Color Constancy for Multiple-Illuminant Scenes using Retinex and SVR

Weihua Xiong and Brian Funt; School of Computing Science Simon Fraser University; Burnaby, Canada

Abstract

Scenes lit by multiple colors of illumination provide a problem for color constancy and automatic white balancing algorithms. Many of these algorithms estimate a single illuminant color, but since when there are multiple illuminants, there is in fact not a single correct answer. For automatic white balancing and color-cast removal in digital images, multiple illuminants mean that a single, image-wide adjustment of colors may not yield a good result, since the adjustment that makes one image area look better, may simultaneously make another look worse. Retinex is one method that adjusts colors on a pixel-by-pixel basis, and so inherently addresses the multiple-illumination problem, but it does not always produce a perfect overall color balance. On the other hand, illumination estimation by Support Vector Regression (SVR), produces quite good overall color balance for single-illuminant scenes, but does not adjust the colors locally. By combining Retinex and SVR in to a hybrid Retinex+SVR method, some of these problems can be overcome. Experiments with both synthetic and real images show promising results.

Introduction

Many scenes involve multiple sources of illumination. One very common example occurs when one is indoors and looks across the room and through a window to the outdoors. The indoor illumination will generally be reddish in comparison to the bluish illumination provided by the sky. These situations can lead to very strange looking digital photographs. If the camera is correctly balanced for the indoor illumination, the window will often look far too blue. The problem is that a single color balance setting is insufficient. The colors must, to some extent at least, be adjusted locally to account for the local variation in scene illumination.

The majority of the illumination-estimation methods [1-5] that have been developed for automatically color balancing images make a single estimate of the scene illumination. They, therefore, are susceptible to the situations such as the too blue window. Retinex is an exception in that it makes a separate illumination estimate for each pixel. Although this is a strength of Retinex, it can also be a weakness in that the illumination estimate is strongly influenced by the colors in each pixel's neighborhood.

Our goal is to gain the benefits of both the local and global approaches by merging them into a single process. In particular, we use Retinex to make local color adjustments and then apply the Support Vector Regression (SVR) method to the Retinex-processed image to adjust the overall color balance. In scenes with strong differences of illumination, our hypothesis is that because it makes local adjustments, Retinex will attenuate the differences in illumination, and as result SVR will be able to make a better global adjustment for the illumination. The experiments describe below show that this hybrid method works better than either SVR or Retinex alone.

Overview of Retinex and SVR

In the current context, the most important feature of Retinex is that it estimates the illumination color (or equivalently the surface color) locally at each pixel by making comparisons between the pixel and other image pixels. Comparisons to nearby pixels are given more weight than distant pixels, but every pixel effects the results at every other pixel to some extent. Retinex processing is carried out on each color channel independently. The McCann99 [6] version of retinex is based on a multi-resolution image pyramid. McCann99 initially makes spatial comparisons at the coarsest resolution and then propagates the results down to the next higher resolution level of the pyramid. Since results at the coarsest resolution eventually propagate down to influence the finest resolution, distant pixels influence each other to some extent. Another important aspect of Retinex is that it includes a 'reset' operation. The reset means that locally the most reflective surface is assumed to be 'white'. This assumption has the effect that other pixels' colors are adjusted relative to this white.

Support Vector Regression is machine learning technique which is an extension of support vector classification [13]. The use of the term 'regression' in SVR can be thought of by analogy to linear least-squares regression; however, SVR is not based on a least-square error measure and it works for non-linear functions. SVR is a learning technique in that from a set of training data, it determines a function that interpolates the data. SVR has been used to estimate [2] by separately interpolating the r and g chromaticities as functions of the image's binarized color histogram. Other histogram-based illumination-estimation using neural networks [1], correlation matrices [3], and KL-divergence [5], although quite different in their specific details, are similar in their overall approach and yield similar results. Potentially any one of them could be substituted for SVR in what follows.

Implementation Details

For the Retinex implementation, we use the Matlab version of McCann99 Retinex [6]. For Support Vector Regression we use the "3D" method described in [2] which is based on binarized histograms of the image pixels' (L, r, g) *where* $L = R + G + B$ and $r=R/L$ and $g = G/L$. . We quantize L into 25 equal steps, and r and g into 50 steps so the 3D histograms consist of 62,500 (25x50x50) bins. After training, SVR provides an estimate of the *rg*-chromaticity of the overall scene illumination based on the binarized image histogram submitted to it.

SVR requires a training set. We created a training set of 56,730 histograms by random subsampling of colors from images contained in the 11,346 "grayball" image database [7].

Each image is processed first with McCann99 Retinex. The binarized Lrg color histogram of the resulting image is then passed to SVR which returns the estimate of the illumination chromaticity. The SVR estimate is not actually an estimate of the true illumination, but rather an estimate of the illumination relative to the post-Retinex-processed image. The SVR illumination estimate is used in a diagonal von Kries transformation to correct the post-Retinex image in order to adjust it to have the colors it would have had if the original scene had been imaged under the canonical illumination. This Retinex-SVR image is then compared with the ground-truth image of the same scene imaged under the canonical white illumination

We evaluate Retinex-SVR performance at each pixel in terms of the distance between measured in *rg*-chromaticity (r=R/(R+G+B), g=G/(R+G+B)) space and in terms of the angle in degrees between colors in RGB space. These errors are defined by the following formulas, where subscript 'p' indicates the result after Retinex-SVR and 'g' indicates the ground-truth image.

$$Ed = \sqrt{(r_p - r_g)^2 + (g_p - g_g)^2} \quad \textbf{(1)}$$

$$Ea = \cos^{-1}\left[\frac{(r_p, g_p, b_p) \circ (r_g, g_g, b_g)}{\sqrt{r_p^2 + g_p^2 + b_p^2} * \sqrt{r_g^2 + g_g^2 + b_g^2}}\right] \quad \textbf{(2)}$$

We also compute three statistics on the distribution of errors across all the pixels in an image: the median, the RMS (root mean square) and the mean of the top 1/2 percentile of the largest errors, denoted MMax. In contrast to a single maximum error, MMax is a more representative measure of the methods failures. RMS of the errors from N pixels is given by the standard formula:

$$RMS = \frac{1}{N}\sqrt{\sum_{i=1}^{N} E_i} \quad \textbf{(3)}$$

The Wilcoxon signed-rank test with a 0.01 threshold for accepting or rejecting the null hypothesis is also used to evaluate difference between error distributions [8].

Synthetic Image Experiments

Our first experiments are based on synthetic images that model a scene with two quite distinct illuminants lighting different parts of the scene. We generate synthetic scenes composed of patches of different reflectance by randomly selecting reflectances from the 1995 available in the database described by Barnard [9]. The patches are divided into two sections by an irregular boundary representing where the illumination changes. RGB values for the patches are calculate by using two illumination spectra, CIE A on the left, CIE D65 daylight on the right, and sensor sensitivity functions of the SONY DXC-930 camera color balanced equal-energy white. The ground-truth image is generated using equal-energy white illumination over the whole scene. The sensitivity functions were normalized for this white light. All of the spectra and sensitivity functions were downloaded from the Simon Fraser University color database [10].

Figure 1 shows the results of SVR, Retinex and Retinex+SVR processing. The top left Mondrian is the input image with a white line superimposed demarcating the boundary between the two illuminations to make it easier to see. The line is not part of the actual input image. SVR applied to the input Mondrian estimates the illumination's rgb-chromaticity as [0.375, 0.298, 0.308], in other words, as quite reddish in comparison to white [0.333, 0.333, 0.333]. This successfully removes some of the reddish cast from the left side of the image, but introduces more blue to the right side (Figure 1, bottom row on the left). On the other hand, when SVR is applied to the Retinex-processed image (Figure 1, middle row on the left), it estimates the "illumination" as a bluish [0.296, 0.315, 0.389]. In this second case, there was no actual illumination; rather it is SVR's estimate of what the illumination would be if the Retinex output were actually an unprocessed input image. Since the Retinex result is too blue in comparison to the ground-truth Mondrian (top right), correcting the colors based on SVR's estimate improves the image so that now the bottom right (Retinex+SVR) and top right (ground truth) images are very similar. Numerical results are tabulated in Tables 1 and 2. The Wilcoxon signed-rank test applied to the angular error indicates that for this image the performance difference is significant and that Retinex+SVR outperforms Retinex, and Retinex outperforms SVR.

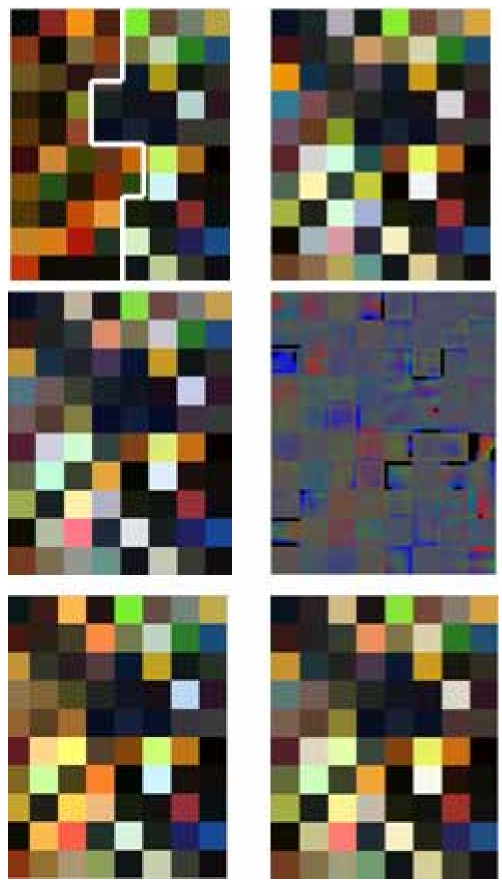

Figure 1 Synthetic image results. Top left: input image with a white line superimposed to indicate the illumination boundary. Top right: ground-truth image under equal energy white light. Middle left: Retinex result. Middle right: Retinex illumination map; Bottom left: SVR result. Bottom right: Retinex+SVR result.

	Distance (* 10^2)			Angular		
	MMax	RMS	Med	MMax	RMS	Med
Retinex+ SVR	15.19	5.15	2.28	15.74	6.27	3.23
Retinex	19.59	7.58	3.28	18.41	7.86	4.43
SVR	30.44	10.97	4.66	33.15	13.08	5.96

Table 1 Comparison of MMax (see text for definition), RMS and median error on a per-pixel basis between the ground-truth image values and the processed image values for processing by Retinex+SVR, McCann99 Retinex alone, and SVR alone.

	Retinex+SVR	Retinex	SVR
Retinex+SVR		+	+
Retinex	-		+
SVR	-	-	

Table 2 Comparison of the different methods via the Wilcoxon signed-rank test with 0.01 as the threshold applied to the angular errors. A "+" means the algorithm listed in the corresponding row is better than the one in the corresponding column. A "-" indicates the opposite.

Real Image Experiments

The first set of real-image experiments is based on some real scenes we constructed in the lab containing two distinct illuminants similar to those found indoors and outdoors. A bluish illuminant was created by placing a light blue filter in front of a tungsten lamp. The reddish illuminant was a Solux 4100K tungsten bulb connected to a dimmer. By adjusting the dimmer, the color temperature of the light drops significantly. These scenes were photographed using a Sony DSC V1 camera. To obtain the ground-truth image, a white reflectance standard was introduced at the side of the scene, and an additional image was taken under unfiltered tungsten light. The RGB channels were then scaled in order to make the reflectance standard perfectly white (i.e., R=G=B=255).

The first test scene is shown in Figure 2a. It contains some books, boxes, and a Mini Macbeth ColorChecker and is lit with reddish light from the left and bluish light from the right. Figure 2b shows the same scene imaged under white light. In addition to using white light, the resulting image was further white balanced by scaling the RGB channels so that the image of a calibrated white reflectance results in R=G=B.

Figure 2c shows the retinex result with the intensity at each pixel adjusted to match the input image in Figure 2a. Although retinex processing affects the luminance as well as the chromaticity of each pixel, here we are interested only in its effect on chromaticity and are restoring the luminance (R+G+B) to match that of the input image. The SVR result, which is also adjusted to preserve pixel luminance, is shown in Figure 2d. Since SVR makes the same color adjustment across the whole image, anything it does must inevitably be a compromise. In this case, SVR has removed some of the blue cast from the input image, but this introduces some orange cast in other parts of the image. On the other hand, the Retinex+SVR result shown in Figure 2e contains neither a blue nor an orange cast. SVR determined the single value for the illumination in rgb-chromaticity as a slightly bluish [0.306, 0.308, 0.385] in comparison to white [0.333, 0.333, 0.333]. When applied to the Retinex-processed image, SVR's estimate is [0.324, 0.341, 0.327].

The numerical results presented in Tables 3 and 4 show that retinex and SVR perform with similar accuracy for this image, while the Retinex+SVR hybrid outperforms each of the others taken individually.

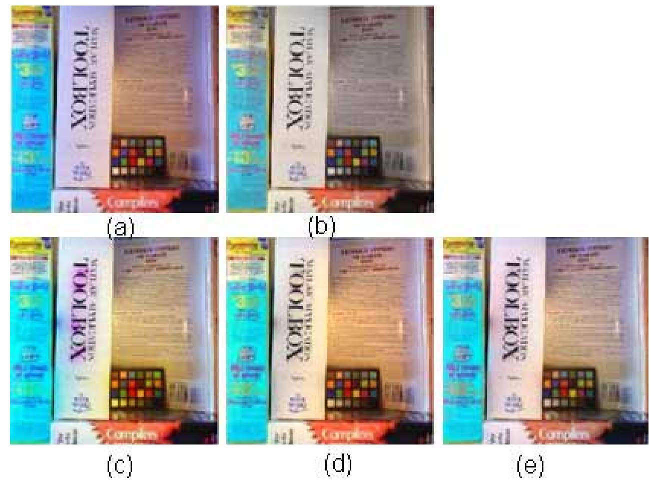

Figure 2 Two-illuminant books scene: (a) input image with reddish light coming from the left and bluish from the right; (b) ground-truth image captured under white light matching the camera's white point; (c) Retinex result (d) SVR result (e) Retinex+SVR result.

	Distance (* 10^2)			Angular		
	MMax	RMS	Med	MMax	RMS	Med
Retinex+ SVR	59.13	12.27	5.84	49.53	12.66	6.46
Retinex	60.72	14.77	9.73	15.07	13.27	12.29
SVR	59.21	14.58	8.38	53.62	16.71	10.18

Table 3 Comparison for the two-illuminant books scene of MMax (see text for definition), RMS and median errors measured on a pixel-by-pixel basis between the ground-truth image values and the processed image values for processing by Retinex+SVR, Retinex alone, and SVR alone

	Retinex+SVR	Retinex	SVR
Retinex+SVR		+	+
Retinex	-		=
SVR	-	=	

Table 4 Comparison of the different methods via the Wilcoxon signed-rank test for the two-illuminant books scene. A "+" means the method listed in the corresponding row is better than the one in the corresponding column; a "-" indicates the opposite; and a "=" indicates they are indistinguishable.

We designed a second scene in the lab intended in this case to model the situation of being indoors in a room with a window to the outdoors. The scene shown in Figure 3a consists of a toy human figure 'outdoors' seen through a window. The mountain scene on the left is a picture on the wall 'indoors'. The colored ball is also indoors. The outdoor objects are lit with sky blue light, while the indoor ones are lit by reddish-orange light. Figure 3b is the ground truth image with pixel intensities adjusted to match those of the input image. The Retinex result in Figure 3c shows that Retinex reduces the magnitude of the difference between the two illuminants, but the overall color balance is to yellow. SVR determines the single value for the illumination in *rgb*-chromaticity as a slightly reddish [0.343, 0.335, 0.322].

On the other hand, when SVR is applied to the Retinex-processed image, SVR's estimate is [0.346, 0.358, 0.297]. SVR provides better overall color balance in Figure 3d, but the outdoor part becomes even bluer. The Retinex+SVR result, Figure 3e, has the indoor section reasonably well balanced and has reduced, but not eliminated the outdoor blue. Numerical results are presented in Tables 5 and 6.

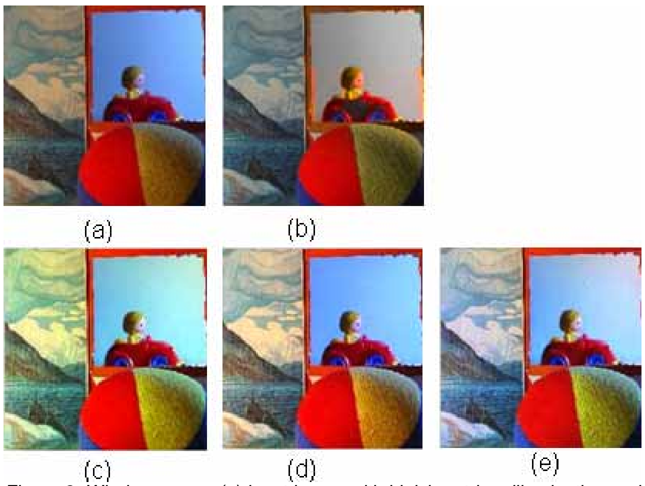

Figure 3 Window scene: (a) input image with bluish outdoor illumination and red-orange indoor illumination. (a) input image (b) ground-truth image captured under white light that matches the camera's white point; (c) Retinex result (d) SVR result (e) Retinex+SVR result

	Distance (* 10^2)			Angular		
	MMax	RMS	Med	MMax	RMS	Med
Retinex+ SVR	36.74	7.71	3.55	42.24	8.99	5.04
Retinex	50.05	11.98	4.85	56.33	13.38	8.39
SVR	40.17	9.76	6.26	43.26	10.94	7.83

Table 5 Comparison of MMax (see text for definition), RMS and median errors measured on a pixel-by-pixel basis between the ground-truth image values and the processed image values for processing by Retinex+SVR, Retinex alone, and SVR alone.

	Retinex+SVR	Retinex	SVR
Retinex+SVR		+	+
Retinex	-		-
SVR	-	+	

Table 6 Comparison of the different methods via the Wilcoxon signed-rank test for the window scene. A "+" means the method listed in the corresponding row is better than the one in the corresponding column. A "-" indicates the opposite.

In addition to laboratory scenes, we processed images of other typical scenes. The advantage of the laboratory scenes is that it is possible to obtain a ground truth image with which to evaluate the error in illumination estimation. Outside the laboratory, it is difficult to make enough measurements of the illumination distribution to obtain the ground truth image. During a subjective evaluation of several hundred images, we found that in many cases there is little difference in the overall image quality between Retinex, SVR and Retinex+SVR. This is in part because the majority of scenes do not contain dramatic differences in incident illumination. However, in the cases where the scene clearly contains quite different illuminants, Retinex+SVR is superior. An example of one such scene and the results of the three methods is shown in Figure 4. In this example, Retinex has again reduced the difference in illumination, but has left the image with a slight blue cast that Retinex+SVR removes. Although our goal has been to remove the color shifts created by multiple colors of illumination, success in reaching that goal does not guarantee that this will lead to preferred image renderings. However, being able to compute the illumination distribution should prove useful when calculating a preferred rendering.

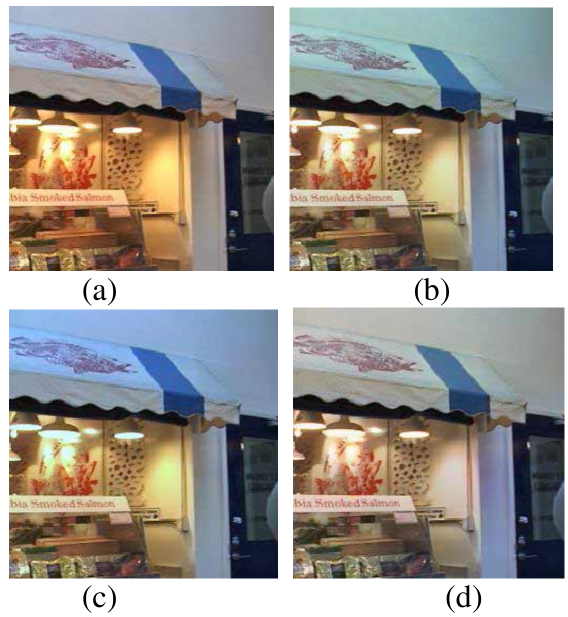

Figure 4 Typical natural image with two illuminations: (a) input image; (b) Retinex result; (c) SVR result; (d) Retinex+SVR result

Retinex Iteration Time

McCann99 Retinex is a multi-resolution algorithm and one of its key parameters [11] is the number of iterations it performs at each resolution. We determined the optimal setting for Retinex+SVR by plotting the error as a function of the number of iterations. Figure 5 shows the plot for the case of the two-illuminant window scene. The plots for other scenes were similar with the minimum error found at 4 iterations. All our experiments were thus based on 4 iterations.

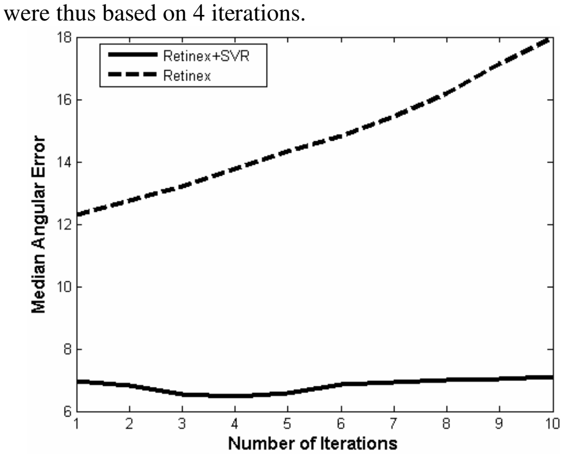

Figure 5 Median angular error as a function of the number of iterations Retinex used at each resolution. This plot is for the two-illuminant window scene; however, for other scenes the results are qualitatively similar.

Conclusion

Many images are of scenes with at least two distinct illuminants. These images present a challenge for automatic white balancing algorithms because there is no single right answer. Retinex is one method that makes local adjustments for the illumination, but it does not always get the overall color balance correct. We proposed a hybrid Retinex+SVR method and shown, at least for the limited set of images it is possible to create in the laboratory, that it works better than either SVR or Retinex working separately.

SVR is not the only illumination-estimation method that could be hybridized with Retinex. Since there are several learning-based illumination methods [1-5] of similar accuracy, it is reasonable to suppose that any one of them could be substituted for SVR in this context with similar results.

Our goal was to remove the color effects of illumination; however, as Hubel [12] has argued perhaps in terms of creating an interesting image it is best to preserve the illumination effects. We have not addressed the problem of preferred reproduction directly, but assume that any additional information that can be extracted from an image concerning the distribution of illumination color will at some point be helpful in creating pleasing images.

Acknowledgment

Funding was provided by the Natural Sciences and Research Council of Canada.

References

[1] V. Cardei, B. Funt, K. Barnard, "Estimating the Scene Illumination Chromaticity Using a Neural Network," *Journal of the Optical Society of America A*, Vol. 19, No.12, pp. 2374-2386, (2002)

[2] B. Funt, W.H. Xiong, "Estimating Illumination Chromaticity via Support Vector Regression," *Proc. 12th Color Imaging Conference,* pp. 47-52 , (2004)

[3] G.D. Finlayson, S.D. Hordley, P.M. Hubel, "Color by Correlation: A Simple, Unifying Framework for Color Constancy," *IEEE Trans. Pattern Anal. Mach. Intel*. Vol. 23, No.11, pp.1209-1221, (2001)

[4] D. Forsyth, "A novel algorithm for color constancy"*, Int. J. of Comp. Vis.,* Vol. 5, pp. 5-36 , (1990)

[5] C. Rosenberg, M. Hebert, and S. Thrun, "Color constancy using KL-divergence", *Proc. Eighth IEEE International Conference on Computer Vision (ICCV '01)*, Vol. 1, July, pp. 239 – 246, (2001)

[6] B. Funt, F. Ciurea., J. McCann., "Retinex in Matlab," *Journal of the Electronic Imaging*, pp.48- 57 , (2004)

[7] F. Ciurea and B. Funt, A Large Image Database for Color Constancy Research*, Proc. 11th Color Imaging Conference, Society for Imaging Science and Technology*, Nov. pp.160-163, (2003).

[8] S.D. Hordley, and G.D. Finlayson, "Re-evaluating colour constancy algorithms", *Proc. 17th ICPR*, Vol. 1, pp. 76-79, (2004)

[9] K. Barnard, L. Martin, B. Funt, and A. Coath., "A Data Set for Color Research," C*olor Research and Application*, Vol. 27, No. 3, pp.148-152 , (2002)

[10] Simon Fraser University, Computational Vision Lab Data, *www.cs.sfu.ca/~colour/data*

[11] F. Ciurea, and B. Funt, "Tuning Retinex Parameters," *Journal of Electronic Imaging* 13(1), Jan, pp. 58-64, (2004).

[12] P. M. Hubel, "The Perception of Color at Dawn and Dusk", *Proc. 7th Color Imaging Conference*, pp.48-51, (1999)

[13] A. Smola and B. Schölkopf, "A tutorial on support vector regression", *Statistics and Computing*, 2003