Pacific Graphics 2009
S. Lee, D. Lischinski, and Y. Yu
(Guest Editors)

Volume 28 (2009), Number 7

Automatic Correction of Saturated Regions in Photographs using Cross-Channel Correlation

Syed Z. Masood Jiejie Zhu Marshall F. Tappen

University of Central Florida
School of Electrical Engineering and Computer Science, Orlando, FL
{smasood, jjzhu, mtappen}@eecs.ucf.edu

Abstract

Incorrectly setting the camera's exposure can have a significant negative effect on a photograph. Over-exposing photographs causes pixels to exhibit unpleasant artifacts due to saturation of the sensor. Saturation removal typically involves user intervention to adjust the color values, which is tedious and time-consuming. This paper discusses how saturation can be automatically removed without compromising the essential details of the image. Our method is based on a smoothness prior: neighboring pixels have similar channel ratios and color values. We demonstrate that high quality saturation-free photos can be obtained from a simple but effective approach.

Categories and Subject Descriptors (according to ACM CCS): I.4.8 [Image Processing and Computer Vision]: Scene Analysis—Photometry

1. Introduction

When photographing a scene, incorrectly setting the camera's exposure can have a significant negative effect on quality of the photo. Too short of an exposure leads to an overly dark image with possibility of quantization artifacts and noticeable sensor noise. Over-exposing photographs leads to the sensor being saturated and parts of the photograph being over-exposed. While significant research has been done in removing sensor noise, a problem for under-exposed images ([APS98], [Kim99]), little help is available for over-exposed images. In this paper, we present a method for correcting saturated areas in photographs.

Saturated areas in a photograph are difficult to work with because the saturation represents a true loss of data – the textures that one would expect to see in a saturated region are replaced with a flat patch. However, as we will demonstrate, in many cases only one or two color channels in an image are saturated. Our method operates by using assumptions about the smoothness of images to estimate the correlation between the saturated and unsaturated color channels. Once the correlation between color channels has been established, it is possible to estimate the appearance of the pixel values in saturated channels.

Figure 1 demonstrates how this process works. The red channel of the pixels area on the cheek in Figure 1(a) is saturated. To recover these red-channel values, the ratio between the red and green and red and blue channels at the saturated pixels is estimated from non-saturated neighbors. For this figure, we only show the recovered red:green ratio (see Figure 1(b)). This neighboring pixel information is weighted based on how close they are with respect to the unsaturated channel of the saturated pixel. Once these ratios have been estimated, both the green and blue channels are used to predict the pixel values where the red channel is saturated.

As we will show, removing the saturation can be accomplished using straightforward quadratic models similar to those used for colorization [LLW04], denoising [TLAF07], and image upsampling [Fat07]. One of the advantages of this approach is that the system can perform robustly in real-world conditions. In fact, the images in Figures 1, 4, 8 and 9 were taken from photographs that were captured with no knowledge of this project. As part of this demonstration, our results are focused on images of faces. We argue that correcting photographs of faces are the most important application of our saturation removal techniques. While it is possible to adjust and optimize, through multiple exposures, the exposure settings for a fixed landscape scene, many photographs of humans are unique moments that cannot be recreated.

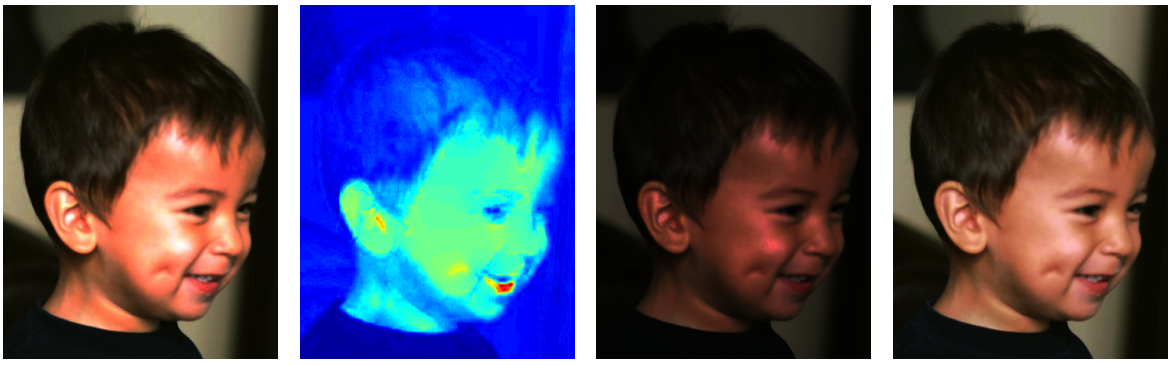

(a) Saturated image (b) Recovered Red:Green ratio (c) Saturation-free result (d) Visual result

Figure 1: *Process of our pixel saturation removal approach. The saturation-free image appears darker as the entire image has been linearly scaled to accommodate the wider range of pixel values necessary to represent the regions where the saturation has been corrected. A pixel-wise nonlinear adjustment to (c), similar to gamma correction, improves visual quality of the final image.*

Thus, the ability to correct images in these unique situations is most vital. Having said this, we show that our method also works for saturated non-face images.

Section 2 discusses related work in this area of research. Section 3 explains our methodology in detail. Section 4 shows our results and experiments conducted and discusses as to how our method is an improvement over other available techniques. It also details the issues we encountered and how they were addressed. Finally, Section 5 provides a brief summary of the paper.

2. Previous Work

Our work is most related to the previous work of Zhang *et al.* [ZB04] and Wang *et al.* [WWZ*07].

In Zhang's saturation removal process, a single, Bayesian is constructed to estimate the correct values of the saturated pixels. This model captures the correlation between color channels, similar to our approach. Unlike our approach, a single model is used for all pixels. This limits the performance of the system because the ratio between color channels varies across the image. This can be seen in Figure 2(a), which shows the ratio between the red and green channels. In this image it can be readily seen that the ratio is not constant, where most of the ratios are variant. Our approach of using smoothness prior to estimate a varying ratios between color channels enables us to remove much more significant saturation effects than those removed in [ZB04].

Wang's approach, which is interactive rather than automatic, is to transfer texture details for over-exposed and under-exposed regions by manually selecting patches. Although promising results are presented, the approach requires that either the image contains repeated texture patterns or additional texture samples under different illuminations. Given the smoothness constraint, our approach can transfer texture details from neighboring pixels, which does not require any user interaction.

Nayar *et al.* [NM00] used a series of images with varying exposures to compute a *High Dynamic Range* image that is free of any saturation. Although this approach satisfies the need for pixel saturation removal, it requires several exposured images of the scene and thus, unlike our method, cannot recover a saturation-free image from only one single saturated image.

The problem of removing pixel saturation is similar to demosaicing in the sense that both rely on color channel relationships. The idea of cross-channel correlation over a local neighborhood has long been used for color reconstruction and color enhancement. Kimmel [Kim99] and Muresan *et al.* [MLP00] discussed the relationship of different color channels and how locally constant color channel ratios can be used to reconstruct color images. Local constant can be viewed as a special case of our smoothness constraint with averaging weights. Duin [Dui96] and Adams *et al.* [APS98] emphasized the importance of using spatial relationship to interpolate colors and recognize them. Tai *et al.* [TWCL08] and Meylan *et al.* [MS04] employed Gaussian filters and logarithmic curves of neighboring pixels respectively to enhance contrast in an image. In addition, Huang *et al.* [HM99] and Weiss *et al.* [WF07] reported statistical models of the correlation in a local region from natural images.

The work of Nachlieli *et al.* [HNS09] and Kryszczuk *et al.* [KK03] is related in that it focuses on automatic color correction for face images corresponding to human visual preferences, though this work focuses on the case where all of the color channels have been observed. Our goal is to improve the quality of image where color information has been lost due to saturation.

Our approach also takes advantages of cross-channel correlation and regards it as a smoothness constraint. We estimate the true values of saturated pixels by deriving two cost functions which force smoothness for color channel ratios and color values. The key success of applying this smoothness constraint is that in many cases only one or two color channels in an image are saturated, therefore, our approach can propagate neighboring information correctly to saturated regions.

3. Algorithm

We work in the *RGB* color space and assume pixel values (in R, G, B, respectively) larger than 235 (maximum allowed value is 255) are saturated. This is because, in general, a signal level of digital camera output is considered to be saturated at 235 in a full resolution image [Yoh95] as the linear response is destroyed beyond this limit.

For each pixel, we label saturated channels using this threshold. This results in three types of pixels as the input to our algorithm: nonsaturated, partially saturated (saturated in some but not all channels) and totally saturated (saturated in all channels) pixels. To recover the true values of saturated channels, we first recover the correct color ratios for each pixel. These ratios are then used in calculating true values for saturated pixels. Color ratios and the true values are estimated by minimizing two separate cost functions, each utilizing neighboring information.

In our experiments, we use raw images taken by a CCD camera as the input image. For postprocessing rendered images (images transformed with non-linear camera response curves), we provide a practical approach to first linearize these images using Radial Basis Functions [Bis95], and then generate saturation-free images using the proposed method.

3.1. Estimating Color Ratio

The smoothness constraint in estimating color ratios tells us that the ratios between color channels at a pixel are close to those observed at its neighbors. Figure 2 shows this strong correlation between color channel ratios of the original image with that of pixel-shifted images. The pixel shifts are of magnitude $1, 3$ and 5 along the vertical, horizontal and diagonal directions respectively.

We take advantage of this correlation to estimate the ratio between different color channels in areas where one of the channels has been saturated. Using a quadratic cost-function the ratios in unsaturated areas can be propagated into saturated areas. Using $\alpha_p = R_p/G_p$ (red:green ratio at pixel p) as an example, the ratio is estimated by finding the vector α that minimizes

$$\sum_p w_p^c(\alpha_p - \alpha_p^O)^2 + \left(\alpha_p - \sum_{q \in N(p)} w_{pq}\alpha_q\right)^2. \qquad (1)$$

(a) Original image

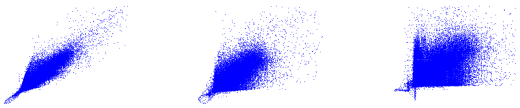

(b) 1-pixel vertical shift (c) 3-pixel horizontal shift (d) 5-pixel diagonal shift

Figure 2: *Example of correlation of Red:Green ratios between the original image and different pixel-shifted images. Image shifts are of magnitude* $1, 3$ *and* 5 *along the vertical, horizontal and diagonal directions respectively. Y axis denotes ratios from original image. X axis denotes ratios from shifted images.*

Here the sum is over all pixels p in the image. In the inside summation, $q \in N(p)$ denotes that q is a neighbor of p. The first term in the sum constraint constrains α_p to match the observed ratio α_p^O if the pixel is not saturated. For unsaturated pixels, w_p^c will be a large value, and will be 0 for saturated pixels. This technique was used for convenience, though Dirichlet boundary conditions could also be used, similar to [PGB03]. Other ratios ($\beta_p = R_p/B_p, \gamma_p = G_p/B_p$) can be calculated similarly.

The function w_{pq} is the weighting function that is designed to not propagate color ratios across edges in the image, similar to [LLW04]. This weighting function describes how similar the estimated color ratios should be at pixels p and q:

$$w_{pq} = \frac{1}{k}\sum_i (\exp(\frac{-(Y_i(p) - Y_i(q))^2}{2\sigma_{p_i}^2})) \qquad (2)$$

In this notation, the summation index represents summation across all color channels that are unsaturated. The variables $Y_i(p)$ and $Y_i(q)$ denote unsaturated values of channel i at pixel p and from a 3×3 neighborhood of p respectively. The quantity σ_{p_i} is the variance of the pixel values in a 3×3 neighborhood of p. This variance is computed in channel i and is only computed using unsaturated pixels. When computing the variance, saturated pixels are ignored. The vari-

able k is the total number of unsaturated channels at pixel p.

This weighting function is designed to induce a strong correlation between color ratios if two pixels have a similar appearance in the unsaturated color channels. In the case that all three color channels are saturated, w_{pq} is set to a very small value to keep the system well-defined. As Section 3.2 will show, the color ratio estimates from pixels where all channels are saturated will be ignored.

3.2. Estimating Color Values

Having estimated the correct color ratios for each channel, we estimate the true image intensities by minimizing a cost function. Using pixels in the red channel as an example, we denote our estimate of the unsaturated values as $\mathbf{R}^*$. The estimate is produced by this minimization:

$$\mathbf{R}^* = \arg\min_{\mathbf{R}} \sum_p \Big((w_p^g (R_p - \alpha_p G_p)^2 + w_p^b (R_p - \beta_p B_p)^2 + w_p^n (R_p - R_p^O)^2) + w_p^s \sum_{q \in N(p)} (R_p - R_q)^2 \Big), \tag{3}$$

where r_p denotes the value of $\mathbf{R}$ at pixel p. The variable w_p^g is set to 1 if R_p (Red channel of pixel p) is saturated and G_p is not saturated. The weight w_p^b is similarly assigned if R_p is saturated and B_p is not saturated. If the red channel of p is not saturated, we set w_p^n to 1, which makes the value close to its original value R_p^O. The pixel values in other color channels, $\mathbf{G}$ and $\mathbf{B}$ are computed in a similar fashion.

The final term in Equation (3) is for pixels where all three channels are saturated. We have observed that even when one channel is not saturated in most of the image, all three color channels will be saturated for a small number of pixels. In this case, we assume that pixel intensities should be locally smoothed. By setting the weight w_p^s to 1 only if all channels are saturated, the final term in (3) implements this local smoothness assumption. As in Equation (1), this summation is over all pixels q that are neighbors to pixel p.

Equation 3 can be viewed as a classical energy minimization function which is composed of a data term (the first three terms) and a smoothness term (the last term). Because the cost function in Equation 3 is quadratic, it can be optimized in closed form.

The optimization is performed by first expressing Equation (3) in matrix form:

$$\mathbf{R}^* = \arg\min_{\mathbf{R}} \; (\mathbf{R} - \mathbf{h}^g)^T W^g (\mathbf{R} - \mathbf{h}^g) + (\mathbf{R} - \mathbf{h}^b)^T W^b (\mathbf{R} - \mathbf{h}^b) + (\mathbf{R} - \mathbf{R}^O)^T W^n (\mathbf{R} - \mathbf{R}^O) + (F\mathbf{R})^T W^s (F\mathbf{R}). \tag{4}$$

In this formulation, the term $\mathbf{h}^g$ is a vector such that the pth pixel in $\mathbf{h}^g$ is equal to $\alpha_p G_p$. The other vectors $\mathbf{h}^b$ and $\mathbf{R}^O$ have a similar correspondence with Equation (3). The matrix W^g is a diagonal matrix such that the pth element along the diagonal is equal to w_p^g, from Equation (3). The matrices W^b, W^n, and W^s are constructed similarly. Finally, the matrix F computes the difference between neighboring pixels in the image. In this matrix, there are eight rows for each pixel. In each of the eight rows corresponding to a pixel p, the column corresponding to p will be equal to 1. For each pixel q that is a neighbor to p, there will be a -1 in the column corresponding to q.

Differentiating this cost function leads to a sparse set of linear constraints, similar to [LLW04]. In practice, we solve this system using solver routines in MATLAB, though other solvers can be used [LLW04].

4. Results

In this section, we will demonstrate our system both quantitatively and qualitatively. We will first show how our approach is able to accurately estimate the correct ratios between color-channels in saturated regions using an image where the ground-truth values are known. We will also show how this leads to quantitative improvements. In all our experiments, we use a 3×3 neighborhood. Next, we will demonstrate system performance on images where the ground-truth is unknown. As mentioned earlier, we have focused on images with faces as these will be the most valuable application of this technique. Finally, we will show how simpler techniques, such as pixelwise scaling or colorization [LLW04], cannot desaturate images with comparable quality as those obtained with our approach.

We use Matlab's built in least squares solver (back slash) for sparse linear systems to compute the pseudo-inverse. In the examples, the image sizes are typically 500×400, requiring around 30 seconds to process.

4.1. Quantitative Evaluation

The performance of our approach can be evaluated quantitatively using the auto-exposure bracketing feature in modern SLR cameras. Using this feature, we obtained two images of the same subject with exposures that differed by 4 stops. The brighter image, shown in Figure 3(f) is saturated on the forehead and part of the nose, while the darker image (Figure 3(e)) contained no saturated pixels.

Figure 3(a) shows the ratio between the red and green channels at each pixel in the darker image, which has no saturation, while Figure 3(b) shows the ratio between these channels in the brighter image, which has saturated pixels. Notice that on the forehead and tip of the nose that the ratios are lower in the saturated image than in the unsaturated image. This occurs because the saturated red values are not has high as they should be.

Figure 3(c) shows the ratios estimated using the technique described in Section 3.1. This is much closer to the true ratios. Using these ratios, an estimate of the unsaturated image can be computed, as described in Section 3.2. Figure

(a) Ground truth ratio (b) Saturated ratio (c) Recovered ratio (d) Error comparison

(e) Ground truth (f) Saturated (g) Recovered by our approach (h) Enlarged comparison

Figure 3: *First row shows per-pixel comparison of Red:Green ratio from ground truth, saturated and saturation-free image by our approach. Second row compares the colored images. The last figure in first row shows a numerical comparison of averaged errors of Red channel to ground truth. We can see our approach removes the saturated pixels correctly and achieve better visual results. The last figure in second row shows enlarged part of forehead. The first is the ground truth, the second is the saturated image, the last is our result. We can see our result recovers the underlying texture details well.*

3(g) shows the result of this process. This image reveals much more texture information for the saturated regions and hence is more visually pleasing than the corresponding saturated image. We can observe that our saturation-free image is much closer to the ground truth image and recovers underlying texture for the saturated regions on the forehead and nose (see Figure 3(h)). Quantitatively, the error for the red channel of the saturation-free image is around 90% correct against to the ground truth (see Figure 3(d)).

4.2. Performance on General Photographs

Figures 8 and 9 show how our approach is able to significantly enhance the visual quality of a collection of photographs taken from personal photo collections. Most of these photos were taken by individuals unfamiliar with this research project.

For the results in Figure 8, we process the entire photograph to demonstrate how our method can perform robustly on larger-scale scenes. In Figure 9, we show the results produced by our method on images focused on the face. In these images our method is able to realistically remove the effects of the over-exposure on faces.

In both of these figures, the results of our system are shown in the right-most two columns of these figures. In the next-to-last column, the recovered image is scaled so that the full range of the estimated unsaturated image will lie in the range $[0, 255]$. Because the estimated image will have intensity values beyond the range of a typical image, our method, like [WWZ*07], can be thought of as estimating a high dynamic range image from the input low dynamic range image. The last column of these figures shows the estimated unsaturated image after correcting the image with a pixelwise non-linearity, similar to gamma correction. This correction is solely to enhance the perceptual quality of the estimated HDR images and make it easier to view the quality of the results.

4.3. Comparison with Baseline Methods

In this section, we show that similar results cannot be achieved just using existing techniques, such as pixelwise non-linearities and colorization.

Figure 4(b) shows the result when the saturated image in Figure 4(a) is modified with a pixelwise non-linearity. This non-linearity chosen manually for maximum perceptual image quality. However, as can be seen in Figure 4(b), the effects of the saturation are not removed. Our result shows the saturation has been significantly in Figures 4(c) and 4(d).

(a) Input image (b) Linear adjust (c) Our result (d) Our result visualized

Figure 4: *Results of comparison with pixel-wise non-linearity. (b) is obtained by manually adjusting the input (a). Notice that the image in (b) still appears saturated. Our result in (c) is saturation-free and is best visualized after a pixelwise non-linear adjustment as shown in (d). It thus gives a a much more realistic appearance as compared to the saturated image.*

(a) Saturated input (b) Colorization [LLW04] (c) Colorization using our result

Figure 5: *(a) Original saturated image. (b) Result obtained using colorization [LLW04] directly on the saturated image. (c) Result after pixel-wise non-linear adjustment to our saturation-free image. The results obtained using our method reveal much more texture for the saturated regions resulting in a realistic looking image.*

These images show a much more natural skin texture and appearance.

We also evaluated whether the colorization technique proposed by Levin et al. in [LLW04] could be used to remove the effects of saturation. The basic idea of this technique is the utilization of the luminance channel to estimate how similar the chrominance of neighboring pixels should be. These weights are then used to estimate the chrominance channels of the image. We applied colorization to remove saturated pixels by marking all pixels where at least one channel was saturated as pixels that needed to be colored.

Unfortunately, a direct application of colorization to the saturated image does not give us good results because the luminance value of the saturated pixels is still incorrect. As can be seen in Figure 5(b), the results still appear flat However, as can be seen in Figure 5(c), the results using our approach have a realistic, natural appearance.

4.4. Non-face Images

Figure 6 shows a case on non-face images. The saturation in the flower petals is mainly caused by strong sun illumination,

(a) Saturated Image (b) Our Result

Figure 6: *Results from non-face image. Our method is able to correct saturation on the flower petals, caused by strong sunlight.*

which is typical in outdoor environment. We can see that our saturation-free result is more realistic and visually pleasing.

4.5. Limitations

The primary limitation of our method is that at least one of the color channels cannot be saturated. If all three of the color channels are saturated then there is no information

Figure 7: *1st column are saturated images. 2nd column are our results. Our system fails to handle cases where a significant portion of the image is totally saturated (1st row) or where totally saturated pixels are edges pixels (2nd row).*

about the textures present in the saturated region. In these cases, the only solution is the application of in-painting techniques, such as [CPT03], [BSCB00], and [TCLT07]. Another limitation of our approach is handling totally saturated pixels along the edges. Recovering such pixels is cumbersome as neighboring pixels exhibit contrasting colors and thus the smoothness constraint fails to recovers the true color values.

The first case in Figure 7 shows an example of the type of image that are method cannot correct. In this image, all of the channels are saturated in a significant portion of the image. The second case in Figure 7 shows an example of a non-face image that our method has difficulty in correcting. As observed, the saturated pixels along and around the petal edges use smoothness to recover the color values. Since the petals(yellow) and grass(green) have contrasting colors, the final result does not present true color values for the saturated edge pixels.

5. Conclusion

We presented an approach to do saturation removal: given an image with saturated pixels, our approach removes saturated artifacts and produces a nice saturation-free image. Results show that it works well for general images. We also introduced a practical method to process images transformed with non-linear camera response curves. By first linearizing such images, the approach also generates descent saturation-free images.

Acknowledgements

This work was supported by grant HM-15820810021 through the NGA NURI program.

References

[APS98] ADAMS J., PARULSKI K., SPAULDING K.: Color processing in digital cameras. *IEEE Micro 18*, 6 (1998), 20–30.

[Bis95] BISHOP C. M.: *Neural Networks for Pattern Recognition*. Oxford University Press, 1995.

[BSCB00] BERTALMIO M., SAPIRO G., CASELLES V., BALLESTER C.: Image inpainting. In *SIGGRAPH '00: Proceedings of the 27th annual conference on Computer graphics and interactive techniques* (New York, NY, USA, 2000), ACM Press/Addison-Wesley Publishing Co., pp. 417–424.

[CPT03] CRIMINISI A., PÉREZ P., TOYAMA K.: Object removal by exemplar-based inpainting. *Computer Vision and Pattern Recognition, IEEE Computer Society Conference on 2* (2003), 721.

[Dui96] DUIN R.: The influence of spatial pixel relations on the image recognition performance. In *Proc. ASCI'96, 2nd Annual Conf. of the Advanced School for Computing and Imaging (Lommel, Belgium)* (June 5-7, 1996), pp. 248–252.

[Fat07] FATTAL R.: Image upsampling via imposed edges statistics. *ACM Transactions on Graphics (Proceedings of SIGGRAPH 2007) 26*, 3 (2007).

[HM99] HUANG J., MUMFORD D.: Statistics of natural images and models. In *In CVPR* (1999), pp. 541–547.

[HNS09] H. NACHLIELI R. BERGMAN D. G. C. S. B. O. G. R., SHAKED D.: Skin-sensitive automatic color correction.

[Kim99] KIMMEL R.: Demosaicing: Image reconstruction from color ccd samples. *IEEE Trans. Image Processing 8* (1999), 1221–1228.

[KK03] K.M. KRYSZCZUK A. D.: Color correction for face detection based on human visual perception metaphor. In *Proc. Workshop on Mult-Modal User Authentication* (Santa Barbara, 2003).

[LLW04] LEVIN A., LISCHINSKI D., WEISS Y.: Colorization using optimization. *ACM Trans. Graph. 23*, 3 (2004), 689–694.

[MLP00] MURESAN D., LUKE S., PARKS T.: Reconstruction of color images from ccd arrays.

[MS04] MEYLAN L., SÃIJSSTRUNK S.: Color image enhancement using a retinex-based adaptive filter. vol. 2, pp. 359–363.

[NM00] NAYAR S. K., MITSUNAGA T.: High dynamic range imaging: Spatially varying pixel exposures. In *Proc. IEEE CVPR* (2000), pp. 472–479.

[PGB03] PÉREZ P., GANGNET M., BLAKE A.: Poisson image editing. *ACM Tranactions on Graphics 22*, 3 (2003), 313–318.

[TCLT07] TING H., CHEN S., LIU J., TANG X.: Image inpainting by global structure and texture propagation. In *MULTIMEDIA '07: Proceedings of the 15th international conference on Multimedia* (New York, NY, USA, 2007), ACM, pp. 517–520.

[TLAF07] TAPPEN M. F., LIU C., ADELSON E. H., FREEMAN W. T.: Learning gaussian conditional random fields for low-level vision. In *IEEE Conference on Computer Vision and Pattern Recognition (CVPR07)* (2007).

[TWCL08] TAI S.-C., WANG N.-C., CHANG Y.-Y., LU Y.-C.: A two-stage contrast enhancement algorithm for digital images. In *CISP '08: Proceedings of the 2008 Congress on Image and*

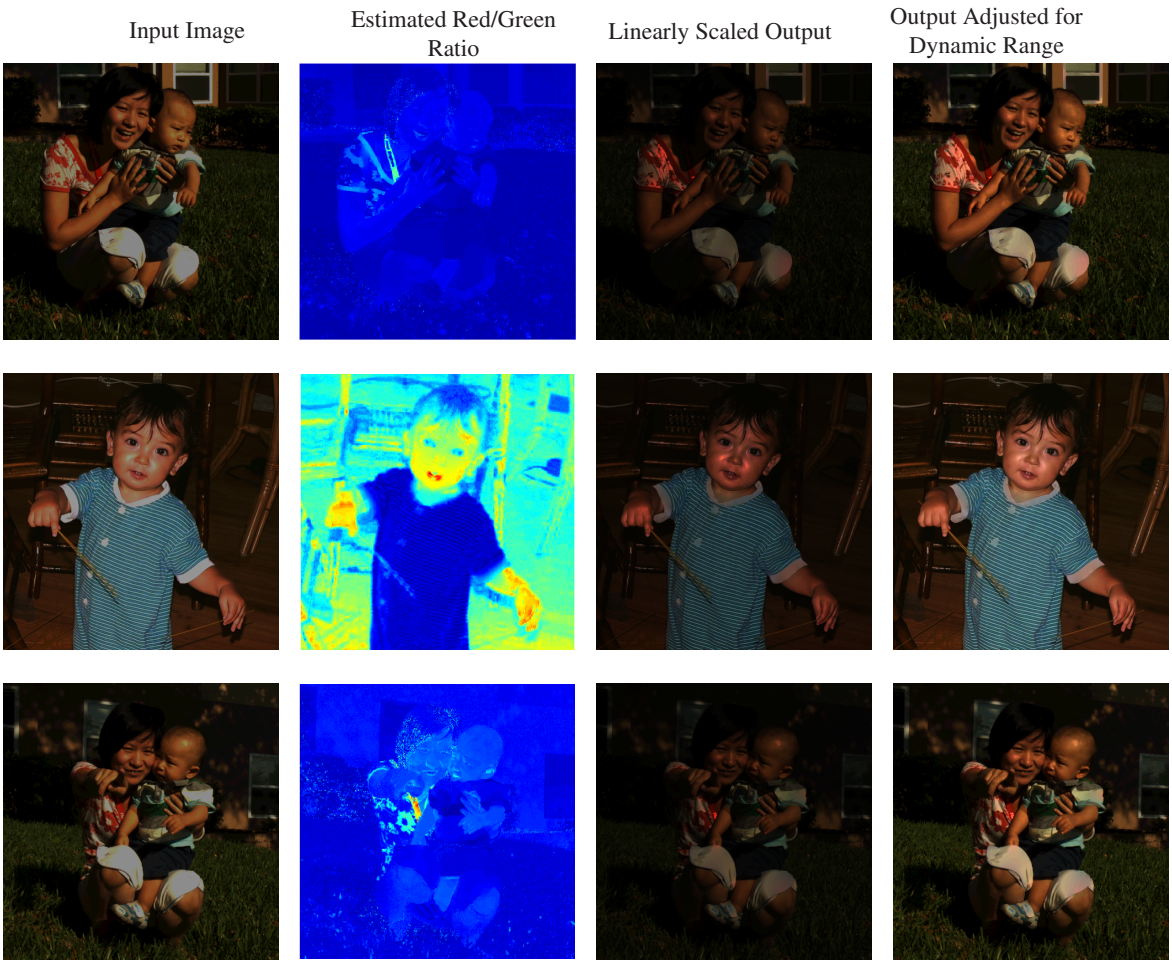

Figure 8: *1st column: Original saturated image. 2nd column: Recovered red:green ratio using the original saturated image. 3rd column: Saturation-free image generated using our method 4th column: Final result after pixel-wise nonlinear adjustment to the saturation-free image.*

Signal Processing, Vol. 3 (Washington, DC, USA, 2008), IEEE Computer Society, pp. 256–260.

[WF07] WEISS Y., FREEMAN W. T.: What makes a good model of natural images. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition* (2007).

[WWZ*07] WANG L., WEI L.-Y., ZHOU K., GUO B., SHUM H.-Y.: High dynamic range image hallucination. In *Rendering Techniques 2007: 18th Eurographics Workshop on Rendering* (June 2007), pp. 321–326.

[Yoh95] Yohkoh analysis guide : Instrument guide.

[ZB04] ZHANG X., BRAINARD D.: Estimation of saturated pixel values in digital color imaging. *JOSA-A 21*, 12 (December 2004), 2301–2310.

Figure 9: *1st column: Original saturated image. 2nd column: Recovered red:green ratio using the original saturated image. 3rd column: Saturation-free image generated using our method. 4th column: Final result after pixel-wise nonlinear adjustment to the saturation-free image.*