

Light Mixture Estimation for Spatially Varying White Balance

Eugene Hsu¹

Tom Mertens²

Sylvain Paris³

Shai Avidan³

Frédo Durand¹

¹Massachusetts Institute of Technology, Computer Science and Artificial Intelligence Laboratory

²Hasselt University – tUL – IBBT, Expertise Centre for Digital Media, Belgium

³Adobe Systems, Inc.



(a) Original image



(b) Traditional white balance



(c) Our result

Figure 1: Image (a) shows a photograph under mixed flash and indoor lighting. Traditional white balance (b) produces unnatural results because it cannot deal with spatially varying light color. Here, the yellow color cast is still visible, and parts of the face have a blue cast. By estimating the relative contribution of lights at each pixel, our technique is able to reproduce colors more faithfully (c).

Abstract

White balance is a crucial step in the photographic pipeline. It ensures the proper rendition of images by eliminating color casts due to differing illuminants. Digital cameras and editing programs provide white balance tools that assume a single type of light per image, such as daylight. However, many photos are taken under mixed lighting. We propose a white balance technique for scenes with two light types that are specified by the user. This covers many typical situations involving indoor/outdoor or flash/ambient light mixtures. Since we work from a single image, the problem is highly under-constrained. Our method recovers a set of dominant material colors which allows us to estimate the local intensity mixture of the two light types. Using this mixture, we can neutralize the light colors and render visually pleasing images. Our method can also be used to achieve post-exposure relighting effects.

Keywords: image processing, computational photography, white balance, color constancy

1 Introduction

The human visual system has the ability to interpret surface colors independently of surrounding illumination. This is known as *color constancy* [Wandell 1995]. In the photographic pipeline, the white balance function emulates this ability to ensure the natural rendition of images. Without proper white balance, many typical scenes show undesirable color casts. For instance, a photograph taken under incandescent lighting will appear unnaturally orange.

All modern cameras and photo editing packages include some functionality to restore natural white balance. These tools generally assume a single illuminant. Unfortunately, many scenes are illuminated by a mixture of lights with different color temperatures. A typical example is a flash photograph in which the foreground subject has a cold blue cast and the background is lit by warmer lights (Figure 1). Classical white balance techniques assume a single light color, yielding unsatisfactory results. To deal with these scenes, professional photographers use gel filters to bring all lights to the same color temperature. This allows them to apply traditional white balance techniques afterwards. While this is an effective approach, it often requires an impractical amount of time and effort.

We present a white balance technique for scenes with two light types that we assume are provided by the user. This two-light scenario is motivated by practical scene configurations involving indoor/outdoor or flash/ambient lighting. Our technique estimates the relative contribution of each light color at each pixel, which we call the mixture. This is a severely ill-posed problem because we have only one image as input. We overcome this difficulty by assuming that the scene is dominated by a small set of material colors. Using this assumption, we propose a new voting scheme and a new interpolation method that retrieves the light mixture at every pixel. We show that this information makes it possible to white balance the input image and also to apply simple relighting effects.

Contributions. To solve the white balance problem for two types of light, our paper introduces the following contributions:

- We present a method to estimate the light mixture per pixel, assuming the light colors are given by the user (Figure 2).
- To achieve this goal, we first evaluate the material colors at a sparse set of pixels. We use a voting scheme that assumes a small set of material colors [Omer and Werman 2004] (§ 4).
- To extend the mixture to the entire image, we formulate white balance as a matting problem in chromaticity space and solve it with the matting Laplacian [Levin et al. 2006] (§ 5).
- We show that recovering the light mixture at each pixel allows for the independent control of light colors and intensities (§ 7).



Figure 2: An overview of our approach. Given an input image (a), we extract dominant material colors using a voting technique. Here, it identifies pixels corresponding to white (b) and brown (c) material colors. This information is used to estimate the local relative contributions of each light (d). The voting scheme only labels reliable pixels. The unreliable pixels, shown in blue, are inferred using an interpolation scheme (e). The mixture information can then be used to compute a visually pleasing white balance (f).

2 Related Work

Many simple techniques have been proposed for white balance and the related problem of color constancy. The white-patch technique assumes that all scenes contain an object with white reflectance, and the gray-world technique assumes that the average color in a scene is neutral [Buchsbaum 1980]. More modern techniques exploit geometric models of color space [Forsyth 1990; Finlayson and Hordley 2000], statistical knowledge about lights and colors [Brainard and Freeman 1997; Finlayson et al. 2001], novel observations about scenes [van de Weijer and Gevers 2005], and natural image statistics [Gijssenij and Gevers 2007]. These are successful on a variety of real-world photographs, but they all assume that only a single illuminant type is present.

An option to extend color constancy methods to mixed lighting is to let users segment images into regions lit by a single type of light. Image editors such as Adobe Photoshop offer selection tools to restrict the spatial extent of color correction filters. Lischinski and colleagues [2006] show a scribble interface that can successfully correct localized color casts. We instead aim for an automatic process with less localized correction, since illumination may affect large disconnected portions of an image.

A few automatic methods have been proposed to handle scenes with mixed lighting. These are the most related to our work. Barnard and colleagues [1997] describe a technique that removes the color variations between different types of lights in the scene before applying a gamut-based color constancy technique [Finlayson 1995]. However, this method employs a classical assumption of smooth illumination [Land and McCann 1971] which is often violated at geometric discontinuities. Kawakami and colleagues [2005] describe a related method for large objects in outdoor scenes. Unlike our technique, they do not require the light colors beforehand. However, this automation is achieved through more restrictive assumptions.

They assume that all shadows are hard and that lights are black-body radiators; in contrast, our technique handles both hard and soft shadows from a greater variety of light source types. Ebner [2004] introduces a method based on a localized gray-world assumption. His method is effective when objects are predominantly gray, but it struggles when large colored objects are introduced. Additionally, his assumption of smooth illumination can result in color halo artifacts, as shown in Figure 8.

3 Image Model

Assumptions. In general form, the white balance problem is highly underconstrained because each pixel corresponds to an unknown illuminant and reflectance spectrum. Our technique focuses on a specific scenario described by the following hypotheses:

- The interaction of light can be described using RGB channels only, instead of requiring full spectra.
- Surfaces are Lambertian and non-fluorescent, so image color is the product of illumination and reflectance in each channel.
- Color bleeding due to indirect illumination can be ignored.
- There are two illuminant types present in the scene and their colors are known beforehand.
- Scenes are dominated by only a small number of material colors. In other words, the set of reflectance spectra is sparse.

The first hypothesis is similar to von Kries chromatic adaptation, but in RGB space. The use of a more advanced color space such as that of Chong and colleagues [2007] might also be fruitful, but we opt for the simpler approach. The next two hypotheses are common among existing white balance techniques for single illuminants.

The last two hypotheses are crucial since they differentiate our approach from prior work. Many mixed lighting configurations contain only two light types; for instance, it is common for indoor photographs to mix flash/ambient or indoor/outdoor lighting. The assumption of few material colors draws from the work of Omer and Werman [2004] in image analysis and the work of Levin and colleagues [2006] in matting.

Image Formation. For a single illuminant, our assumptions yield the image formation model $\mathbf{I} = k\mathbf{R}\mathbf{L}$. Here, \mathbf{I} is a 3×1 vector containing the observed RGB pixel color, \mathbf{R} is a diagonal 3×3 matrix representing the spectral reflectance at some pixel, and \mathbf{L} is a 3×1 vector containing the illuminant color. The scalar k represents factors that influence light intensity such as shadowing, surface orientation, and illuminant power. Note that there is a scale ambiguity between these terms; an object twice as dark lit by a light twice as bright will produce the same result. We can extend this to two illuminants \mathbf{L}_1 and \mathbf{L}_2 with respective intensities k_1 and k_2 :

$$\mathbf{I} = \mathbf{R}(k_1 \mathbf{L}_1 + k_2 \mathbf{L}_2). \quad (1)$$

In our formulation, the spectral reflectance term \mathbf{R} and the intensities k_1 and k_2 are unknown and vary per pixel. The light colors \mathbf{L}_1 and \mathbf{L}_2 are known and constant over the scene.

White Balance. White balance removes color casts induced by colored light sources. It produces an image as if all the lights are white, that is, if $\mathbf{L}_i = \mathbf{1} = [1, 1, 1]^T$ for all i . We can model it by defining a matrix $\mathbf{W} = \text{diag}(W_r, W_g, W_b)$ as follows:

$$\mathbf{W}\mathbf{I} = \mathbf{W}\mathbf{R}(k_1 \mathbf{L}_1 + k_2 \mathbf{L}_2) = \mathbf{R}(k_1 \mathbf{1} + k_2 \mathbf{1}). \quad (2)$$

This yields the following expression for color channels $c \in \{r, g, b\}$:

$$W_c = \frac{k_1 + k_2}{k_1 L_{1c} + k_2 L_{2c}} = \frac{1}{\alpha L_{1c} + (1 - \alpha) L_{2c}} \text{ with } \alpha = \frac{k_1}{k_1 + k_2}. \quad (3)$$

In other words, we can perform white balance without knowing k_1 and k_2 ; only their relative proportions are needed. Thus, our goal is to solve for the mixture α at each pixel.

4 Material Color Estimation

We first recover a representative set of material colors in the scene. This will be used to estimate local light mixtures. Recovering material colors under the influence of unknown spatially varying light mixtures is a highly underconstrained problem. To arrive at a practical method, we follow Omer and Werman [2004]. They observe that natural scenes are dominated by a small set of material colors and apply this assumption to image analysis tasks such as segmentation. We verify this hypothesis with images illuminated by white light. For these images, $\mathbf{I} = k\mathbf{1}\mathbf{R}$, which means that the observed pixel color corresponds to the material color up to a scale factor. As illustrated in Figure 3, it is typical in such images for material chromaticity values to cluster around a few values. But under mixed illumination, the density peaks are smeared out because the measured colors can be any linear combination of two components (Equation 1). Since this combination is arbitrary, no single pixel can yield a reliable estimate of material color. We overcome this difficulty by gathering estimates from as many pixels as possible.

Voting Scheme. We begin by densely sampling the set of possible material colors. Since there is a scale ambiguity in the image formation model (Equation 1), we do not sample albedo values which differ only by a scale factor. We divide this space into a 32×32 grid with logarithmic spacing so that r/b and b/r are similarly sampled, thereby yielding an even distribution of colors.

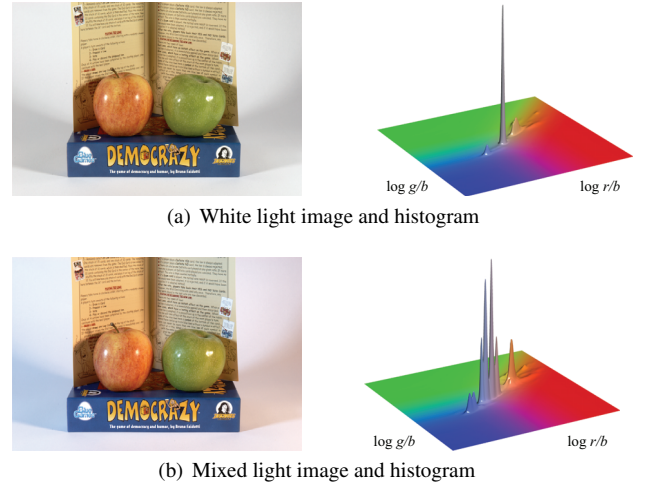


Figure 3: For photographs under white lighting, chromaticity values tend to cluster around a few values (a). Under mixed lighting, these peaks are smeared out lines (b) because the image is a blend of two components.

Given a material color \mathbf{R}_0 and pixel with observed color \mathbf{I} , we can solve Equation 1 because there are two unknowns (k_1 and k_2) and three linear relationships (one per channel). Since k_1 and k_2 are attenuation factors, we clamp them to be non-negative. If the residual $\|\mathbf{I} - \max(0, k_1) \mathbf{R}_0 \mathbf{L}_1 - \max(0, k_2) \mathbf{R}_0 \mathbf{L}_2\|$ is less than some threshold, we say that the pixel *votes* for the material color \mathbf{R}_0 . Intuitively, a pixel votes for a material color only if its observed color can be explained by lighting variations. This voting scheme is similar in spirit to the Hough transform [Hough 1962]. We use a threshold of 0.02 (with color channel values between 0 and 1).

Set Estimation. We greedily process the material colors, starting with the one with the most votes. The pixels that vote for this color are marked. The remaining unmarked pixels then vote for the second most popular material color. This process continues until the number of votes for the top sample is too small. Our implementation stops at 4 percent of the total number of pixels. This process yields a small set of material colors $\{\mathbf{R}_n\}$.

At the end of this step, each marked pixel has voted for a material color which explains its observed color well. A caveat is that some pixels can be accounted for by multiple material colors in the set $\{\mathbf{R}_n\}$. To detect these ambiguous pixels, we test every marked pixel against the material color set. If the residual is small for more than one of the material colors, we unmark the pixel. For the remaining marked pixels, we have a reliable estimate of k_1 and k_2 from which we can compute mixture values α according to Equation 3. The unmarked pixels are considered unreliable and their computed mixture values are discarded. In the following section, we propose an interpolation scheme to fill in these missing values.

Discussion. This technique works well for many typical images, but certain pathological situations may occasionally arise and cause errors. For instance, two pixels with different material colors might have the same observed color due to coincidental lighting conditions. Our technique will mistakenly group these pixels together during voting, but they often end up being unmarked because they cast multiple votes. It is also possible that certain material colors do not exhibit enough mixture variation for voting to work, such as in the case of a uniformly lit wall. However, most scenes contain sufficiently variable lighting due to factors like falloff and shadows.

5 Mixture Interpolation

Our voting scheme yields the mixture of lights for a subset of a pixels. Here, we describe how we extend these values over the entire image using interpolation (in other words, to fill in the blue pixels in Figure 2d). Edge-aware schemes (e.g., [Levin et al. 2004; Lischinski et al. 2006]) are common for these tasks. We instead apply the matting Laplacian [Levin et al. 2006] in chromaticity space for superior results. These techniques are compared in Figure 4.

White Balance as Matting. In this section, we show that the light mixture α can be treated similarly to the foreground/background mixture in matting. But instead of the RGB space used for matting, we use the $[r/b, g/b]$ chromaticity space described in § 4. This could in theory yield instabilities for small values of b , but in our experiments we found that most natural images have nonzero values in all color channels.

Define $\mathbf{I}' = [I_r/I_b, I_g/I_b] = [I'_{rb}, I'_{gb}]$. Then Equation 1 yields the following expression for $I_{cb} \in \{I_{rb}, I_{gb}\}$:

$$I'_{cb} = \frac{I_c}{I_b} = \frac{R_c}{R_b} \left(\frac{k_1 L_{1c} + k_2 L_{2c}}{k_1 L_{1b} + k_2 L_{2b}} \right) = R'_{cb} \left(\frac{k_1 L_{1c} + k_2 L_{2c}}{k_1 L_{1b} + k_2 L_{2b}} \right). \quad (4)$$

We use the scaling degree of freedom to impose $L_{1b} = L_{2b} = 1$ and distribute the material chromaticity \mathbf{R}' :

$$\mathbf{I}' = \alpha \mathbf{R}' \mathbf{L}'_1 + (1 - \alpha) \mathbf{R}' \mathbf{L}'_2. \quad (5)$$

This is the relationship that we seek: image chromaticities \mathbf{I}' are a linear blend of the material chromaticity \mathbf{R}' multiplied by the light chromaticities \mathbf{L}'_1 and \mathbf{L}'_2 . This is similar to the classical foreground/background mixture problem in matting and suggests that a matting technique can be used to interpolate our data.

The Matting Laplacian. Levin and colleagues employ the *color line model* in their matting algorithm [2006]. A set of pixels follows this model if their colors lie on a line in RGB space. By assuming that both the foreground and background of an image are locally in this configuration, they demonstrate that the pixel opacities β can be obtained by minimizing the quadratic $\beta^T \mathbf{M} \beta$, where \mathbf{M} is the *matting Laplacian*. Element ij of \mathbf{M} is defined by considering pairs of pixels i and j in windows w_k centered around pixels k :

$$\sum_{k | ij \in w_k} \left[\delta_{ij} - \frac{1}{|w_k|} (\mathbf{I}_i - \mu_k) \left[\Sigma_k + \frac{\varepsilon}{|w_k|} \mathbf{E}_3 \right]^{-1} (\mathbf{I}_j - \mu_k) \right]. \quad (6)$$

In this expression, \mathbf{I}_i and \mathbf{I}_j are the colors at pixels i and j ; δ_{ij} is 1 if $i = j$ and 0 otherwise, μ_k and Σ_k are the mean and variance of pixel colors in window w_k , \mathbf{E}_3 is the 3×3 identity matrix, ε is a regularizing constant and $|w_k|$ is the number of neighborhood pixels. See the article of Levin and colleagues for a complete derivation.

For white balance, Equation 5 shows that the chromaticity image \mathbf{I}' is a linear combination of $\mathbf{R}' \mathbf{L}'_1$ and $\mathbf{R}' \mathbf{L}'_2$ controlled by the light mixture α . Thus, if we assume that $\mathbf{R}' \mathbf{L}'_1$ and $\mathbf{R}' \mathbf{L}'_2$ locally lie on a line in chromaticity space, the matting Laplacian can be adapted to solve for α . We define element ij of matrix \mathbf{M}' as follows (using notation from Equation 6 for two dimensions instead of three):

$$\sum_{k | ij \in w_k} \left[\delta_{ij} - \frac{1}{|w_k|} (\mathbf{I}'_i - \mu'_k) \left[\Sigma'_k + \frac{\varepsilon}{|w_k|} \mathbf{E}_2 \right]^{-1} (\mathbf{I}'_j - \mu'_k) \right]. \quad (7)$$

Intuitively, the color line model in chromaticity space can be interpreted as follows. Since \mathbf{L}'_1 and \mathbf{L}'_2 are constant with respect to the scene, it follows that values of \mathbf{R}' also locally form a line. This in

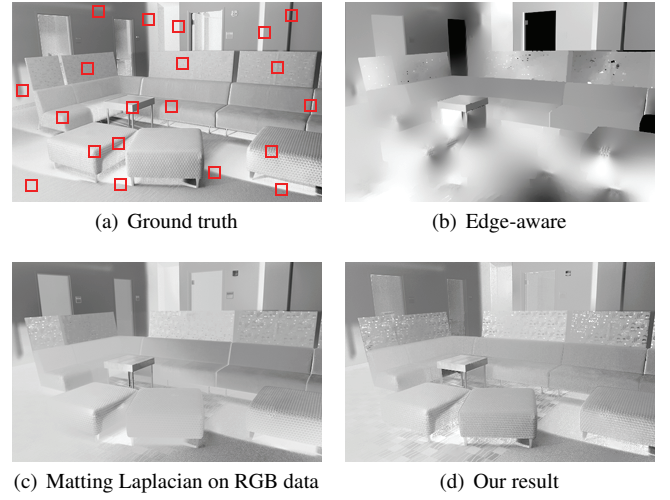


Figure 4: Given a ground truth image (a), we constrain a few points (b) and reconstruct the remainder of the image through interpolation. Edge-aware interpolation gives poor results (b). The matting Laplacian in RGB space does much better (c); applying it in chromaticity space further enhances the recovered data (d).

turn implies that values of \mathbf{R} locally form a plane containing the origin. To summarize, Equation 7 provides a closed-form solution for interpolating light mixtures under the assumption that material colors locally follow a *color plane model*.

Solving the Mixture. The voting scheme in Section 4 provides a set of mixture constraints which we gather into a vector α^* . We use this in a data term to guide the interpolation process. The final objective function that we minimize is:

$$J(\alpha) = \alpha^T \mathbf{M}' \alpha + \lambda (\alpha - \alpha^*)^T \mathbf{D} (\alpha - \alpha^*). \quad (8)$$

Here, \mathbf{D} is a diagonal matrix that selects the marked pixels from the voting step; that is, element d_{ii} is 1 if pixel i has a reliable mixture estimate, and 0 otherwise. The scalar λ balances smoothness of interpolation and fidelity to the data term. This quadratic function can be minimized using a standard sparse matrix solver. In our implementation, we use $\lambda = 10^{-2}$ and $\varepsilon = 10^{-6}$ for the regularization in Equation 7. The particular choice for λ yields soft constraints on α , thereby letting the Laplacian regularize the results. This is useful for fixing small errors made by our voting scheme caused by missing colors or ambiguous pixels.

6 Application

The previous steps provide us with a way to estimate the relative light contributions α at each pixel of a single input image. This can be used to perform white balance using Equation 2. In this section, we discuss several practical issues that arise when applying our technique.

Input. Our image formation model assumes linear image values. For processed files such as JPEGs, it is necessary to achieve linear gamma by inverting the transformations applied by the camera or software. However, this process may introduce artifacts due to quantization and pixel saturation. The best results are achieved from images that are inherently linear, such as those obtained directly from digital camera raw files.

Light Colors. Our approach assumes that the two light colors in the scene are known. Here, we propose a simple interface to estimate these colors. Many scenes contain objects with neutral material colors. We ask the user to click on several of these objects. From our image formation model, we know that their pixel values are of the form $k_1 \mathbf{L}_1 + k_2 \mathbf{L}_2$ with $k_1 \geq 0$ and $k_2 \geq 0$. We choose the pixels with the most extreme color casts and assume that they correspond to $(k_1, k_2) = (0, 1)$ and $(k_1, k_2) = (1, 0)$. This allows us to solve for the light colors up to a scale factor. A typical usage scenario is shown in Figure 6.

Unfortunately, not all scenes contain neutral objects. For these situations, a user can borrow measurements from other photos taken under identical lighting conditions. As with traditional white balance techniques, the same calibration can often be reused for multiple photos. However, it is important to note that our technique does not perform well if light colors are incorrectly specified.

Scalability. Both the voting and the interpolation schemes are computationally intensive and not tractable at full resolution. We overcome this problem by working on downsampled images with larger dimension 640. We upsample the light mixtures using joint bilateral upsampling [Kopf et al. 2007] and perform white balance by applying Equation 2 at full resolution.

Light Filtering. With the white-balanced image \mathbf{WI} and the light mixture α , we can recover the images corresponding to each of the two lights separately:

$$\alpha \mathbf{WI} = \left(\frac{k_1}{k_1 + k_2} \right) \mathbf{R} (k_1 \mathbf{1} + k_2 \mathbf{1}) = k_1 \mathbf{R}. \quad (9)$$

In other words, $\alpha \mathbf{WI}$ is a white-balanced image of the scene lit only by the first light. Equivalently, $(1 - \alpha) \mathbf{WI}$ corresponds to the second light. With these two components at hand, we can apply any effect to either of the lights separately, as shown in Figure 5.

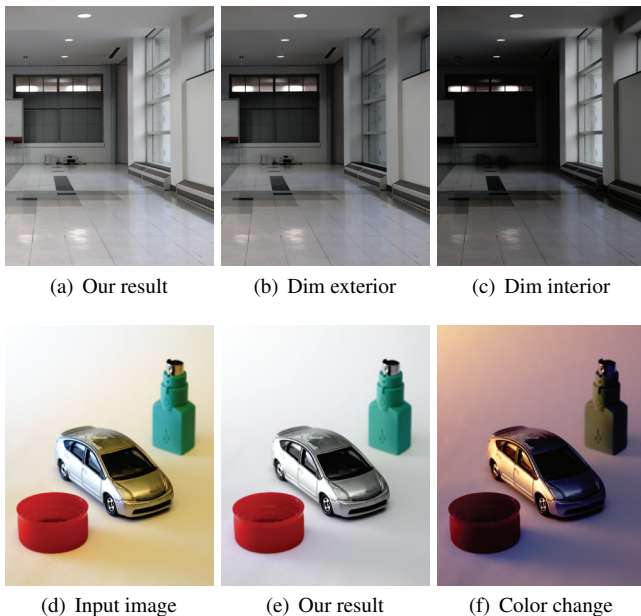


Figure 5: We can achieve additional postprocessing effects. Using the mixture, lights can be dimmed (b,c) and colored (f) separately.




Figure 6: To specify lights, a user clicks on parts of an image with neutral material color (a). We interpret the two most extreme color casts as the light colors. After this step, our technique requires no more user interaction to produce high-quality results (b).

7 Results

We have applied our method to a variety of input images. Our technique significantly reduces the color casts introduced by mixed lighting and produces visually pleasing images. In general, we found that our method yields its most successful results in indoor environments, as shown in Figure 8. Such scenes often satisfy the assumption of sparse material color.

In Figure 7, we validate our technique on two synthetic examples. Here, the input images are generated by combining two images of the same scene taken under different single-light conditions. This enables the computation of ground truth mixture values and white balance results. On the first set of images, the highest errors appear on the left apple (which has a dense mixture of material colors) and in the fold of the paper (where there is some color bleeding). These regions violate the hypotheses of our image formation model. Nevertheless, the output image is not visually objectionable despite the presence of numerical error.

On the second set of images in Figure 7, the highest errors appear on the stuffed toy. This is due to incorrect estimation of the green material color during the voting process. Even so, our method removes the unsightly color variation from the original input. We show that the numerical difference from ground truth, while nonzero, is roughly constant over the toy. Also note how our technique removes the coloration in the specular highlights, even though these specularities violate our image model.

8 Discussion

Although the assumptions of our image model may be violated in some situations, we have shown that our method is robust for a variety of typical images. Our technique produces pleasing results even in the presence of specularities and interreflections. However, it is not foolproof since it needs to see the same material color under different lighting mixtures. This may not hold in certain photographic situations. For instance, if all light sources very diffuse, the mixture α may be nearly constant over the scene; this presents a highly ambiguous situation to the voting scheme. Scenes that exhibit a strong foreground-background separation may also cause problems if there is no mixing of the lights (that is, if α is a binary mask). Fortunately, we have found that these cases are rare in practice.

Our technique requires a small amount of user interaction to specify light sources. The required effort is negligible when processing several images. However, for large batches of photos, it is still desirable to achieve full automation. A potential solution might be to constrain the possible illuminants, perhaps to a canonical set [Finlayson et al. 2001] or to blackbody radiators [Kawakami et al. 2005]. We hope to address this problem in future work.

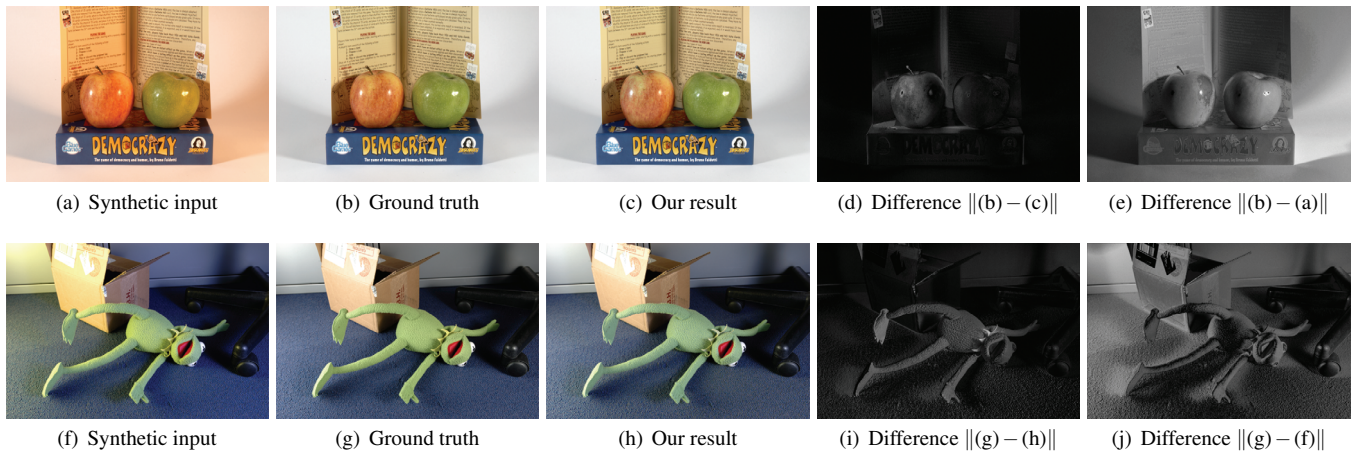


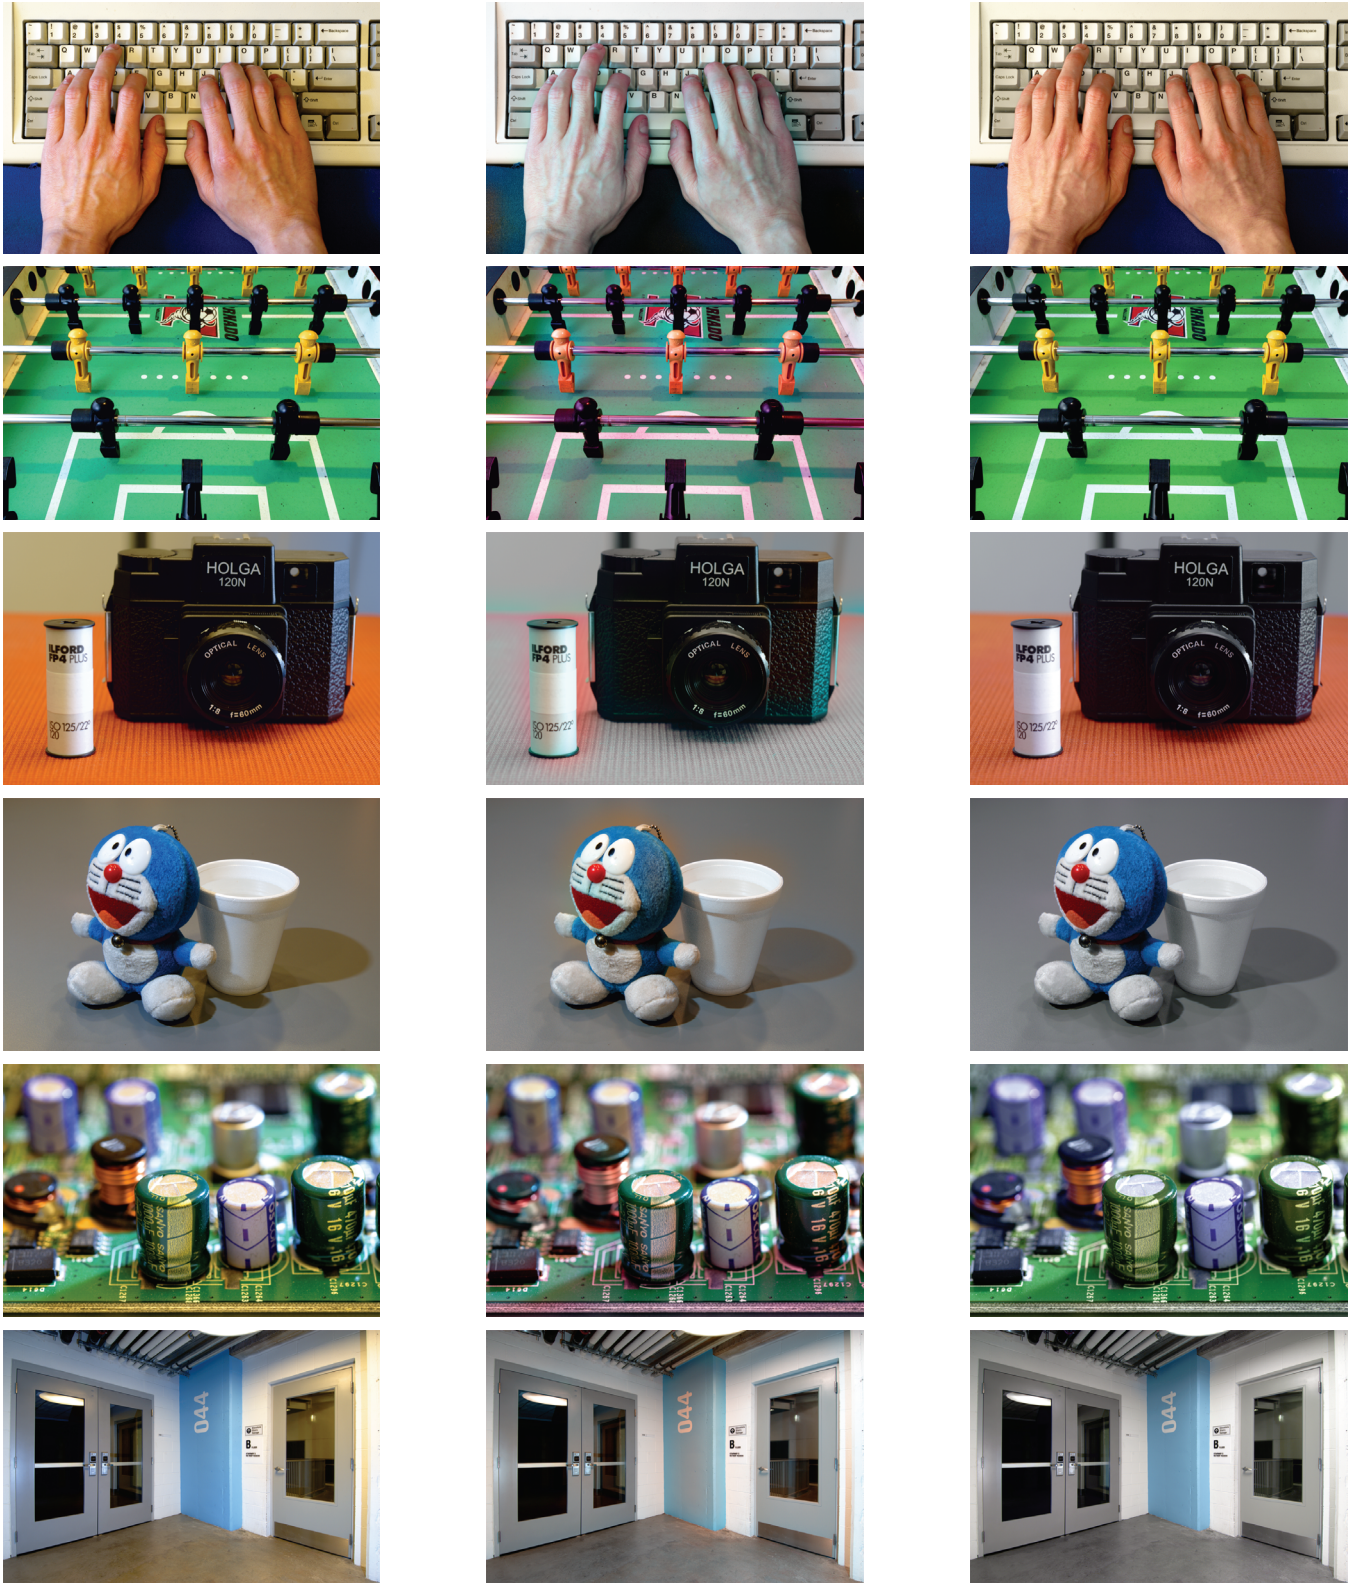
Figure 7: For evaluation, we build synthetic input images (a,f) by combining two exposures under single illuminants. This allows us to compare ground truth results (b,g) to those achieved by our method (c,h). The difference images (d,i) show various sources of error in our technique. For comparison, we also show the difference images between the input and the ground truth images (e,j). Although our results do not match ground truth perfectly, the output of our technique is preferable to the input.

Acknowledgements

We would like to thank the Adobe Creative Technology Labs, the MIT Computer Graphics Group, and Bill Freeman for their helpful advice. Frédo Durand acknowledges a Microsoft Research New Faculty Fellowship and a Sloan Fellowship. Tom Mertens acknowledges a research fellowship from the Belgian American Educational Foundation. Part of the research at Expertise Centre for Digital Media is funded by the European Regional Development Fund and the Flemish government.

References

- BARNARD, K., FINLAYSON, G. D., AND FUNT, B. V. 1997. Color constancy for scenes with varying illumination. *Computer Vision and Image Understanding* 65, 2 (Mar.), 311–321.
- BRAINARD, D. H., AND FREEMAN, W. T. 1997. Bayesian color constancy. *Journal of the Optical Society of America* 14, 7 (July), 1393–1411.
- BUCHSBAUM, G. 1980. A spatial processor model for object colour perception. *Journal of The Franklin Institute* 310, 1 (July), 1–26.
- CHONG, H., GORTLER, S., AND ZICKLER, T. 2007. The von Kries hypothesis and a basis for color constancy. In *IEEE International Conference on Computer Vision*, 1–8.
- EBNER, M. 2004. Color constancy using local color shifts. In *European Conference on Computer Vision*, 276–287.
- FINLAYSON, G. D., AND HORDLEY, S. D. 2000. Improving gamut mapping color constancy. *IEEE Transactions on Image Processing* 9, 10 (Oct.), 1774–1783.
- FINLAYSON, G. D., HORDLEY, S. D., AND HUBEL, P. M. 2001. Color by correlation: A simple, unifying framework for color constancy. *IEEE Transactions on Pattern Analysis and Machine Intelligence* 23, 11 (Nov.), 1209–1221.
- FINLAYSON, G. D. 1995. Color constancy in diagonal chromaticity space. In *IEEE International Conference on Computer Vision*, 218–223.
- FORSYTH, D. A. 1990. A novel algorithm for color constancy. *International Journal of Computer Vision* 5, 1 (Aug.), 5–36.
- GIJSENII, A., AND GEVERS, T. 2007. Color constancy using natural image statistics. In *IEEE Computer Vision and Pattern Recognition*, 1–8.
- HOUGH, P. V. C., 1962. Method and means of recognizing complex patterns. U.S. Patent 3,069,654.
- KAWAKAMI, R., IKEUCHI, K., AND TAN, R. T. 2005. Consistent surface color for texturing large objects in outdoor scenes. In *IEEE International Conference on Computer Vision*, 1200–1207.
- KOPF, J., COHEN, M. F., LISCHINSKI, D., AND UYTENDAELE, M. 2007. Joint bilateral upsampling. *ACM Transactions on Graphics* 26, 3 (July), 96:1–96:5.
- LAND, E. H., AND MCCANN, J. J. 1971. Lightness and Retinex theory. *Journal of the Optical Society of America* 61, 1 (Jan.), 1–11.
- LEVIN, A., LISCHINSKI, D., AND WEISS, Y. 2004. Colorization using optimization. *ACM Transactions on Graphics* 23, 3 (Aug.), 689–694.
- LEVIN, A., LISCHINSKI, D., AND WEISS, Y. 2006. A closed form solution to natural image matting. In *IEEE Computer Vision and Pattern Recognition*, 61–68.
- LISCHINSKI, D., FARBMAN, Z., UYTENDAELE, M., AND SZELISKI, R. 2006. Interactive local adjustment of tonal values. *ACM Transactions on Graphics* 25, 3 (July), 646–653.
- OMER, I., AND WERMAN, M. 2004. Color lines: Image specific color representation. In *Computer Vision and Pattern Recognition*, 946–953.
- VAN DE WEIJER, J., AND GEVERS, T. 2005. Color constancy based on the grey-edge hypothesis. In *IEEE International Conference on Image Processing*, 722–725.
- WANDELL, B. A. 1995. *Foundations of Vision*. Sinauer Associates, Sunderland, MA.



(a) Average white balance

(b) Local Color Shifts [Ebner 2004]

(c) Our result

Figure 8: White balance comparisons. In column (a), we show results of global white balance using the average light source color. Here, the color casts are never fully removed, especially in shadow regions. Column (b) shows results from Ebner’s method [2004] which tends to produce colored halos because it assumes smooth illumination. The last column (c) shows our results. By comparison, our method is able to restore the original colors more faithfully and produce a more appealing result. Notice that there are no color casts left, even in shadow areas.