employ a multi-scale fusion algorithm. We have shown that by choosing the appropriate weight maps and inputs, an image fusion strategy can be used to effectively decolorize images. We performed an extensive evaluation against the recent decolorization operators. Moreover, our operator is able to transform color videos into a decolorized version that preserve the original saliency and appearance in a consistent manner.




Fig. 8. IQA evaluation of operators. In the top are shown the results obtained by applying the contrast-based measure of Aydin et al. [28] between the original color and decolorized images. In the bottom part are displayed the graphics that plot the average ratio (over the complete set of images) of the pixels with the contrast changed after applying the corresponding transformation. Note that the *green* is related with blurring while *blue* and *red* are related with sharpening.

To future work, we would like to experiment the potential of our operator for several other applications but also to perform a perceptual evaluation of the recent color-to-grayscale techniques.

References

1. Ancuti, C.O., Ancuti, C., Bekaert, P.: Effective single image dehazing by fusion. IEEE Int. Conf. on Image Processing (ICIP) (2010)
2. Agarwala, A., Dontcheva, M., Agrawala, M., Drucker, S.M., Colburn, A., Curless, B., Salesin, D., Cohen, M.F.: Interactive digital photomontage. ACM Trans. Graph (SIGGRAPH) (2004)
3. Perez, P., Gangnet, M., Blake, A.: Poisson image editing. ACM Trans. Graph (SIGGRAPH) (2003)
4. R., R.B.: The art and science of digital compositing. Morgan Kaufmann (1999)
5. Grundland, M., Vohra, R., Williams, G.P., Dodgson, N.A.: Cross dissolve without cross fade : Preserving contrast, color and salience in image compositing. Computer Graphics Forum (EUROGRAPHICS) (2006)

6. Burt, P.J., Hanna, K., Kolczynski, R.J.: Enhanced image capture through fusion. In IEEE Int. Conf. on Computer Vision (1993)
7. Mertens, T., Kautz, J., Reeth, F.V.: Exposure fusion: A simple and practical alternative to high dynamic range photography. *Computer Graphics Forum* (2009)
8. Rasche, K., Geist, R., Westall, J.: Re-coloring images for gamuts of lower dimension. *Computer Graphics Forum (EUROGRAPHICS)* (2005)
9. Gooch, A.A., Olsen, S.C., Tumblin, J., Gooch, B.: Color2gray: salience-preserving color removal. *SIGGRAPH, ACM Trans. Graph.* **24** (2005) 634–639
10. Bala, R., Eschbach, R.: Spatial color-to-grayscale transform preserving chrominance edge information. In: *Color Imaging Conf.* (2004)
11. Smith, K., Landes, P.E., Thollot, J., Myszkowski, K.: Apparent greyscale: A simple and fast conversion to perceptually accurate images and video. *Computer Graphics Forum* (2008)
12. Neumann, L., Cadik, M., Nemcsics, A.: An efficient perception-based adaptive color to gray transformation. In: *Proc. of Comput. Aesthetics.* (2007) 73–80
13. Grundland, M., Dodgson, N.A.: Decolorize: Fast, contrast enhancing, color to grayscale conversion. *Pattern Recognition* **40** (2007)
14. Kim, Y., Jang, C., Demouth, J., Lee, S.: Robust color-to-gray via nonlinear global mapping. *ACM Trans. Graph (SIGGRAPH ASIA)* (2009)
15. de Queiroz, R.L., Braun, K.: Color to gray and back: color embedding into textured gray images. *IEEE Trans. on Image Processing* **15** (2006) 1464–1470
16. Nemcsics, A.: Color space of the coloroid color system. In: *Color Research and Applications. Volume 12.* (1987)
17. Alsam, A., Drew, M.S.: Fast multispectral2gray. In *Journal of Imaging Science and Technology* (2009)
18. Socolinsky, D., Wolff, L.: Multispectral image visualization through first-order fusion. *IEEE Transactions on Image Processing* (2002)
19. Fairchild, M., Pirrotta, E.: Predicting the lightness of chromatic object colors using cielab. *Color Research and Application* (1991)
20. Nayatani, Y.: Relations between the two kinds of representation methods in the helmholtz-kohlrausch effect. *Color Research and Application* (1998)
21. Achanta, R., Hemami, S., Estrada, F., Ssstrunk, S.: Frequency-tuned salient region detection. In *IEEE CVPR* (2009)
22. Burt, P., Adelson, T.: The laplacian pyramid as a compact image code. *IEEE Transactions on Communication* (1983)
23. Rahman, Z., Woodell, G.: A multi-scale retinex for bridging the gap between color images and the human observation of scenes. In *IEEE Trans. on Image Proc.* (1997)
24. Durand, F., Dorsey, J.: Fast bilateral filtering for the display of high-dynamic-range images. *ACM Trans. Graph (SIGGRAPH)* (2002)
25. Farbman, Z., Fattal, R., Lischinski, D., Szelinski, R.: Edge-preserving decompositions for multi-scale tone and detail manipulation. *ACM Trans. Graph (SIGGRAPH)* (2008)
26. Subr, K., Soler, C., Durand, F.: Edge-preserving multiscale image decomposition based on local extrema. *ACM Trans. Graph (SIGGRAPH ASIA)* (2009)
27. Fattal, R., Lischinski, D., Werman, M.: Gradient domain high dynamic range compression. *ACM Trans. Graph (SIGGRAPH)* (2002)
28. Aydin, T., Mantiuk, R., Myszkowski, K., Seidel, H.S.: Dynamic range independent image quality assessment. *ACM Trans. Graph (SIGGRAPH)* (2008)
29. Cadik, M.: Perceptual evaluation of color-to-grayscale image conversions. *Computer Graphics Forum* **27** (2008)