

Correcting Over-Exposure in Photographs

Dong Guo, Yuan Cheng, Shaojie Zhuo and Terence Sim
School of Computing, National University of Singapore, 117417
{guodong, cyuan, zhuoshao, tsim}@comp.nus.edu.sg

Abstract

This paper introduces a method to correct over-exposure in an existing photograph by recovering the color and lightness separately. First, the dynamic range of well exposed region is slightly compressed to make room for the recovered lightness of the over-exposed region. Then the lightness is recovered based on an over-exposure likelihood. The color of each pixel is corrected via neighborhood propagation and also based on the confidence of the original color. Previous methods make use of ratios between different color channels to recover the over-exposed ones, and thus can not handle regions where all three channels are over-exposed. In contrast, our method does not have this limitation. Our method is fully automatic and requires only one single input photo. We also provide users with the flexibility to control the amount of over-exposure correction. Experiment results demonstrate the effectiveness of the proposed method in correcting over-exposure.

1. Introduction

Over-exposure is a loss of highlight details in some bright regions of a photograph. It occurs when the light falling on the camera sensor exceeds what the sensor can capture. Over-exposure happens very often in daily-life photography due to the High Dynamic Range (HDR) of the scene. In photography, the term “dynamic range” is used to describe the ratio between the brightest and darkest measurable light intensities. The dynamic range of common digital cameras is very limited, usually 1000:1, which is much less than that of the real-world scenes. High contrast scenes, such as outdoor environment under direct sun light, may have a very high dynamic range, from 10^5 up to 10^9 . In such scenes, it is very difficult to make everything well-exposed; over-exposure is almost inevitable.

In practice, the photographers sometimes decrease exposure value (EV) to prevent over-exposure. But decreasing too much EV will make the photo dim and suffer from sensor noise.

Some works on High Dynamic Range (HDR) image cap-



Figure 1. Correcting over-exposure. (a) Input photograph, taken in an outdoor scene, with some portion over-exposed. (b) Detected over-exposed regions, marked in blue lines. (c) The result of our method, where the highlight of over-exposed regions is successfully reduced while the color is faithfully corrected.

turing, such as [2] [6], aim to fully capture the whole dynamic range. With tone mapping techniques, such as [3] [7] [1], the HDR images are mapped to Low Dynamic Range (LDR) ones, thus avoiding over-exposure. However, HDR cameras are too expensive, while other existing HDR capturing solutions usually require multiple shots with different exposure values. Such multiple inputs methods are restrictive, because they require the scene to be static. Furthermore, HDR capture works only for new photographs; it can not correct over-exposure in existing photographs.

In this paper, we present a method to correct over-exposed regions (Fig. 1(b)) in a single existing photograph (Fig. 1(a)). In our result (Fig. 1(c)), the highlight of over-exposed regions is greatly reduced yet the contrast is still preserved. Meanwhile the color of these regions is faithfully corrected.

The intensity of over-exposed regions is clipped at the maximum value (*e.g.* 255 in images with 8-bit per channel), thus appearing uniformly white. Therefore, a natural way to recover the over-exposed regions is to first estimate the actual value, *e.g.* a work on estimating HDR from LDR [8], and then compress the estimated HDR back to LDR image. However, in most cases, it is difficult to accurately estimate the actual value from a region if its information is completely lost. This is because the actual value might be slightly higher than the maximum value or boost up to a huge one (such as light from the sun). Instead of estimating the actual value and re-mapping to LDR, we present a method that slightly compresses the dynamic range of well-exposed regions while expanding the dynamic range of over-exposed regions. This directly produces an image with the over-exposure corrected.

Contribution. Our proposed method is effective in correcting over-exposed regions in existing photographs. The method is fully automatic and only requires a single input photo. The user has the flexibility to decide the amount of over-exposure correction.

2. Related Work

There is not much previous work directly addressing over-exposure correction. The closest are works by Zhang and Brainard [11] and Masood *et al.* [5]. In Zhang and Brainard’s work, the ratios between different color channels are used to recover the over-exposed channels. The ratios are estimated based on that of pixels around over-exposed regions. However, the assumption of spatial-invariant ratios in their work is inapplicable in real cases. Thus, Masood *et al.* utilize spatial-variant ratios in estimating pixel values in the over-exposed channels. Both the two works can only handle partial over-exposure *i.e.* one or two color channels are over-exposed. Regions of full over-exposure *i.e.* all three channels are all over-exposed are left untouched. However, in real photographs, partial over-exposed regions are quite limited in nature. In most cases, they only appear around full over-exposed regions as intermediate regions. In contrast, our algorithm works with both partial over-exposed and full over-exposed regions.

Some previous works focused on hallucinating HDR from an LDR image, such as Wang *et al.* [10] and Rempel *et al.* [8]. Wang *et al.* used texture synthesis algorithm to fill the detail texture in over-exposed regions. Users have to specify the clue where the texture of the over-exposed region is similar. The lightness of the over-exposed region is estimated by a Gaussian ellipsoid based on the neighbors around the over-exposed region. In the fashion of texture synthesis it is always required that similar regions should be available in the same photograph or other possible photographs. Also users’ hints for texture synthesis requires a lot labor work if there are too many over-exposed regions.

In contrast, the work of Rempel *et al.* aims to enhance the visualization of an LDR image on an HDR display device. A smooth brightness enhancement function is applied on and around the over-exposed region to boost up the dynamic range of the original image. However, color correction was not considered in this work.

HDR imaging can be used to capture an HDR scene without over-exposure. With HDR compression, such as [7] [3] [1], an HDR image can be compressed into an LDR image. This kind of tone-mapped image could be well-exposed anywhere, depending on the tone-mapping function. However, HDR cameras are priced too high. Other systems such as [2] [6] can composite multiple LDR photographs of the same scene with different exposure values. Thus, they require both the camera and the scene be static and the illumination be unchanged.

Instead of HDR capturing, some other works tackle the over-exposure with additional information. Zhang and Sim [12] proposed a method that can recover over-exposed regions by transferring details from a corresponding Near-Infrared (NIR) image. Thus their method may deal with scene with motion. However, it is still quite possible that both visible and near-infrared images are over-exposed simultaneously. Also special equipment is needed.

Both HDR and NIR imaging techniques are designed for capturing new photographs. They can not correct over-exposed region in existing photographs. In contrast, in this paper, we focus on correcting over-exposure with only one existing photograph.

3. Over-Exposure Correction

In an over-exposed region, the pixel values are clipped at the maximum value, usually 255. Thus, an over-exposed region becomes uniform at about 255 in all or some channels¹. Fig. 1(a) shows an example. The girl’s portrait was taken in an outdoor scene with strong sun-light from her left side. Although taken with auto exposure mode, her left face is still over-exposed, marked in blue lines in Fig. 1(b). In the following part, we use Ω to denote the over-exposed region, and $\neg\Omega$ to denote the rest of image.

There are two aspects in correcting over-exposure in Ω *viz.* lightness recovery and color correction. The actual lightness of over-exposed regions should be at least the maximum value. Thus, a natural way to correct over-exposure is to “hallucinate” the lightness first. As introduced in Wang *et al.* [10], a Gaussian ellipsoid is used to fit the boundary of an over-exposed region so as to guess the lightness inside. However, in actual fact, it might be incorrect if strong light sources exist *e.g.* the sun in outdoor scene. The lightness level is rather difficult to estimate.

¹Sometimes, due to the compression algorithm of JPEG or other format, this value might be slightly lower than 255.




Figure 2. The workflow of our method. Two aspects are included, lightness recovery and color correction. Lightness recovery is through L^* channel using over-exposure likelihood \mathcal{P} . Color correction is through a^* , b^* channels using color confidence Ψ . Both \mathcal{P} and Ψ are derived from an over-exposure map \mathcal{M} , which is generated from the input. See text for details.

Yet, for subsequent display purpose, the hallucinated lightness should be compressed back to LDR. This inspired us to design an algorithm to recover the over-exposed regions directly in a low dynamic range image. This, on one hand, avoids directly estimating the lightness, and on the other hand, directly makes good use of the original information captured by cameras.

During color correction, color in Ω is corrected via neighborhood propagation based on both neighborhood similarity and the confidence of the original color.

To separately deal with lightness and color, the input image is first converted to CIELAB colorspace, where the L^* channel represents lightness and a^* , b^* channels represent the color. In the rest of this paper, we use L to represent L^* channel and $C = (a \ b)^T$ to represent a^* and b^* channels of the input image. \tilde{L} , \tilde{C} are defined similarly to represent the L^* and a^*b^* channels of the result image, respectively.

The workflow of our algorithm is shown² in Fig. 2. The input is a photograph captured by a digital camera, with over-exposed regions. First, over-exposed regions are detected, denoted by an over-exposed map. The over-exposed likelihood and color confidence are generated from the map. Lightness recovery and color correction are performed us-

ing the two probability maps. The recovered lightness and corrected color are combined in the output image.

3.1. Over-exposure detection

Previous works usually use a simple scheme to detect over-exposure: If the value of a pixel is equal or larger than a threshold, the pixel is considered over-exposed. Usually, the threshold is set to 254 to eliminate the effects of the error due to the compression algorithm. However, such a hard threshold does not handle well the gradual transition from over-exposure regions to their immediate neighbors. The color of these neighbors is desaturated ($\|C\|_2$ becomes smaller) and lightness increases. Thus, we create an over-exposed map \mathcal{M} which is in $(0, 1)$ denoting how much a pixel is affected by over-exposure. The map \mathcal{M} is defined based on the L and C values. The pixel is more likely considered over-exposed, if L is larger or $\|C\|_2$ is smaller. Thus, \mathcal{M}_i is defined as

$$\mathcal{M}_i = \frac{1}{2} \left(\tanh \left(\delta \cdot \left((L_i - L_T) + (C_T - \|C_i\|_2) \right) \right) + 1 \right), \quad (1)$$

where L_T and C_T denote the boundary value of the over-exposure region, which makes $\mathcal{M} = 0.5$, and δ controls how fast \mathcal{M}_i grows with larger L_i or smaller $\|C_i\|_2$. We use $\delta = 1/60$, $L_T = 80$ and $C_T = 40$ in our experiments.

²Some images in Fig.2 3 4 are rendered in false color for better visualization. Warmer color (red) denotes higher value.

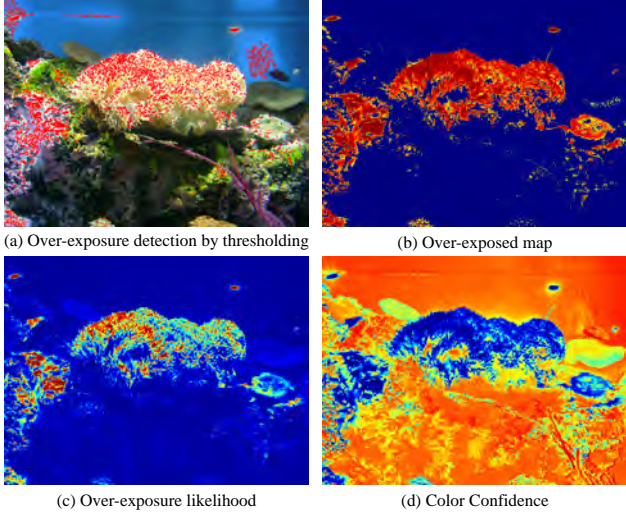


Figure 3. Illustration of (a) over-exposure detection by simple thresholding, (b) over-exposed map (only showing $\mathcal{M} > 0.5$), (c) over-exposure likelihood, and (d) color confidence. The input image is shown in Fig. 6.

We show an example of \mathcal{M} in Fig. 3. Fig. 3(b) shows the area $\mathcal{M} > 0.5$, which covers much more area than the detection result by simple thresholding method (Fig. 3(a)). To be more clear, the region Ω is defined as all pixels whose $\mathcal{M} \geq 0.5$ and $\neg\Omega$ for $\mathcal{M} < 0.5$. Ω defines the area that is seriously affected by the light in the scene and needs correction.

3.2. Lightness recovery

An over-exposed region is rather flat due to clipping. To make room for the recovered lightness of Ω , we use a tone mapping technique to compress the dynamic range of $\neg\Omega$. Once the dynamic range of $\neg\Omega$ is compressed, Ω could be expanded to fill the gap.

Here we introduce *over-exposure likelihood* \mathcal{P} to measure how likely the pixel in region Ω is still over-exposed in the output image. \mathcal{P} is defined based on \mathcal{M} , *i.e.*

$$\mathcal{P}_i = \frac{1}{K} \cdot \frac{1}{1 - \mathcal{M}_i}, \quad (2)$$

where K is a normalization factor to make $\max_i \mathcal{P}_i = 1$.

In a sense, \mathcal{P} reflects relative value of the actual lightness in Ω . An example of \mathcal{P} is shown in Fig. 3(c). \mathcal{P} is zero in most part, and up to 1 only when \mathcal{M} is very large.

For dynamic range compression in $\neg\Omega$, we adapt the method proposed by Fattal *et al.* [3]. Specifically, the gradient of the image is attenuated non-linearly: larger gradient is compressed more than smaller ones. The attenuation factor in $\neg\Omega$ is a power of the magnitude of the gradient. Gradient in Ω is kept unchanged to make the recovered light-

ness smooth in most places and keep details if any. The attenuation function $\mathbf{z}(\cdot)$ is defined as

$$\mathbf{z}(\nabla L_i) = \begin{cases} \left(\frac{\|\nabla L_i\|}{\alpha} \right)^{1-\beta} \cdot \nabla L_i & \text{if } i \in \neg\Omega \\ \nabla L_i & \text{otherwise} \end{cases}, \quad (3)$$

where α and β are two parameters to control the compression ratio of the image. α is to control the minimal gradient that is compressed. It is usually set to the 0.1 times the average gradient magnitude. So, β is left to control the compression ratio. This is a user adjustable parameter to control the overall effect of our results. We use $\beta = 0.9$ in most of our experiments. The choice for β will be discussed in Section 4.

Now that the desired gradients are obtained, the first objective is to keep the gradients of result image as similar as possible with these gradients. This leads to the energy of

$$\mathcal{E}_1 = \sum \|\nabla \tilde{L}_i - \mathbf{z}(\nabla L_i)\|_2^2 \quad (4)$$

should be minimized.

On the other hand, we need to keep the lightness in Ω as close as possible to the original value, with different likelihood \mathcal{P} . Thus, we define the second energy

$$\mathcal{E}_2 = \frac{1}{|\Omega|} \sum_{i \in \Omega} \mathcal{P}_i |\tilde{L}_i - L_i|^2, \quad (5)$$

where $|\Omega|$ denotes the number of elements in Ω .

The lightness of $\neg\Omega$ tends to be lower due to the compression of its dynamic range, while the lightness in Ω tends to keep its original high value with different likelihood. As a result, the lightness in Ω is modified according to \mathcal{P} , which represents the relative lightness in Ω . To recover the lightness, an overall energy

$$\mathcal{E}_L = \mathcal{E}_1 + \lambda \mathcal{E}_2, \quad (6)$$

is to be minimized with a hard constraint that

$$\tilde{L} = L \quad \text{if } L < \min(L) + r(\max(L) - \min(L)) \quad (7)$$

where $r = 0.1$. λ is to balance gradient energy \mathcal{E}_1 and value energy \mathcal{E}_2 . Smaller λ means the lightness of Ω is more affected by dynamic range compression in $\neg\Omega$. We use $\lambda = 5$ in our experiments which produce good results. The hard constraint (7) means the low-lightness regions are kept unchanged. Minimizing (6) is equivalent to solving a banded linear system, which can be solved efficiently. An example of recovered lightness is shown in Fig. 2, labeled with \tilde{L} .

3.3. Color correction

The color in or around Ω is more or less affected by over-exposure. We can use the over-exposed map \mathcal{M} to represent how confident a pixel color is. Ψ is defined as

$$\Psi_i = 1 - \mathcal{M}_i. \quad (8)$$

An example of Ψ is shown in Fig. 3(d).

We attempt to estimate less confident color from more confident one, propagating pixel from pixel via neighborhood similarity. A similar work by Levin *et al.* [4] aims to colorizing a gray-scale image. The color of each pixel is propagated from the user stroke (color) via neighbor pixels. The similarity in color is based on the gray value similarity. For color correction in this problem, the similarity is based on the lightness difference as well as the original color information that is confident.

The color of each pixel is similar to its neighbors with a similarity weight. Also, it is similar to its original color with a confidence Ψ . Thus, the color is corrected by minimizing \mathcal{E}_c , which is defined as

$$\sum_i \left[(1 - \Psi_i) \left\| \tilde{C}_i - \sum_{j \in \mathcal{N}_i} w_{i,j} \tilde{C}_j \right\|_2^2 + \Psi_i \left\| \tilde{C}_i - C_i \right\|_2^2 \right], \quad (9)$$

where \mathcal{N}_i denotes the neighborhood of a pixel i .

For a pixel $i \in \Omega$, its confident value Ψ_i is nearly zero. The second term of \mathcal{E}_C is omitted, *i.e.* the original color value of this pixel is ignored. Thus the color for this pixel is propagated from its neighboring pixels. In contrast, for a pixel around Ω , its confident value $\Psi_i < 1$, both the two terms have influence in determining its color \tilde{C}_i . For a pixel far from Ω , its original color is dominating. The original color tends to be unchanged.

Our strategy for setting the weights w_{ij} is similar to that in bilateral filter [9]. The weight is product of several Gaussian functions, *i.e.*

$$w_{ij} = G(i-j)G(D_{\tilde{L}}(i,j))G(D_a(i,j))G(D_b(i,j)), \quad (10)$$

where $D_{\tilde{L}}(\cdot, \cdot)$, $D_a(\cdot, \cdot)$ and $D_b(\cdot, \cdot)$ denote the distance of corrected L^* , and original a^* and b^* channels, respectively. The first Gaussian measures the spatial distance, while the second one measures the lightness difference. In other words, pixels that are nearer tend to have more similar color; pixels whose lightness is similar tend to have more similar color. The third and fourth Gaussian functions measure the influence of the original color difference.

4. Results and Discussion

With different amount of well-exposed region being compressed, different levels of over-exposed region is corrected. In Fig. 4, we show a series of results based on different β values. Smaller β results in more compression on $-\Omega$ and thus makes more space for recovering Ω . As a result, Ω in result appears darker. When $\beta = 1$, the L channel stays untouched, while color channels are still corrected. The result is slightly better than the input. As β decreases, the exposure reduces yet still keeping the relative contrast. A

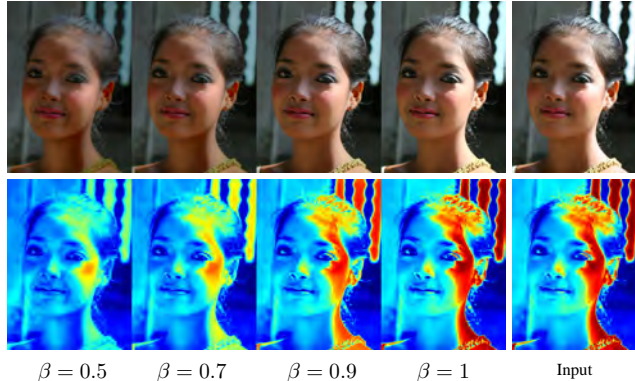


Figure 4. Results comparison with different β . Top row, recovered and input image. Bottom row, recovered and original L channel.

too low β may cause the $-\Omega$ too dim, which is also undesirable. Usually, β ranging from 0.8 to 0.9 yields a good result. $\beta = 0.9$ is the value we used to obtain most results.

A comparison of our result with that of Masood *et al.* is shown in Fig. 5. As there are regions with all three color channels over-exposed, *e.g.* on the right side wall (marked in red circle). Their method failed to recover these regions³. In contrast, our method reduced strong sun-light and corrected the color of the over-exposed region.

More results are shown in Fig. 6. In the flower example, many petals are over-exposed, while the bee on the flower is well-exposed. In our result, the strong reflection on petals is suppressed and the color is perfectly corrected. The bee still appears well-exposed. In another example, the body of coral is largely over-exposed. The color becomes pale. In contrast, in the result, it looks natural and well-exposed. In the kid and girl examples, faces under strong sunlight become too bright. In our results, the lightness is reduced and color of skin is faithfully corrected. The arm of the Buddhist statue reflects strong sunlight, resulting in over-exposure in the photo. We successfully corrected the over-exposure while still keeping the shining effects on the arm.

5. Conclusion and Future Work

In this paper, we have presented a method of correcting over-exposure on an existing photograph. Instead of recovering the actual lightness of the over-exposed regions and then compressed back into the image range, we directly estimate the value in the output image. The compression of well-exposed regions makes room for the over-exposed regions to expand the dynamic range. An over-exposure likelihood is employed to derive the lightness of over-exposed regions in the result image. Color correction is based on the color from boundary of over-exposed regions, the similarity

³The result image by Masood *et al.* was generated by the code provided on the authors' website.



Figure 5. Comparison of our result with that of Masood *et al.* [5]. Artifacts in (b) are indicated in red circles, due to their limitation in handling regions with all three color channels over-exposed. In our result (c), the over-exposed regions are corrected successfully.



Figure 7. Failure case of our method. The boundary between the face and the background is missing due to serious over-exposure of the photograph. Our color correction method may propagate the color from the face into the background.

of pixel neighborhood, and the confidence of the original color. Good results have demonstrated the effectiveness of our method.

Limitation: For severely over-exposed photographs, the boundary may become unclear between two adjacent objects. Our color correction method may propagate the color across objects, which is undesirable. Fig. 7 shows such an example. Due to severe over-exposure, the old man’s face and the window in background are connected. As a result, the window is colored by red from the old man’s face. In contrast, although in the same photo, the face at right bottom is successfully corrected because the boundary of the face is very clear. A possible way to overcome this limitation is to ask user to provide boundary clues.

Blooming effects usually come with over-exposure. Both the lightness and color of regions around the over-exposed region are affected. Our current method can correct somehow the color of the blooming effects, yet nothing for lightness. One possible future work would be fixing the blooming effects in over-exposure correction.

6. Acknowledgment

We thank the reviewers for their valuable comments, Xiopeng Zhang, Hui Ji, Ning Ye for their insightful discus-

sions, and Daniel Schwen, Han Lu, Jing Sun, and Xianjun Wang for providing their photographs. We acknowledge the generous support of NUS.

References

- [1] H.-T. Chen, T.-L. Liu, and T.-L. Chang. Tone reproduction: A perspective from luminance-driven perceptual grouping. In *Proc. CVPR*, 2005.
- [2] P. E. Debevec and J. Malik. Recovering high dynamic range radiance maps from photographs. In *Proc. ACM SIGGRAPH*, pages 369–378, 1997.
- [3] R. Fattal, D. Lischinski, and M. Werman. Gradient domain high dynamic range compression. *ACM Trans. Graphics*, 21(3):249–256, 2002.
- [4] A. Levin, D. Lischinski, and Y. Weiss. Colorization using optimization. *ACM Trans. Graphics*, 23(3):689–694, 2004.
- [5] S. Z. Masood, J. Zhu, and M. F. Tappen. Automatic correction of saturated regions in photographs using cross-channel correlation. In *Proc. Pacific Conference on Computer Graphics and Applications*, 2009.
- [6] T. Mitsunaga and S. K. Nayar. Radiometric self calibration. In *Proc. CVPR*, 1999.
- [7] E. Reinhard, M. Stark, P. Shirley, and J. Ferwerda. Photographic tone reproduction for digital images. *ACM Trans. Graphics*, 21(3):267–276, 2002.
- [8] A. G. Rempel, M. Trentacoste, H. Seetzen, H. D. Young, W. Heidrich, L. Whitehead, and G. Ward. Ldr2hdr: on-the-fly reverse tone mapping of legacy video and photographs. *ACM Trans. Graph.*, 26(3):39, 2007.
- [9] C. Tomasi and R. Manduchi. Bilateral filtering for gray and color images. In *Proc. ICCV*, 1998.
- [10] L. Wang, L.-Y. Wei, K. Zhou, B. Guo, and H.-Y. Shum. High dynamic range image hallucination. In *Proc. EGSR*, 2007.
- [11] X. Zhang and D. H. Brainard. Estimation of saturated pixel values in digital color imaging. *J. Opt. Soc. Am. A*, 21(12):2301–2310, 2004.
- [12] X. Zhang, T. Sim, and X. Miao. Enhancing photographs with near infrared images. In *Proc. CVPR*, 2008.



Figure 6. Results of correcting over-exposure. From left to right, top to down, flower, coral, kid, plant, statue, girl, and leaves. For each pair, top row is the input image, and bottom row is our result. Over-exposed regions are indicated by red circles.