

Gradient Domain Color Restoration of Clipped Highlights

Mushfiqur Rouf

Cheryl Lau

Wolfgang Heidrich

The University of British Columbia

Abstract

Sensor clipping destroys the hue of colored highlight regions by misrepresenting the relative magnitude of the color channels. This becomes particularly noticeable in regions with brightly colored light sources or specularities. We present a simple yet effective gradient-space color restoration algorithm for recovering the hue in such image regions. First, we estimate a smooth distribution of the hue of the affected region from information at its boundary. We combine this hue estimate with gradient information from channels unaffected by clipping to restore clipped color channels.

1. Introduction

Colored light sources are ubiquitous in modern environments, with examples ranging from sodium street lights to neon signs, warning and exit signage, as well as colored LEDs used both as indicator lights and for architectural lighting. Photographing scenes with such colored lights is challenging – the light sources themselves are often orders of magnitude brighter than reference white surfaces in the scene. Due to the limited dynamic range of image sensors, color channels are then clipped independently of each other, based on the color of the light.

In the case of colored light sources, the clipping not only alters the intensity of the affected image region but also its hue, which can affect the mood of the scene or make its rendition unrealistic. While standard multi-exposure high dynamic range (HDR) imaging techniques such as the work by Debevec and Malik [6] can restore both the intensity and the hue, we aim at restoring just the correct hue of the clipped regions from *a single photograph*. Our goal is therefore not a new HDR capture technique or heuristic for boosting the dynamic range but instead to devise a simple yet effective method for restoring colors of clipped regions, thereby generating a version of a traditional low dynamic range image with improved color rendition.

Our method is based on the observation that, in many cases, colored lights may only result in clipping *some* of the color channels while leaving others unaffected. We can use this property to restore the hue and brightness of the clipped

channels for certain types of scenes.

As an example, consider the neon sign depicted in Figure 1. From the reflection in the building facade we can infer that the sign itself emits red light, yet the neon tubes are depicted as yellow due to selective clipping of the red and green channels in the photograph. Our algorithm manages to reconstruct the correct color of the neon sign and to restore washed out details in partially clipped regions such as the curtains (Figure 1).

Our method is based on a gradient domain approach. For each image region with at least one clipped color channel, we first estimate smooth hue distributions using data from pixels just outside the clipped region. We combine these hue estimates with gradient information from the color channels unaffected by clipping in order to estimate the gradients of the clipped color channels. We then restore the image by solving a Poisson problem. In an optional preprocessing step, we can also fill in a smooth gradient field in regions where all three color channels are clipped. Doing so avoids discontinuities in the gradient field and resulting Mach bands at the transition from partially to fully clipped image regions.

2. Related Work

Gradient domain image processing has become a powerful technique for image manipulations, starting with Elder and Goldberg’s work on contour domain image editing [10] and continuing with general formulations for image manipulations in the gradient domain [20, 4]. [14] proposed a user guided colorization of photographs using gradients under the assumption that drastic color changes in natural images are usually correlated with strong edges in the greyscale input image. While our method falls within the scope of general gradient domain processing, to the best of our knowledge, our method is the first automatic method to use this tool for restoring clipped highlights.

Clipped signal restoration: For band limited 1D signals, reconstruction algorithms have been proposed for situations where the number of missing samples is low [1] or where a statistical model of an undistorted signal is known [19]. However, neither of these approaches can be trivially extended to images because current natural image statistics

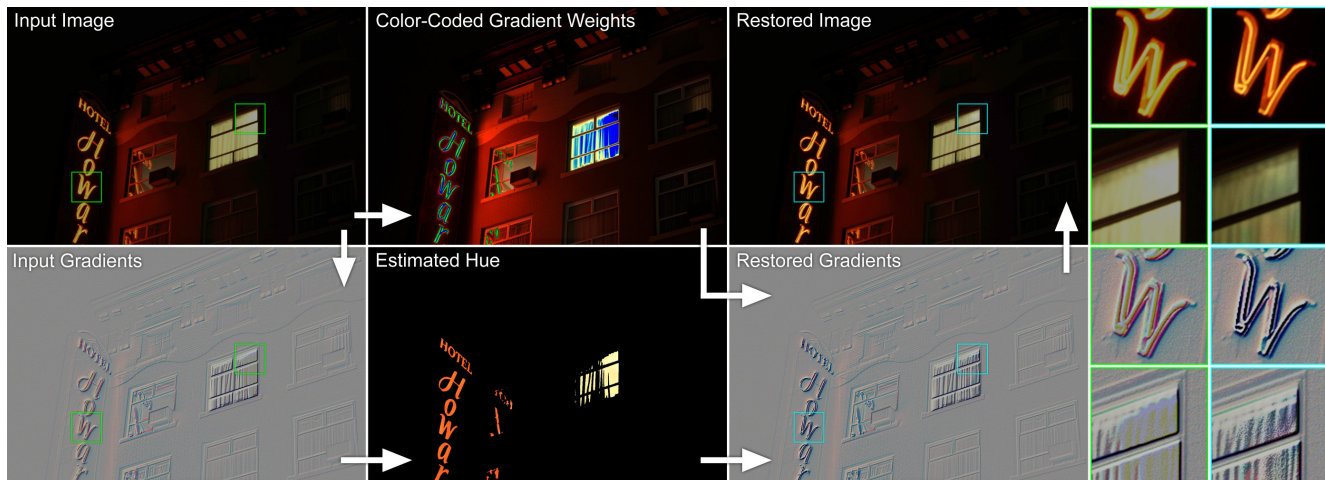


Figure 1: Color restoration in the gradient domain. Discoloration artifacts due to sensor clipping are evident in the input image – notice the color shift in the neon sign and the loss of detail in the curtain region. Our restoration process first estimates the hue of partially clipped pixels. Next, we combine the hue with the gradient field of the input image in order to compute the restored gradients. The weights indicate the level of confidence in each captured value in red, green and blue channels – we have low confidence in values close to zero or one. Then an integration gives the final result. Note the enhancements to the neon sign and the top-right corner of the curtain and corresponding restoration in the gradient domain. Images have been tone-mapped to visualize the restored colors.

are insufficient for restoring clipped pixels. For noisy images, pixel values just above the clipping threshold can be restored [11].

HDR imaging: There has also been a lot of work on merging multiple exposures for HDR imaging, starting with the seminal work by Debevec and Malik [6]. Unlike this approach and later approaches (e.g. [12, 18, 8, 23]), we do not aim to change the photography process to increase the amount of information captured about a scene, but instead aim to extract as much information as possible from a single, given photograph.

LDR to HDR enhancement: Reconstructing an HDR image from a single exposure with clipped values is a challenging problem that yields only approximate solutions based on heuristics or manual user intervention [17, 2, 22, 7]. These methods estimate only the brightness and not the hue of the clipped regions. The goal of our work is complementary to these approaches; we do not seek to boost the dynamic range of an image but simply to faithfully restore the hue of clipped image regions for a better color rendition of colored lights. While this process does extend the pixel values outside the original 3D color gamut, the overall gain in luminance is typically small, and the image remains faithful to a traditional (LDR) photograph. If desired, our method could be combined with any of the mentioned heuristics for dynamic range expansion.

Restoration of clipped colors: The problem we consider in this paper has received some attention before. In the case of color images, pixels that are clipped in one or two color channels can be estimated using cross-channel corre-

lation by modeling the pixel values as a combination one or more 3D Gaussian distribution(s) [24]. Guo et al. [13] recover color and lightness through propagation of information. While they correct cases with all channels clipped, their algorithm involves human intervention. Masood et al. [16] and Elboher and Werman [9] restore highlights in the spatial domain. DCRAW, a popular public domain software package for processing RAW image file formats [5], also has a restoration mode for clipped color channels. As we show in the paper, all these spatial domain approaches are prone to severe discontinuity artifacts which are eliminated with our gradient domain approach.

3. Method

Our method for gradient reconstruction is based on three steps (see Figure 2): an (optional) preprocessing step that smoothly fills in gradients in image regions where all color channels are clipped (Section 3.4), smooth hue estimation for clipped pixels from information just outside the clipping region (Section 3.2), and finally, detail transfer from unclipped to clipped channels (Section 3.3). In partially clipped areas where at least one channel of the input image remains unsaturated, this approach recovers both the hue and the texture; in fully clipped regions we recover the hue only.

All three steps are performed in gradient space and can be reduced to simple gradient manipulations and a sequence of independent Poisson solves. While this is a very simple algorithm, it has the advantage of being easy to implement, and we demonstrate that it is highly effective in producing

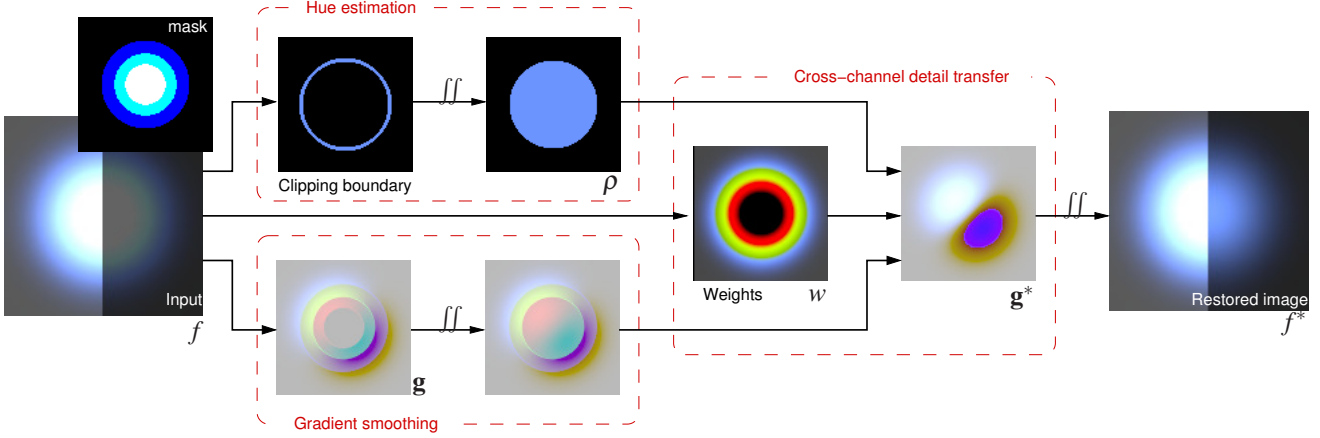


Figure 2: Overview of our method. We compute a gradient field \mathbf{g} given an input f . We estimate a smooth hue distribution ρ over the clipped region from the observed hue at its boundary. Guided by ρ , we then estimate the unknown intrinsic gradient field \mathbf{g}^* with a weighted (w) combination of gradients from unclipped channels. Note that we can only restore the original \mathbf{g}^* at pixels where least one channel remains unclipped. However, in regions with all channels clipped, we can optionally smooth the gradients to avoid abrupt changes visible as Mach bands. All steps, including the final restoration of f^* can be cast into simple Poisson equations.

high quality hue restorations.

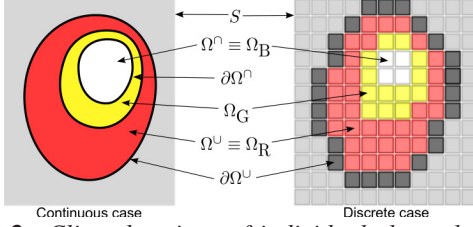


Figure 3: Clipped regions of individual channels are expressed as Ω_R , Ω_G and Ω_B . Ω^U denotes pixels with any channel clipped, while Ω^N denotes pixels with all channels clipped. $\partial\Omega^U$, etc. denote corresponding region boundaries. Note that we have partial data in $\Omega^U \setminus \Omega^N$ and no data in Ω^N ; hence our algorithm reconstructs scene details of f^* only in $\Omega^U \setminus \Omega^N$, and restores color in Ω^N .

3.1. Image formation model

Let $f_k^*(p)$ be the k^{th} color channel of an (unknown) intrinsic image f^* at position $p \in \mathbb{R}^2$. In this intrinsic image, which we seek to restore, color channels are unaffected by clipping and correspond to the *native* color channels of a capture device, i.e. the channels that directly correspond to the color filters of a camera. If we assume a camera with a limited dynamic range $(0 \dots 1]$, the image that is actually captured by this camera is given by the channels

$$f \triangleq \min(1, f^* + n), \quad (1)$$

where n represents a noise term. We now define $\Omega_k = \{p \in \mathbb{R}^2 : f_k(p) = 1\}$ as the set of image positions p where channel k is clipped (Figure 3). Image regions with *at least one* clipped channel are denoted as Ω^U , and regions with *all*

channels clipped as Ω^N :

$$\Omega^U \equiv \bigcup_k \Omega_k \quad \text{and} \quad \Omega^N \equiv \bigcap_k \Omega_k. \quad (2)$$

Finally, we define $\partial\Omega_k$, $\partial\Omega^U$ and $\partial\Omega^N$ to be the boundaries around the corresponding sets.

The fundamental assumption we make in our work is that the hue varies smoothly over Ω^U and can be estimated from the pixels on its boundary $\partial\Omega^U$, for example because glare extends the hue of the clipped regions into the boundary. This assumption is valid for highlights generated by a single colored light source such as an LED or a neon sign, similar to the image in Figure 1. It is, however, violated for scenes such as sunsets where the hue of the sky may not be independent of the luminance. In such scenes, estimating the hue based only on measurements that are dim enough to fall below the clipping threshold mis-estimates the hue and will not result in plausible reconstructions, as we will show in Section 4.

3.2. Hue interpolation

First, we generate a smooth hue estimate of all regions Ω^U containing at least one clipped channel. As mentioned above, we assume that the hue of this region can be interpolated from the (known) hue on its boundary.

In our implementation, the hue ρ is represented as a multichannel image, with the same number of channels and color space as f and f^* . We perform the interpolation by solving a Laplace equation over Ω^U with a Dirichlet boundary condition in $\partial\Omega^U$:

$$\nabla^2 \rho = 0 \quad \text{over } \Omega^U \quad \text{with } \rho|_{\partial\Omega^U} = f|_{\partial\Omega^U}. \quad (3)$$

This is a standard Poisson problem that can be solved very efficiently. Although one could use more sophisticated inpainting techniques to produce more detailed hue maps, we found that our smooth hue estimates work well for a large range of images and are in fact more robust than, for example, the edge-stopping interpolation used by Masood et al. [16] (see discussion in Section 4, Figure 7).

Boundary cleanup. Image noise and sampling artifacts from single-chip cameras with color filter arrays such as Bayer patterns [3] can result in high-frequency hue variations on the boundary that result in distracting artifacts when they serve as the basis for hue estimation. In order to suppress these high-frequency variations, we apply a small 1D bilateral filter *along the boundary* on $f|_{\partial\Omega^U}$ to suppress noise. In our experiments, we use a spatial (domain) sigma of 5 pixels and a range sigma of 0.25 (out of 1). Where possible, we further suggest to use simple linear interpolation for demosaicing the boundary pixels $\partial\Omega^U$, while more sophisticated methods can be used elsewhere in the image.

3.3. Cross-channel detail transfer and color restoration

In a second step, we combine the estimated hue with information from unclipped channels, where available, to estimate the pixel values of f^* . We first discuss the case of image regions where all channels f_j except for $f_k = f_k^*$ are clipped. In this case, the known values from channel f_k^* and the channels ρ_j of the estimated hue uniquely define pixel values of the clipped channels:

$$f_j^* = \frac{\rho_j}{\rho_k} f_k^* = \frac{\rho_j}{\rho_k} f_k. \quad (4)$$

Gradient domain formulation. The spatial reconstruction mentioned above works well for the case of a single channel providing unclipped image data. In regions where multiple channels provide valid data, the competing information must be reconciled with the hue estimates in a spatially smooth fashion (Figure 4). To this end, we first re-cast the problem as a gradient domain reconstruction.

Let $\mathbf{g} = \nabla f$ and $\mathbf{g}^* = \nabla f^*$ be the gradient vector field of the captured and the intrinsic images, respectively. The gradient domain version of Equation 4 can be obtained by computing the gradient of both sides and assuming a locally constant hue ρ :

$$\mathbf{g}_j^* = \frac{\rho_j}{\rho_k} \mathbf{g}_k. \quad (5)$$

Given a gradient estimate, we recover each channel f_k^* by solving a Poisson equation over the clipped region in that channel Ω_k with a Dirichlet boundary condition in $\partial\Omega_k$:

$$\nabla^2 f_k^* = \nabla \cdot \mathbf{g}_k^* \text{ over } \Omega_k \text{ with } f_k^*|_{\partial\Omega_k} = f_k|_{\partial\Omega_k}. \quad (6)$$

In fully clipped regions Ω^\cap , where no scene detail is present in the captured image f , \mathbf{g}^* will be (mis-)estimated as 0, and consequently the reconstructed image will be flat but will contain the estimated hue. The transition from valid gradient data to flat image regions may cause Mach bands. In Section 3.4 we describe a method for filling in smooth gradients before the detail transfer step to avoid this problem.

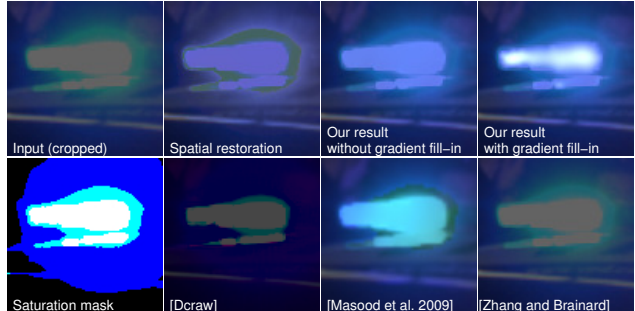


Figure 4: Advantages of our gradient domain method. We show a particular case cropped from Figure 9(f) where the spatial approaches would fail. In comparison, our gradient based approach faithfully restores the color.

Multiple reference channels. If two or more channels remain unclipped, we have multiple, possibly conflicting sources of gradient information. In this case, we use a weighted combination of reference gradients,

$$\mathbf{g}_j^* = \frac{\sum_{k \neq j} w_k \cdot \rho_j / \rho_k \cdot \mathbf{g}_k}{\sum_{k \neq j} w_k}. \quad (7)$$

Since \mathbf{g}^* is a combination of multiple gradient fields, it might not be integrable even though \mathbf{g} is. The Poisson system projects this estimated gradient field onto a feasible space. In order to choose an appropriate weighting function w in Equation 7 above, we observe that:

- Weights should be proportional to the reliability of the gradients. In images exhibiting photon shot noise, smaller pixel values should get lower weights.
- In order to avoid discontinuity artifacts like Mach bands at the border between regions with different numbers of clipped channels, the weighting function should have a smooth profile overall, including close-to-zero slopes near values 0 and 1.

In consideration of these factors, we choose a piecewise cubic weighting function with an off-center peak $m \in [0, 1]$:

$$w_k(p) = \begin{cases} 3 \left(\frac{f_k(p)}{m} \right)^3 - 2 \left(\frac{f_k(p)}{m} \right)^2 + \epsilon & \text{if } f_k(p) \leq m \\ 3 \left(\frac{1-f_k(p)}{1-m} \right)^3 - 2 \left(\frac{1-f_k(p)}{1-m} \right)^2 + \epsilon & \text{otherwise} \end{cases} \quad (8)$$

ϵ is used to avoid zero weighting which can cause division-by-zero. In our implementation, $m = 0.65$ and $\epsilon = 10^{-3}$.

Figures 4 and 7 show comparisons of our gradient based method with a spatial reconstruction using the same channel weights, as well as several spatial methods. Note how the spatial reconstructions suffer from discontinuous changes in hue while our gradient-based approach provides a smooth reconstruction.

3.4. Gradient smoothing for fully clipped regions

As mentioned above, gradient fields in fully clipped regions Ω^\cap are flat because all sensor values are clipped to the clipping threshold. Derivative discontinuities in the gradient field at the boundaries $\partial\Omega^\cap$ of these areas may become visible in the reconstruction results as Mach bands (Figure 4, row 1, column 3). To suppress such artifacts, we use an optional pre-processing stage, in which we generate gradients for one of the channels over Ω^\cap . We only apply this method if there is one channel k whose clipping region is completely contained within the clipping regions of the other channels, i.e. $\Omega^\cap = \Omega_k$.

This smooth gradient infilling can again be cast as a set of two sequential Poisson problems with Dirichlet boundary conditions, this time in log space (quantities with $\hat{\cdot}$ are computed on log images):

$$\nabla^2 \hat{\mathbf{g}}_k^* = 0 \quad \text{over } \Omega_k \quad \text{with } \hat{\mathbf{g}}_k^*|_{\partial\Omega_k} = \nabla \log f_k|_{\partial\Omega_k}, \quad (9)$$

$$\nabla^2 \hat{f}_k^* = \nabla \cdot \hat{\mathbf{g}}_k^* \quad \text{over } \Omega_k \quad \text{with } \hat{f}_k^*|_{\partial\Omega_k} = \log f_k|_{\partial\Omega_k}. \quad (10)$$

The linear space channel k can then be recovered as $f_k^* = \exp(\hat{f}_k^*)$. The motivation for solving this problem in log space is that it corresponds to a generalization of fitting a Gaussian to the gradients on the boundary Ω^\cap , as can be seen by analyzing a 1D example (Figure 5). Given a clipped input signal (blue) in linear domain (Figure 5, left), we first take the log of this signal and then solve for a gradient (red in Figure 5, right), which will vary linearly over the clipping region. By integrating this gradient up with a second Poisson solve, we obtain a log image channel in which the intensity varies quadratically over the clipped region. In linear space, this corresponds to a Gaussian extrapolation. In 2D images, true Gaussian extrapolations are obtained for circularly shaped regions in which boundary gradients are rotationally symmetric. Other configurations result in asymmetric reconstructions, which are, however, still smooth everywhere. With gradients defined continuously over the image domain, the reconstruction smoothly restores colors in clipped regions (Figure 4, top right).

3.5. Discretization

Our derivation so far has been based on continuous images and gradients. To work with digital images, we discretize the resulting systems in a straightforward fashion, using 4-connected pixel neighborhoods N_p . The boundaries are defined as unclipped pixels with at least one clipped pixel in their neighborhood.

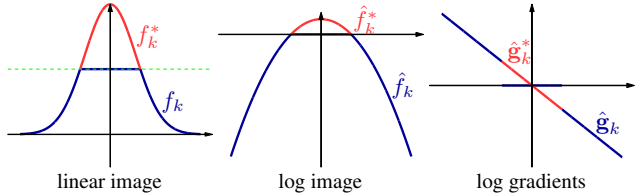


Figure 5: Gaussian infilling. The input signal (blue) is clipped at the clipping threshold (green), resulting in a discontinuous gradient field. Log-space gradient interpolation (red) results in Gaussian infilling of clipped regions.

For gradient estimation we use divided differences over the neighborhoods N_p . The blending weights are first computed independently for each pixel (Equation 8), but since they are applied to gradients estimated over a neighborhood, we low-pass filter the weights over the same neighborhood, using a minimum filter.

4. Results and Analysis

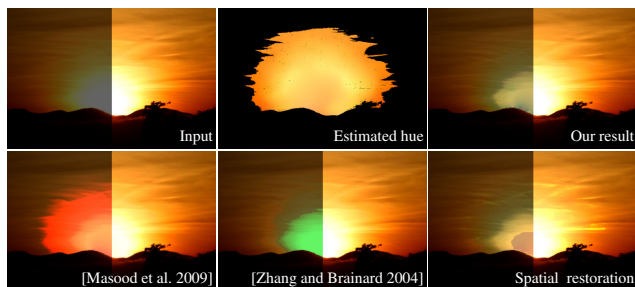


Figure 6: A failure case. Correlation between hue and intensity in the intrinsic image means that the correct hue for the clipped region is not observed anywhere in the image and thus cannot be recovered. The mis-estimation of hue also results in discontinuities between different clipping regions (see text).

We have run our algorithm both on images we captured in RAW mode with different models of Canon SLRs, as well as images obtained from other sources. For the RAW images, we use linear interpolation for demosaicing along the boundaries of the clipped regions and DCRAW [5] for the rest of the image. Images obtained from outside sources are first approximately linearized by applying the inverse of the sRGB gamma curve. Our implementation uses a multi-grid Poisson solver for all subproblems and takes about one minute to solve a 10 megapixel image on an Intel Core 2 Duo machine running at 3GHz.

Figures 4 and 7 show comparisons of our results with [5] and [16], using the respective authors' implementations, and comparisons with [24], using a third party implementation. Figure 4 shows a cropped region of Figure 9(f), depicting flashing police lights. We can see that the spatial methods all generate artifacts at the boundaries between regions with different numbers of clipped channels. Our gradient-

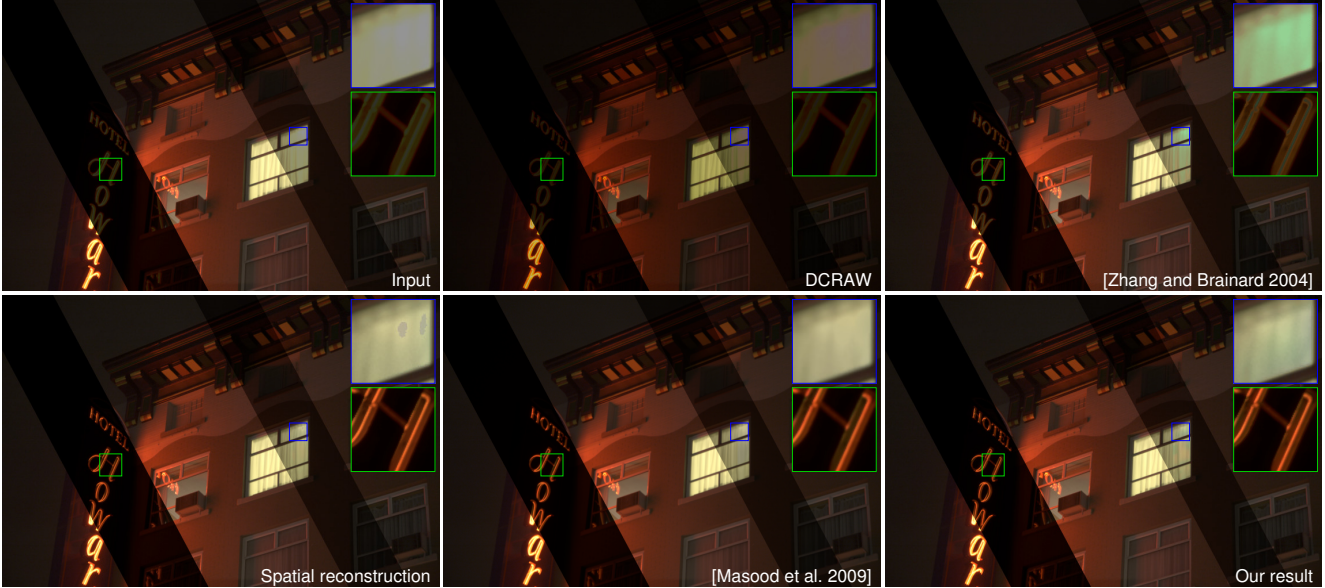


Figure 7: Comparison with other methods. Our method faithfully restores the neon sign (green box) and the curtain (blue box), which are clipped in the input image.



Figure 8: Examples of color restoration with our method. In each pair, the input image is on the left, and our result is on the right, tonemapped with color preservation (see text), showing enhanced brightness and restored scene detail.

based approach avoids these artifacts.

The neon signs in Figure 7 appear yellowish white, although from the reflection on the windows it is evident that the neon signs should be red in color; also note that the upper right corner of the curtain is completely flat due to clipping. DCRAW [5] fails to correct either of these discoloration artifacts. Zhang et al.’s approach [24] reconstructs the curtain well with their global model but fails to reconstruct the neon sign due to the absence of a localized model, which implies that local control is important for such restoration. A spatial reconstruction (Equation 4) restores the neon sign well but shows discontinuity artifacts in the curtain. In this example, the quality of the result by Masood et al. [16] is comparable to ours for both regions.

Figures 8 and 9 contain examples of a variety of scenes including day and night shots, man-made light sources, a sunset scene, and a human face. Since our color restoration produces pixel values outside the 3D gamut of the original image, we choose two different visualizations to il-

lustrate the results. The first is a split-image representation for two different virtual exposures, which is commonly used to visualize HDR images (e.g. [22]). The second is a tone-mapped version of the output using Reinhard’s photographic operator [21] with the color correction from Mantiuk et al. [15]. We emphasize that we consider these representations only as visualizations for print purposes; the full restored color image could also be presented on alternative devices with a larger 3D gamut, could be explored interactively with viewers such as the one provided in the supplemental material, or could simply serve as the input for further manual processing with tools such as Photoshop.

Our method restores scene details washed out due to clipping, including details in the curtain in Figure 1, in the water droplets in Figure 9(c), in the sunny background in Figure 9(d), and on the petals of the skunk cabbage in Figure 8(b). Figure 8(a) shows that the method works well even when unclipped regions with different hues touch. Additional examples are provided in the supplemental material.

One downside of transferring data from unclipped channels is that noise is enhanced when the unclipped channel is very dark. Sunset scenes like Figure 9(h) often have strong red and green components close to the sun but a very small blue component. Since static sensor noise and quantization dominate at small luminance values, when we amplify the unclipped channel, the noise is amplified as well. However, this problem can be alleviated by applying a noise removal step before transferring the gradients.

In Figure 6, we demonstrate a failure case for all existing methods, including our own. In this example, intensity and hue of the intrinsic image are correlated so that the correct hue of the clipped regions is not observed anywhere, and the hue estimation fails. With a mis-estimated hue, the correlation between gradients in the different channels is inconsistently estimated, which results in discontinuities in all methods. However, as the results show, the unclipped channels do provide a lot of information about the cloud structure, and we believe that as future work one could derive subject-specific algorithms to handle such scenes.

5. Conclusion and Future Work

We have presented a novel gradient-space algorithm to restore discoloration artifacts due to clipping. We have demonstrated that our algorithm generates smooth and artifact-free results in many real life situations. We have presented comparisons with recent work and demonstrated the advantages of our gradient-based approach. Since all parts of our algorithm can be cast as simple Poisson problems, the algorithm can be easily implemented and incorporated in modern image processing software.

Our current method assumes that the hue in a region is independent of its intensity, implying that clipped pixels have the same hues as unclipped ones. As we have shown, this is not the case for scenes such as sunsets, where hue and intensity are correlated in a way that cannot be learned from the same image since the same clipping threshold is applied everywhere in the image. However, we believe it should be possible to learn this relationship from other images showing similar scenes. In this way, a collection of short exposure sunset images could be used to fix the colors in our sunset image without altering the specific cloud structure in our image.

References

- [1] J. Abel and J. Smith III. Restoring a clipped signal. In *Proc. Int. Conf. Acoust., Speech, and Signal Process., 1991*, pages 1745–1748, 1991.
- [2] F. Banterle, P. Ledda, K. Debattista, and A. Chalmers. Inverse tone mapping. In *Proc. GRAPHITE '06*, pages 349–356, 2006.
- [3] B. Bayer. Color imaging array, 07 1976. US Patent 3,971,065.
- [4] P. Bhat, L. Zitnick, M. Cohen, and B. Curless. Gradientshop: A gradient-domain optimization framework for image and video filtering. *ACM Trans. Graph.*, 29(2), 2010.
- [5] D. Coffin. <http://www.cybercom.net/~dcoffin/dcrawl/>, 2011. DCRAW - Decoding raw digital photos in Linux.
- [6] P. E. Debevec and J. Malik. Recovering high dynamic range radiance maps from photographs. In *ACM SIGGRAPH 1997*, pages 369–378, 1997.
- [7] P. Diddy, R. Mantiuk, M. Hein, and H. Seidel. Enhancement of bright video features for HDR displays. *Computer Graphics Forum*, 27(4):1265–1274, 2008.
- [8] A. El Gamal. High dynamic range image sensors. In *Tutorial, Int. Solid-State Circuits Conference*, 2002.
- [9] E. Elboher and M. Werman. Recovering color and details of clipped image regions. In *Proc. CGVCVIP*, 2010.
- [10] J. Elder and R. Goldberg. Image editing in the contour domain. *IEEE Trans. PAMI*, 23(3):291–296, 2001.
- [11] A. Foi. Clipped noisy images: Heteroskedastic modeling and practical denoising. *Signal Process.*, 89(12):2609–2629, 2009.
- [12] O. Gallo, N. Gelfand, W. Chen, M. Tico, and K. Pulli. Artifact-free high dynamic range imaging. In *Proc. ICCP*, 2009.
- [13] D. Guo, Y. Cheng, S. Zhuo, and T. Sim. Correcting over-exposure in photographs. In *Proc. CVPR*, pages 515–521. IEEE, 2010.
- [14] A. Levin, D. Lischinski, and Y. Weiss. Colorization using optimization. In *SIGGRAPH '04*, pages 689–694, 2004.
- [15] R. Mantiuk, R. Mantiuk, A. Tomaszewska, and W. Heidrich. Color correction for tone mapping. In *Computer Graphics Forum*, volume 28, pages 193–202, 2009.
- [16] S. Z. Masood, J. Zhu, and M. F. Tappen. Automatic correction of saturated regions in photographs using cross-channel correlation. *Comput. Graph. Forum*, 28(7):1861–1869, 2009.
- [17] L. Meylan, S. Daly, and S. Süssstrunk. The reproduction of specular highlights on high dynamic range displays. In *Proc. CIC*, 2006.
- [18] S. Nayar and S. Narasimhan. Assorted pixels: Multi-sampled imaging with structural models. In *ACM SIGGRAPH'05 Courses*, 2005.
- [19] T. Olofsson. Deconvolution and model-based restoration of clipped ultrasonic signals. *IEEE Trans. Instrumentation and Measurement*, 54(3):1235–1240, 2005.
- [20] P. Perez, M. Gangnet, and A. Blake. Poisson image editing. *ACM Trans. Graph.*, 22(3):313–318, 2003.
- [21] E. Reinhard, M. Stark, P. Shirley, and J. Ferwerda. Photographic tone reproduction for digital images. *ACM Trans. Graph.*, 21(3):267–276, 2002.
- [22] A. Rempel, M. Trentacoste, H. Seetzen, H. 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.
- [23] G. Wetzstein, I. Ihrke, and W. Heidrich. Sensor Saturation in Fourier Multiplexed Imaging. In *Proc. CVPR*, Jun 2010.
- [24] X. Zhang and D. Brainard. Estimation of saturated pixel values in digital color imaging. *JOSA A*, 21(12):2301–2310, 2004.

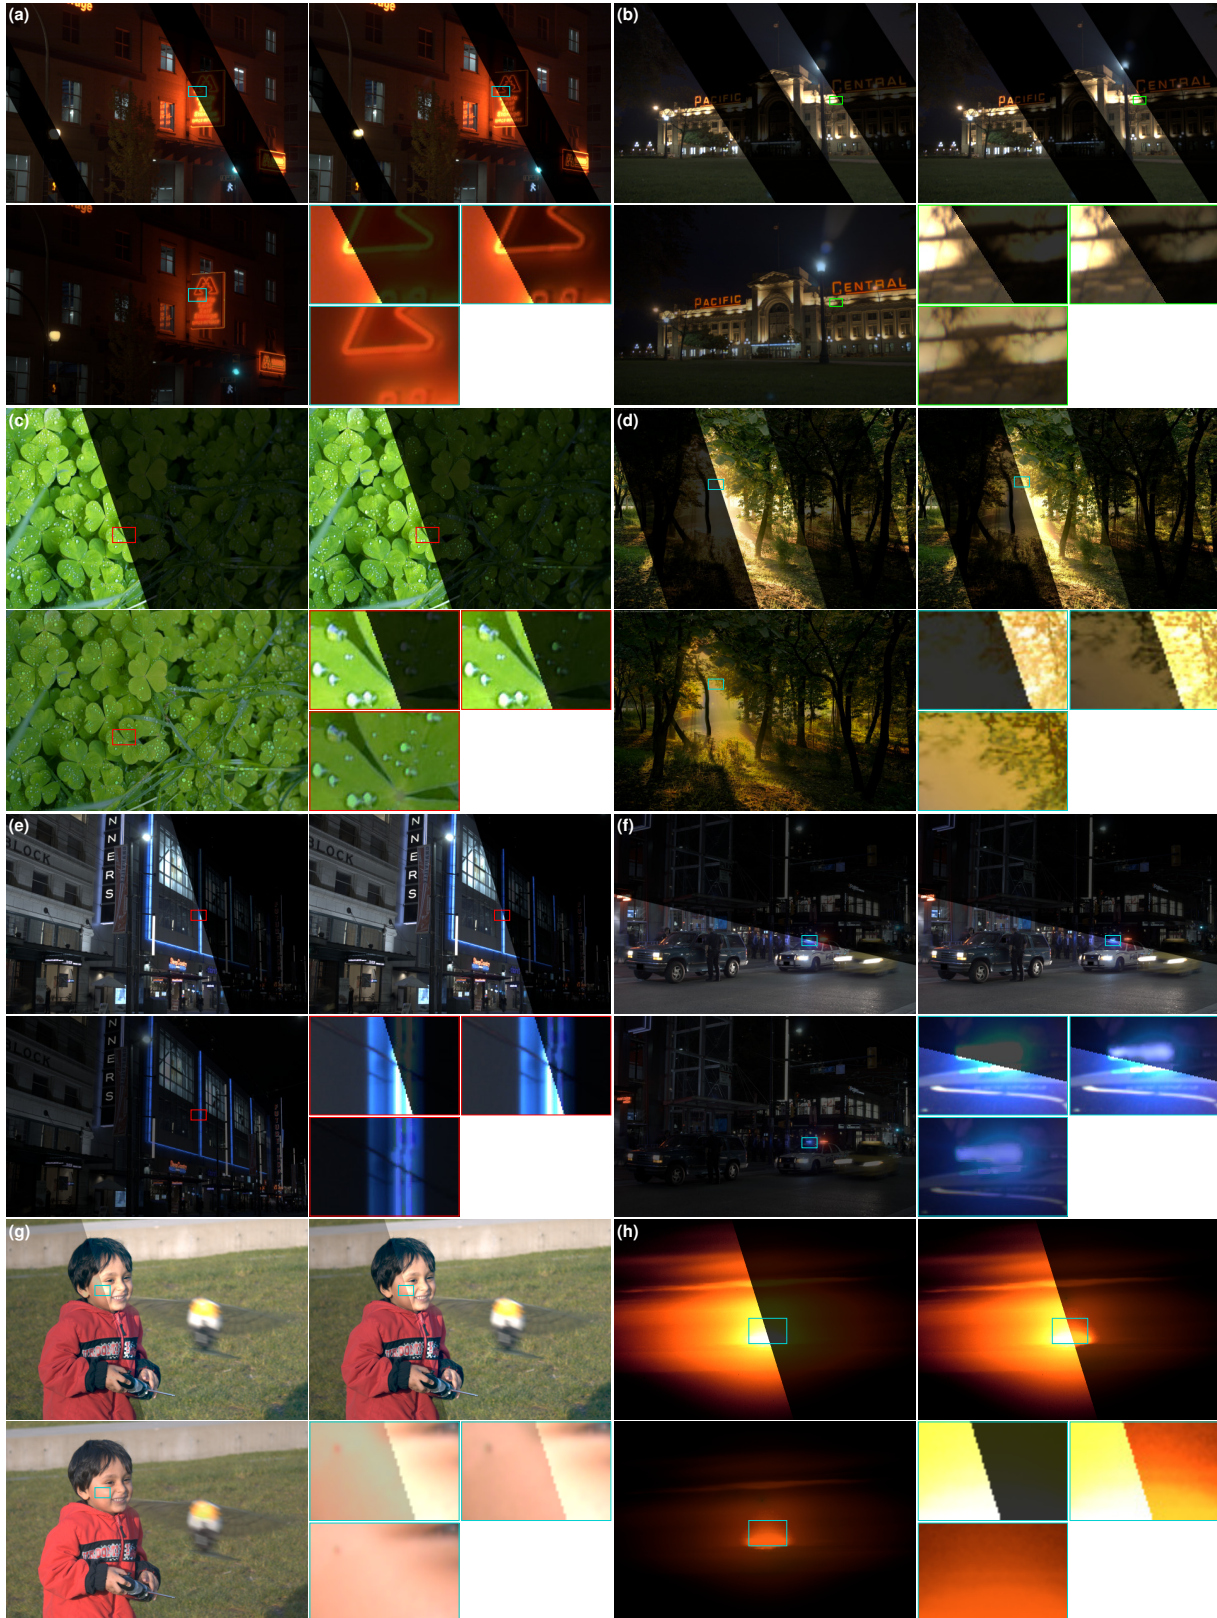


Figure 9: Our algorithm applied to a variety of images. Our method faithfully restores the color, as well as the brightness in most cases, of the clipped image regions. In each group: the input is on the top-left and the result is in the top right, both are shown with split virtual exposures. A tone-mapped image in the lower left shows the restored colors. The lower right shows zoomed regions of all three cases.