Eurographics Symposium on Rendering 2015
J. Lehtinen and D. Nowrouzezahrai
(Guest Editors)

Volume 34 (2015), Number 4

Radiometric Transfer: Example-based Radiometric Linearization of Photographs

Han Li and Pieter Peers

College of William & Mary

Abstract

We present an example-based approach for radiometrically linearizing photographs that takes as input a radiometrically linear exemplar image and a target regular *uncalibrated image of the same scene, possibly from a different viewpoint and/or under different lighting. The output of our method is a radiometrically linearized version of the target image. Modeling the change in appearance of a small image patch seen from a different viewpoint and/or under different lighting as a linear 1D subspace, allows us to recast radiometric transfer in a form similar to classic radiometric calibration from exposure stacks. The resulting radiometric transfer method is lightweight and easy to implement. We demonstrate the accuracy and validity of our method on a variety of scenes.*

Categories and Subject Descriptors (according to ACM CCS): I.3.8 [Computer Graphics]: Applications— I.4.1 [Image Processing and Computer Vision]: Digitization and Image Capture—Radiometry

1. Introduction

The nonlinear relation between recorded pixel brightness in a photograph and the corresponding observed scene radiance is described by the camera-dependent *camera response function*. Recovering the camera response function (i.e., radiometric calibration) and undoing its effect (i.e., radiometric linearization) are critical and important preprocessing steps to many computer graphics and computer vision algorithms such as photometric stereo and reflectance and illumination estimation.

Most radiometric calibration methods focus on recovering the camera response function from an exposure stack of a static scene recorded from a fixed viewpoint and under static lighting conditions. However, such calibration methods are not suited for consumer cameras that offer limited control on exposure or that employ an image-content dependent camera response function. Similarly, such methods are also not suited when direct access to the camera is not possible (e.g., internet photographs). However, one usually has access to (or the opportunity to record) a radiometrically linearized image of the same scene, albeit recorded with a different camera and/or under different lighting or viewing conditions. This potentially rich source of information has not been considered for radiometric calibration, and it raises the question of whether it is possible to *transfer* radiometric information from a radiometric linear image to an uncalibrated photograph of the same scene.

In this paper, we propose a novel example-based radiometric linearization method that takes as input a radiometrically linear photograph of a scene (i.e., exemplar), and a standard (radiometrically uncalibrated) image of the same scene potentially from a different viewpoint and/or under different lighting, and which produces a radiometrically linear version of the latter. Key to our method is the observation that under modest assumptions, the change in appearance of a small local pixel neighborhood in a photograph resides in a 1D linear subspace. This allows us to formulate a fast and lightweight solution that resembles the classic solution for radiometric calibration from exposure stacks.

We demonstrate the qualitative accuracy of our method on a variety of different scenes, and quantitatively validate the robustness and accuracy of various components of our system.

2. Related Work

There exists a large body of prior work on various forms of radiometric calibration.

Probably the most common approach to radiometric cali-

bration estimates the camera response function from an exposure stack of a static scene (captured with a static camera). Debevec and Malik [DM97] estimate a non-parametric camera response model regularized by a smoothness constraint. Mitsunaga and Nayar [MN99] use a flexible polynomial model and only require approximate estimates of the exposures ratios. Grossberg and Nayar [GN04] propose a data-driven model based on a large database of measured camera response functions. However, these methods are limited to static cameras, and several strategies have been proposed to address this to some degree. Grossberg and Nayar [GN03] use intensity histograms and derive a brightness transfer function between two photographs captured with different exposures. Kim and Pollefeys [KP08] exploit epipolar geometry and stereo matching to find corresponding points. In follow up work, Kim et al. [KGFP10] integrate tracking and camera response function recovery to further improve the quality of both. While our input is significantly different, we show that radiometric transfer can be reformulated in a form similar to the above problems, and hence similar solution strategies can be employed.

The recent interest in large photo collections necessitates radiometric calibration methods that can work not only for non-static cameras, but also for non-static scenes and in particular for uncontrolled lighting. Shaque and Shah [SS04] recover the camera response function, modeled by a gamma curve, for fixed-viewpoint photographs under different uncontrolled lighting. Kim and Pollefeys [KFP08] perform radiometric (and vignetting) calibration for static-viewpoint images of outdoor scenes under changing illumination by grouping pixels with similar behavior with respect to changes in illumination. Diaz and Strum [DS11a, DS11b] recover the camera response functions for images in large photo collections by assuming Lambertian surface reflectance and directional or low frequency lighting respectively. They use an inverse rendering approach and leverage the geometry obtained from multiview stereo [FP10]. Kuthirummal et al. [KAGN08] establish prior statistics for large photo collections for recovering the radiometric camera properties. All of these methods require many images of the same scene to either recover the geometry or to exploit statistical properties over many images. In contrast, the proposed radiometric transfer technique only requires a single radiometrically linear exemplar image.

A different strategy is to exploit statistical properties of photographs and cameras to recover the camera response function. Lin et al. [LGYS04] and Lin and Zhang [LZ05] rely on transformations of the linear radiance distribution at edges. Matsushite and Lin [ML07] exploit the symmetry of camera noise distributions. Takamatsu et al. [TMI08] infer camera response functions by looking at noise variances. However, these single image calibration methods are sensitive to deviations from the assumed statistical model. The proposed radiometric transfer technique is complementary to the above calibration methods, and can be used to combine and enhance their results.

Our method is also related to radiometric calibration methods that rely on a calibration object such as a color checker card [CR96], except that in our solution the scene itself acts as the calibration object and no restrictions are imposed on the lighting.

At the heart of our method is the observation that the variations in appearance between corresponding patches can be modeled by a 1D linear subspace. Low-rank structures have been explored in the context of radiometric calibration before. Lee at el. [LMS*13] propose a general minimization framework for radiometric calibration based on low-rank structures. Oh et al. [OLTK15] build on this and propose to utilize low-rank structures to improve the reconstruction of high dynamic range images (assuming a linearized exposure stack).

At a high level, the proposed method also shares similarities with other methods that *transfer* artistic properties between photographs such as color (e.g., [RAGS01, HSGL13]) or style (e.g., [SPB*14]) by establishing a local or global color transformation based on sparse feature matching. However, instead of transferring artistic properties, the proposed method seeks to revert camera-driven intensity transformations to recover the intrinsic radiometric linear image while retaining scene dependent qualities such as lighting and view-dependent effects.

3. Radiometric Transfer

Problem Statement The image formation process in a digital camera can be abstracted as: $I = f(E)$, where I is the resulting image and E is an image that is linearly proportional to the time-average of the incident radiance on the camera's sensor. The exact scale between incident radiance and E depends on various camera characteristics such as exposure time, aperture, light efficiency of the sensor and lenses, etc. The camera response function f, a non-linear mapping between E and the image I, is designed to remap and compress the dynamic range in order to produce visually pleasing images. The goal of radiometric calibration is to undo the effects of f and to recover a radiometrically linear image that is proportional (up to an unknown scale factor) to the time-average image irradiance E.

Radiometric transfer takes as input two images; a radiometrically linear source image E_s, and a regular target image $I_t = f(E_t)$. Both images depict the same subject, but viewed from different viewpoints and under different lighting conditions. The goal of radiometric transfer is to recover $E_t = f^{-1}(I_t)$ by exploiting the knowledge that E_s depicts the same subject.

Fixed Viewpoint We will first consider the case where the viewpoint is the same for both images, but the lighting conditions can differ. In this case, the relation between the known

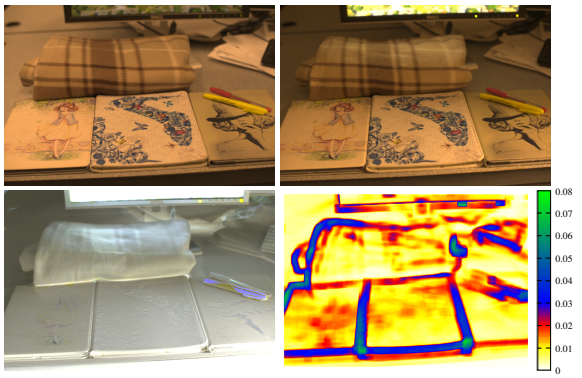

Figure 1: ***Top Row:*** *two radiometrically linear photographs captured under different lighting conditions (office lighting and LCD-panel illumination respectively).* ***Bottom Left:*** *ratio $\kappa(x)$ of the input images.* ***Bottom Right:*** *False color plot of the error on a 1D linear subspace approximation ($\kappa(x) \approx \kappa_x$) for a 33×33 window around each pixel location x.*

source irradiance E_s and the unknown target irradiance E_t can be expressed by their ratio: $E_t(x) = \kappa(x)E_s(x)$. The ratio κ can potentially vary with pixel position x due to changes in the underlying surface normal, material properties, and/or angular variation in lighting (Figure 1, bottom-left). This makes it difficult to directly estimate κ for every pixel from E_s and I_t only. However, we observe that κ is locally slowly varying: $E_t(x') \approx \kappa_x E_s(x')$ for $x' \in \mathcal{N}(x)$, a small neighborhood around x. Hence, we can approximate the ratio $\kappa(x)$ by a constant ratio κ_x for a small neighborhood $\mathcal{N}(x)$ around x:

$$f^{-1}(I_t(x')) \approx E_s(x')\kappa_x, \quad x' \in \mathcal{N}(x). \tag{1}$$

In other words, the appearance space of a small neighborhood of pixels, can be well approximated by a 1D linear subspace, as illustrated in Figure 1, bottom-right. Equation (1) is similar in form to that of classic radiometric calibration from multiple exposures [DM97, KP04], except that in our case the "exposures" (κ_x) are unknown instead of the irradiance image (E_s). Similar as in prior work, we reformulate this expression in the log domain:

$$g(I_t(x')) \approx \log E_s(x') + \log \kappa_x, \tag{2}$$

where $g = \log(f^{-1})$. We characterize the log-inverse of the camera response function using the log-PCA model of Kim and Pollefeys [KP04]:

$$h'_0(I_t(x')) + \sum_{i=1}^{n} c_i h'_i(I_t(x')) = \log E_s(x') + \log \kappa_x, \tag{3}$$

where $n = 25$ is the number of log-PCA terms, h'_i represents the i-th PCA term of the space spanned by the log-inverse camera response functions contained in the DoRF database [GN04], and h'_0 is the mean log-inverse camera response function. While $n = 3$ log-PCA terms already explain 99.6% of the energy [KP04], we use 25 terms to better model irregular and uncommon camera response functions. Each $x' \in \mathcal{N}(x)$ in Equation (3) provides a linear equation in terms of the unknowns $\log \kappa_x$ and the $n = 25$ log-PCA coefficients c_i. Combining the linear equations in a single system, and assuming sufficient variety in pixel values in each patch, allows us to solve for the $n+1$ unknowns using a linear least squares solver. To improve the stability and to ensure sufficient coverage of the available pixel range, we consider the neighborhoods around m different pixel location x_k, $k \in \{0..m-1\}$, and solve for the $n+m$ unknowns: the $n = 25$ coefficients and the m different $\log \kappa_{x_k}$ scale factors (one for each patch).

The final camera response function can then be computed by exponentiating and inverting the obtained function: $f = (\exp g)^{-1}$. However, this is only a *partial* camera response function since the full range of pixel values might not be covered in the the m patches (or even in the target image I). Furthermore, there exists an ambiguity between the partial camera response function and the scale factors κ: $(g+\gamma) - (\log \kappa + \gamma) = \log E_s$ for any γ. Hence, the partial camera response function is only determined up to an unknown scale factor. To expand the range of the recovered partial camera response function, we linearly extrapolate the camera response function below the recovered lower limit of the range to the origin. However, due to the unknown scale factor, we cannot extrapolate beyond the upper limit, and simply cut of the response function at the upper limit of the range.

In some sense, our solution can be seen as inferring the camera response function and "exposures" (i.e., scale factors) from a set of m tiny "image-pairs" (i.e., patches).

Different Viewpoint The above algorithm easily extends to the case where the viewpoint between the source E_s and target image I_t (and thus E_t) differ by introducing an additional function $\phi_{t\to s}$ that maps pixels in a target patch to the corresponding pixels in the source patch. We can then reformulate the local approximation: $E_t(x') \approx \kappa_x E_s(\phi_{t\to s}(x'))$ for $x' \in \mathcal{N}_t(x)$, a small neighborhood around x in the target image, which can be solved using a similar strategy as for the fixed viewpoint.

The direction of the mapping (from target coordinates to source coordinates) is critical, since such a mapping will remap integer pixel coordinates to fractional coordinates, requiring an interpolation to obtain the corresponding pixel intensity. Consequently, such a warping operation will only be correct when executed on radiometrically linear pixel intensities (i.e., the source image).

Patch Selection and Mapping A critical component in the above algorithm is the selection of the patches and the corresponding mapping functions. We desire patches with a rich variation in pixel values that can be reliably corresponded between the source and target images. We propose to use

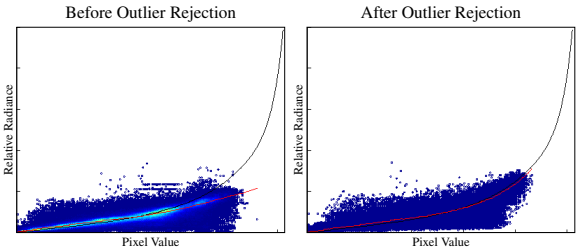

Figure 2: ***Outlier Rejection.*** *Recovered camera response function (red) with corresponding reference camera response function (black) from the* Desk *example shown in Figure 1 before and after outlier rejection. The false color cloud represents the recovered pixel radiance for each pixel in the patches (times the corresponding patch scale factor).*

33×33 pixel-neighborhoods (approximately 1% of the image resolution) around the 200 best matching SIFT correspondences [Low04]. Such correspondences are naturally selected in areas of rich texture, ensuring a rich variety in pixel values in the selected patches. We found that a 33×33 window offers a good balance between providing sufficient linear equations (Equation (3)), providing sufficient overlap between the different patches (in terms of pixel values) to tie the different scale factors together, and minimizing the error introduced by the 1D linear subspace approximation (i.e., larger patches exhibit larger approximation errors). Furthermore, we allow different 33×33 windows to overlap, and include each in the linear system (Equation (3)) with their respective $\log \kappa_{x_k}$ factor.

To obtain a subpixel accurate alignment and to compensate for any non-linear mapping between the source and target patch, we compute SIFT flow [LYT*08] between double sized (i.e., 65×65) windows in the source and target image, and warp the source patch to the target patch. We reject cases for which the inner 33×33 neighborhood contains pixels for which SIFT flow failed to find a corresponding source pixel. We employ a 65×65 initial window to support shifting (to compensate for misalignments) and scaling (due to differences in camera distance) of pixel values from outside the targeted 33×33 window.

Outlier Rejection The above algorithm assumes that the change in pixel values in a small patch can be explained by a 1D linear subspace. However, this is not always the case (see Section 5). Patches that cannot be well represented by the proposed model can adversely affect the quality of the recovered partial camera response function. To remove such outliers, we employ the following two-step strategy. Initially, we compute a candidate camera response function using all patches. We then compute for each patch the fitness of the proposed camera response function:

$$\varepsilon^2(x) = \sum_{x' \in \mathcal{N}_t(x)} \Big(h'_0(I_t(x') + \sum_{i=1}^{n} c_i h'(I_t(x') - \log E_s(\phi_{t \to s}(x')) - \log \kappa_x) \Big)^2. \quad (4)$$

Next, we reject any patches for which its corresponding fitness $\varepsilon^2(x)$ exceeds v times the variance, where v ranges from 2 to 3 depending on how conservative we want outlier rejection to be. Finally, we recompute the camera response function using only the inliers. The key assumption is that the inliers outnumber the outliers, and thus that the initial camera response function can serve as an indicator whether a patch follows the model. Figure 2 shows an example of a recovered camera response function (red) compared to a reference camera response function (black) before and after outlier rejection for the *Desk* example shown in Figure 1. Because we can only recover the partial camera response function up to an unknown scale factor, we apply a global scale factor that minimizes the difference between both. Furthermore, we also plot the recovered relative radiance values of each pixel in each patch (times the scale factor κ_{x_k}). Ideally, the recovered relative radiance should fall on the reference camera response function, but instead it forms a "cloud" around the reference camera response function due to the 1D subspace approximation, alignment errors, and camera noise. The (horizontal) extent of the "cloud" depends on the pixel values present in the patches, and thus this scale depends on the scene and camera settings (e.g., exposure).

4. Validation & Results

Results We demonstrate our method on a variety of scenes (Figure 6). For each example in Figure 6, we show (from left to right): the radiometrically calibrated source image, the resulting radiometric transfer result, a ground truth linearized image, a false color difference image of the former two, and a comparison of the recovered partial camera response function (red) to the ground truth camera response function (black). Pixel values that fall outside the range of the *partial* camera response function are set to white; we also highlight the discarded pixels in the inset. The reference linearized image is computed by applying the ground truth camera response function to the target uncalibrated image. We compute the ground truth camera response function from an exposure stack using the method of Kim and Pollefeys [KP04]. For all of our results, we assume all three color channels share the same camera response function, allowing us to triple the number of equations per patch. To provide a meaningful qualitative and quantitative comparison, we optimize for the optimal *global* scale factor that minimizes the differences between the reference and recovered camera response function; we also apply the same global scale factor to the radiometric transfer result.

Both the source and target images for the examples in Fig-

Table 1: *Summary of the absolute and relative RMSE of the results in Figure 6*

Name	Absolute	Relative
Brick Building	0.000961	1.6%
White House	0.000708	1.3%
Store	0.002698	6.2%
Church	0.001638	2.6%
Desk	0.001827	11.1%
Magazines	0.002920	6.5%
Magazines II	0.002214	7.0%
Flash	0.000585	6.0%

ure 6 are captured with a Nikon 700D camera. We briefly summarize the different scenes (from top to bottom; see also Table 1):

1. The *Brick Building* scene exhibits significantly different lighting and viewpoint between source (sunrise with strong shadows on the building) and the target image (noon on a clear day). The RMSE and relative error of the radiometric transfer are 0.000961 and 1.6% respectively.
2. The *White House* photographs are captured with a greater difference in viewpoint than the Brick Building, and under different lighting conditions: overcast sky (source) and clear sky (target). The RMSE and relative error are: 0.000708 and 1.3%.
3. The *Store* example is captured at approximately the same time of the day, but from a different viewpoint. The target image contains a significant portion not visible in the source image, which is correctly radiometrically linearized. The RMSE is 0.002698, and the relative error is 6.2%.
4. The *Church* is captured at different times of the year (Fall versus Spring). Note that the target image also contains a parked car not present in the source image. With the exception of the pixels outside the recovered range, this new object's radiance is recovered correctly. Additionally, this demonstrates that our method can handle significant changes between the input and target scene. The RMSE of the radiometric transfer is 0.001638, and the relative error is 2.6%.
5. The *Desk* source image is illuminated only by the LCD panel in the background, and the target is illuminated by office lighting. Note that besides differences in lighting and view, the directly visible pixels on the LCD panel are also different. This demonstrates that our method is robust to some degree in change to the scene. The RMSE and relative error of the radiometric transfer result are: 0.001827 and 11.1%. Note, the large relative error is mainly due to the many *dark* pixels.
6. The *Magazines* are captured under identical lighting conditions, but from different viewpoints. This result shows that the proposed method also works for non-diffuse materials (i.e., glossy magazine covers). The RMSE of the radiometric transfer is 0.002920, and the relative error is 6.5%.
7. The *Magazines II* example is captured under similar conditions as the *Magazines*, except that the order of the magazines is different between the source and target images. The RMSE of the result is 0.002214, and the relative error is 7.0%.
8. The *Flash* example demonstrates a radiometric transfer from a scene under ambient lighting to a scene illuminated by the camera flash only, illustrating the robustness of our method to drastic changes in lighting. The RMSE and relative error are: 0.000585 and 6.0%.

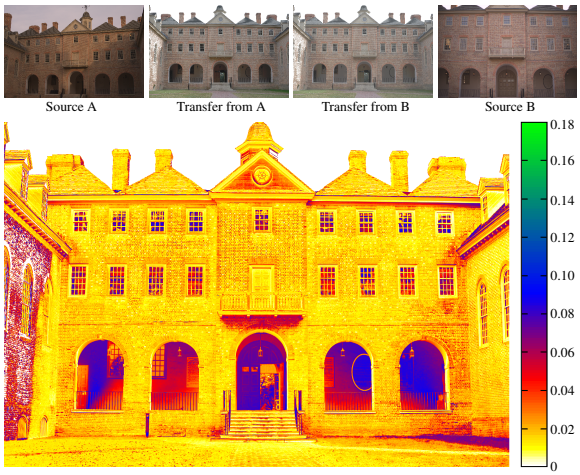

Figure 3: ***Stability Validation.*** *False color difference image between the radiometric transfer results from two different source exemplar images with significantly different lighting and view conditions.*

Stability To validate the stability of radiometric transfer, we compute the radiometric linear version of the target image of the *Brick Building* from two different source images with vastly different view, distance, and lighting conditions. The resulting radiometric linear images are compared in Figure 3. We rely on SIFT flow to compensate for the differences in distance and thus global scale of the image features. As can be seen, the recovered transfer from both source images are visually similar, indicating that our method is robust to different inputs.

Robustness to Camera Model We validate the robustness of the proposed method with respect to different camera models by simulating the acquisition of the target image for each of the camera response functions in the DoRF database [GN04] on the *Brick Building* and *Magazines* scenes from Figure 6. We then use the proposed radiometric transfer method to linearize the simulated target images, and compare the resulting images with the ground truth reference. The mean and variance of the RMS errors over the different camera response functions are 0.0022 and

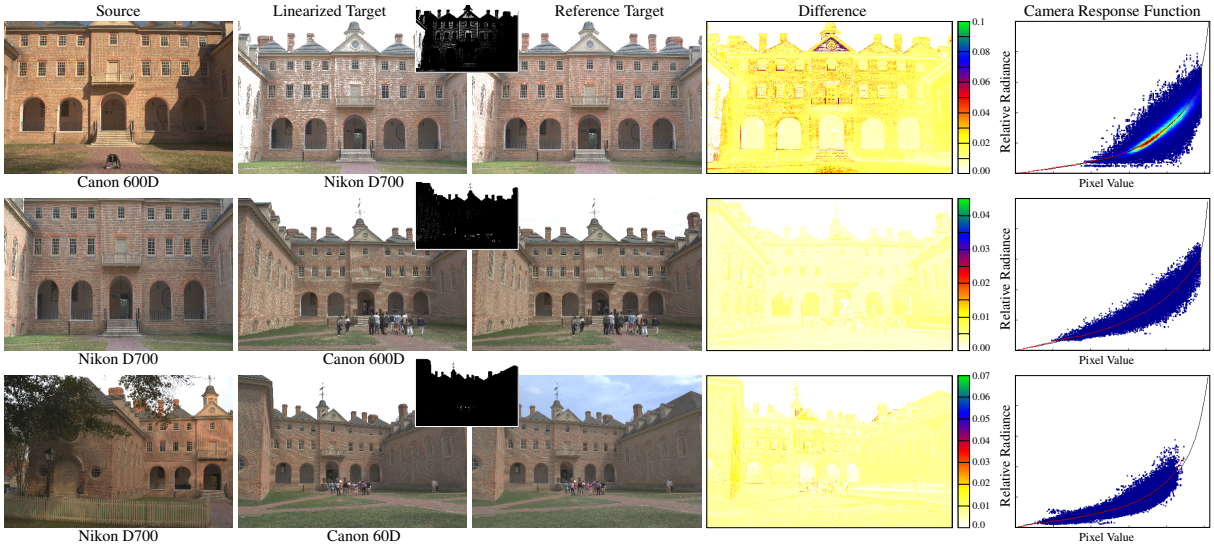

Figure 4: ***Transfer between Different Camera Models****. Robustness validation of radiometric transfer between three different camera models: Nikon D700, Canon 600D, and Canon 60D.*

4.19×10^{-6} for the *Brick Building* example, and 0.0036 and 8.66×10^{-6} for the *Magazines* example. This shows that the proposed method is robust to a wide variety of camera response functions.

Figure 4 shows a cross-validation of captured radiometrically linear/non-linear photographs obtained with three different camera models (Nikon D700, Canon 600D, and Canon 60D) of the *Brick Building* for a wide variety of viewpoints and lighting conditions. As can be seen, our method is able to accurately recover the camera response functions and linearize the target photographs for various combinations of camera-pairs; the relative error is below 3% on all examples.

Robustness to Lighting We validate the robustness of our method with respect to varying lighting conditions using the WILD database [NWN02] which contains a large selection of radiometrically linearized images of an urban scene under a wide variety of weather (and thus lighting) conditions. We select 70 random image pairs from the *clear weather* subset and simulate capture of the target image with a randomly selected camera response function from the DoRF database. Next, we recover the partial camera response function and radiometrically linearize the target image using the other image as the source. We show four selected pairs and the corresponding recovered camera response functions in Figure 5. Of the 70 image pairs, 50 yielded a successful transfer result (Figure 5, first three columns). 20 image pairs did not result in a successful transfer; the last column in Figure 5 shows such a case. The majority of the failure cases is due to an insufficient number of reliable SIFT matches, mainly caused by large dark regions (and thus little texture) due to shadows. The average and variance of the RMS errors are 0.0082 and 2.79×10^{-4} for the full 70 cases, and 0.0053 and 7.55×10^{-6} for the 50 successful transfers.

5. Discussion

1D Linear Subspace Model A key assumption in our method is that the change in appearance of a small image patch (viewed from different viewpoints and under different lighting conditions), spans a 1D linear subspace (after unwarping to correct for geometrical distortions). To better understand under which conditions this assumption holds, we express the appearance of a patch's pixel $E(x,\omega_o)$ in terms of the underlying surface normal $n(x)$, material properties $f_r(x,\omega_o,\omega_i)$, and incident lighting $L(x,\omega_i)$ at the surface location x viewed from a direction ω_o:

$$E(x,\omega_o) = \int_\Omega f_r(x,\omega_o,\omega_i) L(x,\omega_i) \max(n(x)\cdot\omega_i, 0) d\omega_i. \tag{5}$$

The 1D linear subspace assumption essentially factors the patch's appearance in a spatially dependent component $P(x)$ and a position-independent component $\kappa = \int K(\omega_i,\omega_o) d\omega_i$:

$$E(x,\omega_o) \approx P(x) \int_\omega K(\omega_i,\omega_o) d\omega_i. \tag{6}$$

To derive the exact form of the terms κ and $P(x)$, and to better understand the conditions under which this approximation is valid, we consider each of the three components in in Equation (5) separately:

- **Material properties** $f_r(x,\omega_o,\omega_i)$: Assuming that the outgoing direction is constant within a patch, we can factor the material properties in a position dependent albedo

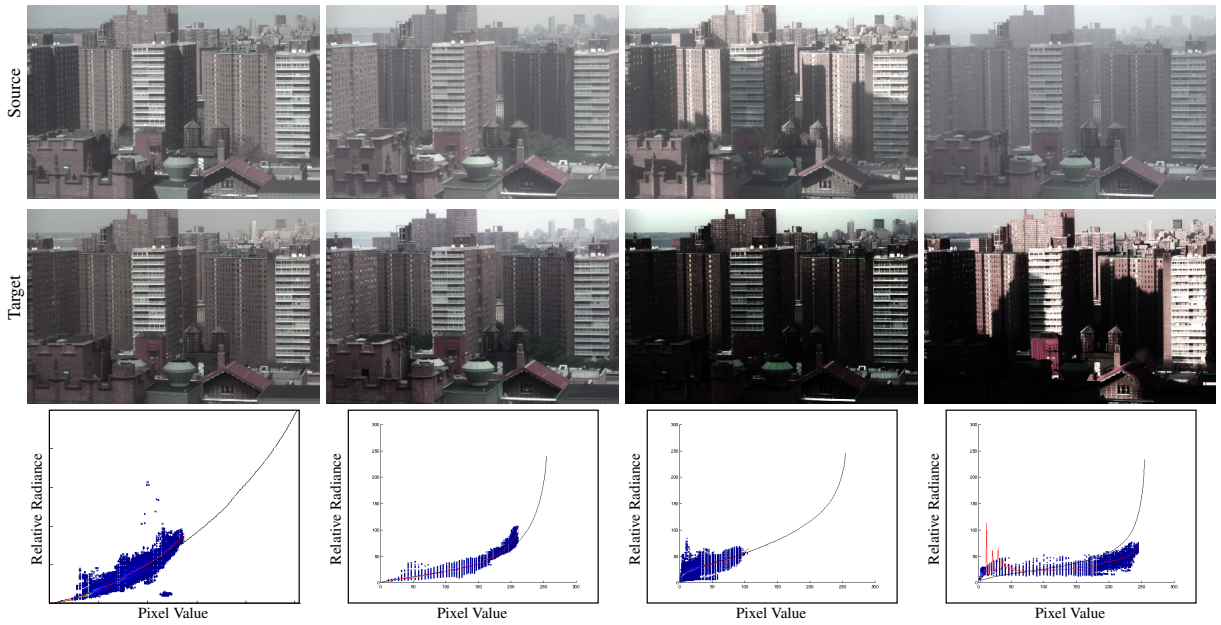

Figure 5: ***Robustness validation of radiometric transfer under different lighting conditions.*** *Four selected results from transfers between simulated captures from the WILD database [NWN02] and a random camera response from the DoRF database [GN04]. The first three columns show successful transfer results under vastly different lighting conditions between the source and target image. The last column shows a failure case where insufficient SIFT features were found.*

function $\alpha(x)$ and a normalized bidirectional reflectance distribution function $\rho(\omega_i, \omega_o)$. While such a factorization is exact for any monochromatic material model such as the common Lambertian surface reflectance model (e.g., with spatially varying albedo), it is only valid for a restricted form of the more general dichromatic surface reflectance model (i.e., a linear combination of diffuse and specular surface reflectance). In particular, this factorization is only valid if relative ratio of diffuse and specular albedo remains constant over the patch: $r \equiv \alpha_s/\alpha_d$, then $\alpha(x) = \alpha_d(x)$, and $\rho(\omega_i, \omega_o) = \rho_d(\omega_i, \omega_o) + r\rho_s(\omega_i, \omega_o)$.

- **Lighting** $L(x, \omega_i)$ can be made position independent by assuming distant lighting: $L(x, \omega_i) \approx L(\omega_i)$. Hence, the lighting is the same for all points in a patch. This excludes situations where a shadow edge crosses the patch or other strong position-dependent changes in the incident lighting. However, this does not imply identical incident lighting over *all* patches.
- **Geometric Term** $\max(n(x) \cdot \omega_i)$. Except for the case where the surface normal is constant over the patch (i.e., $n = n(x)$), the geometric term has both angular as well as positional (surface normal) dependencies. Consequently, the 1D subspace assumption only holds when the surface normals are constant within a patch.

Based on the above analysis, we can refine the the terms K and P in Equation (6) as:

$$P(x) = \alpha(x), \tag{7}$$

$$K(\omega_i, \omega_o) = \rho(\omega_i, \omega_o) L(\omega_i) \max(n \cdot \omega_i). \tag{8}$$

While theoretically only valid under the above assumptions, we found that in practice, small deviations from these assumptions can still be resolved in a least squares sense over many patches. Figure 1 (bottom-right) shows the approximation error for an office scene under different lighting conditions: for each pixel in the image, we compute the RMS error for the optimal scale factor κ in a 33×33 pixel-window. The majority of large errors occur in areas with strong deviations from the assumptions such as geometric discontinuities (e.g., edge of blanket) and across shadow boundaries (e.g., monitor stand).

Limitations The proposed method relies on a sufficient number of reliable SIFT correspondences that cover the majority of the pixel-intensity range in the target image, and SIFT-flow for fine-scale alignment. For scenes with little texture or for scenes recorded from radically different views or under extremely different lighting conditions, radiometric transfer can fail due to an insufficient number of reliable correspondences or inaccurate fine-scale alignment.

While our method is theoretically only valid when the surface normal does not vary within a patch, in practice it is robust to minor variations in surface normal (and thus depth)

due to the least squares fitting of the partial camera response function. However, our method can fail for scenes that consist mostly of fine details with significant depth discontinuities such as photographs of flower beds.

Additionally, the recovered camera response function is limited to the pixel-intensity range contained in the patches. Clearly, oversaturated and undersaturated pixels cannot be linearized; we exclude pixel-intensities outside the $[0.04, 0.96]$ range for both partial camera response curve recovery as well as linearization. Furthermore, the obtained camera response function might only cover a small portion of the intensity domain. However, we found that our algorithm typically finds a valid response function for most of the pixel range in the target image, and hence even with a partial camera response function, we can still obtain a good radiometric linearization.

Finally, we assume that the camera response function is invariant over the image, and that it is the only non-linear transformation applied to the target image. Other non-linear transformations introduced by chromatic aberrations, edge sharpening or adaptive demosaicing can bias the recovery of the camera response function. We currently, rely on statistical averaging over the various patches to mitigate the impact of such additional non-linear enhancements. However, this is not guaranteed and a skewed error distribution could aversely affect the accuracy of the recovered response function.

6. Future Work & Applications

The proposed method assumes that the appearance of an image patch can be well explained by a 1D linear subspace. Expanding this to 2D linear subspaces to better characterize the effects of general dichromatic reflectance functions would be an interesting avenue for future research. Furthermore, the proposed solution is akin to classical radiometric calibration from exposure stacks. This suggests that the proposed method could be adapted to fit in the robust low-rank framework of Lee et al. [LMS*13]. Finally, the pixel-accurate alignment requirement of small image patches is the main source of errors and failure cases. For future work, we would like to investigate methods, such as histogram-based strategies, to alleviate this requirement.

We believe our method paves the way for a number of novel radiometric calibration applications. Most existing work in radiometric calibration focuses on recovering the camera response function from multiple images of a scene taken from similar viewpoints and under similar lighting conditions. This implies direct access to the camera and full control over the target scene. However, this is not the case for large community photo-collections or historical photographs where the set of camera models is too large to fully sample or where the original camera model does not exists anymore, but where the depicted scene is still accessible. We briefly describe how radiometric transfer could be applicable in two such future applications:

- **Large Photo-collections** have recently been recognized as a unique *big data* source of visual information. Typically, the majority of the information in such photo-collections is not radiometrically linear, and prior work, in the absence of convenient and robust linearization methods suited for large photo-collections, has relied on ad-hoc methods to linearize the images (e.g., a gamma-correction [LBP*12]). However, we observe that for large image collections depicting a particular landmark it is likely that a small subset of the dataset is already radiometrically linear (e.g., single image radiometric calibration, known camera model/response function, etc.). Prior work on radiometric calibration of large image collections [DS11b, DS11a, KAGN08] has ignored this potentially rich source of information. One possible strategy for leveraging such information is to gradually grow the set of radiometrically calibrated images using the proposed radiometric transfer method which is robust to changes in view and lighting typically observed in such large photo-collections.
- **Historical Photographs** offer a window into the past. However, the cameras used to capture these images are not available anymore, or the exact process used to develop the film has been lost. Knowledge of the camera response function could help in identifying the camera model or process used to create these photographs. In the case when the scene still exists (e.g., a scenic subject or building), HDR rephotography combined with the proposed radiometric transfer method could recover the historical camera response function.

7. Conclusions

We presented a lightweight method for radiometrically linearizing an uncalibrated target image based on an exemplar calibrated photograph of the same scene recorded from a different viewpoint and under different lighting conditions. Key to our method is the observation that for many patches, their change in appearance (from different viewpoints and lighting) forms a 1D linear subspace. This allows us to reformulate the problem in a form similar to classic radiometric calibration from an exposure stack.

Acknowledgments We wish to thank the reviewers for their constructive feedback, and Kathleen Moore and Zhaoliang Duan for proofreading. This work was partially funded by NSF grants: IIS-1217765 and IIS-1350323, and a gift from Google.

References

[CR96] CHANG Y. C., REID J. F.: RGB calibration for color image-analysis in machine vision. *IEEE Trans. on IP 5*, 10 (Oct. 1996), 1414–1422. 2

[DM97] DEBEVEC P. E., MALIK J.: Recovering high dynamic range radiance maps from photographs. In *SIGGRAPH '97* (1997), pp. 369–378. 2, 3

[DS11a] DIAZ M., STURM P.: Exploiting image collections for recovering photometric properties. In *CAIP* (2011), pp. 253–260. 2, 8

[DS11b] DÍAZ M., STURM P.: Radiometric calibration using photo collections. In *ICCP* (Apr. 2011), pp. 1–8. 2, 8

[FP10] FURUKAWA Y., PONCE J.: Accurate, dense, and robust multiview stereopsis. *PAMI 32*, 8 (Aug. 2010), 1362–1376. 2

[GN03] GROSSBERG M., NAYAR S.: Determining the camera response from images: What is knowable? *PAMI 25*, 11 (Nov 2003), 1455–1467. 2

[GN04] GROSSBERG M., NAYAR S.: Modeling the space of camera response functions. *PAMI 26*, 10 (Oct 2004), 1272–1282. 2, 3, 5, 7

[HSGL13] HACOHEN Y., SHECHTMAN E., GOLDMAN D. B., LISCHINSKI D.: Optimizing color consistency in photo collections. *ACM Trans. Graph. 32*, 4 (July 2013). 2

[KAGN08] KUTHIRUMMAL S., AGARWALA A., GOLDMAN D. B., NAYAR S. K.: Priors for large photo collections and what they reveal about cameras. In *ECCV* (Oct 2008), pp. 74–87. 2, 8

[KFP08] KIM S. J., FRAHM J. M., POLLEFEYS M.: Radiometric calibration with illumination change for outdoor scene analysis. In *CVPR* (2008). 2

[KGFP10] KIM S. J., GALLUP D., FRAHM J.-M., POLLEFEYS M.: Joint radiometric calibration and feature tracking system with an application to stereo. *Computer Vision and Image Understanding 114*, 5 (May 2010), 574–582. 2

[KP04] KIM S. J., POLLEFEYS M.: Radiometric alignment of image sequences. In *CVPR* (2004), pp. 645–651. 3, 4

[KP08] KIM S. J., POLLEFEYS M.: Robust radiometric calibration and vignetting correction. *PAMI 30*, 4 (Apr. 2008), 562–576. 2

[LBP*12] LAFFONT P.-Y., BOUSSEAU A., PARIS S., DURAND F., DRETTAKIS G.: Coherent intrinsic images from photo collections. *ACM Trans. Graph. 31* (2012). 8

[LGYS04] LIN S., GU J., YAMAZAKI S., SHUM H.-Y.: Radiometric calibration from a single image. In *CVPR* (2004), pp. 938–945. 2

[LMS*13] LEE J.-Y., MATSUSHITA Y., SHI B., KWEON I. S., IKEUCHI K.: Radiometric calibration by rank minimization. *PAMI 35*, 1 (Jan. 2013), 144–156. 2, 8

[Low04] LOWE D. G.: Distinctive image features from scale-invariant keypoints. *IJCV 60*, 2 (Nov. 2004), 91–110. 4

[LYT*08] LIU C., YUEN J., TORRALBA A., SIVIC J., FREEMAN W. T.: Sift flow: Dense correspondence across different scenes. In *ECCV* (2008), pp. 28–42. 4

[LZ05] LIN S., ZHANG L.: Determining the radiometric response function from a single grayscale image. In *CVRP* (2005), pp. 66–73. 2

[ML07] MATSUSHITA Y., LIN S.: Radiometric calibration from noise distributions. In *CVPR* (2007). 2

[MN99] MITSUNAGA T., NAYAR S.: Radiometric self calibration. In *CVPR* (Jun 1999), vol. 1, pp. 374–380. 2

[NWN02] NARASIMHAN S. G., WANG C., NAYAR S. K.: All the images of an outdoor scene. In *ECCV* (2002), pp. 148–162. 6, 7

[OLTK15] OH T.-H., LEE J.-Y., TAI Y.-W., KWEON I. S.: Robust high dynamic range imaging by rank minimization. *PAMI* (2015). 2

[RAGS01] REINHARD E., ASHIKHMIN M., GOOCH B., SHIRLEY P.: Color transfer between images. *IEEE Comput. Graph. Appl. 21*, 5 (Sept. 2001), 34–41. 2

[SPB*14] SHIH Y., PARIS S., BARNES C., FREEMAN W. T., DURAND F.: Style transfer for headshot portraits. *ACM Trans. Graph. 33*, 4 (July 2014). 2

[SS04] SHAQUE K., SHAH M.: Estimation of the radiometric reponse functions of a color camera from differently illuminated images. In *ICIP* (2004). 2

[TMI08] TAKAMATSU J., MATSUSHITA Y., IKEUCHI K.: Estimating camera response functions using probabilistic intensity similarity. In *CVPR* (2008). 2

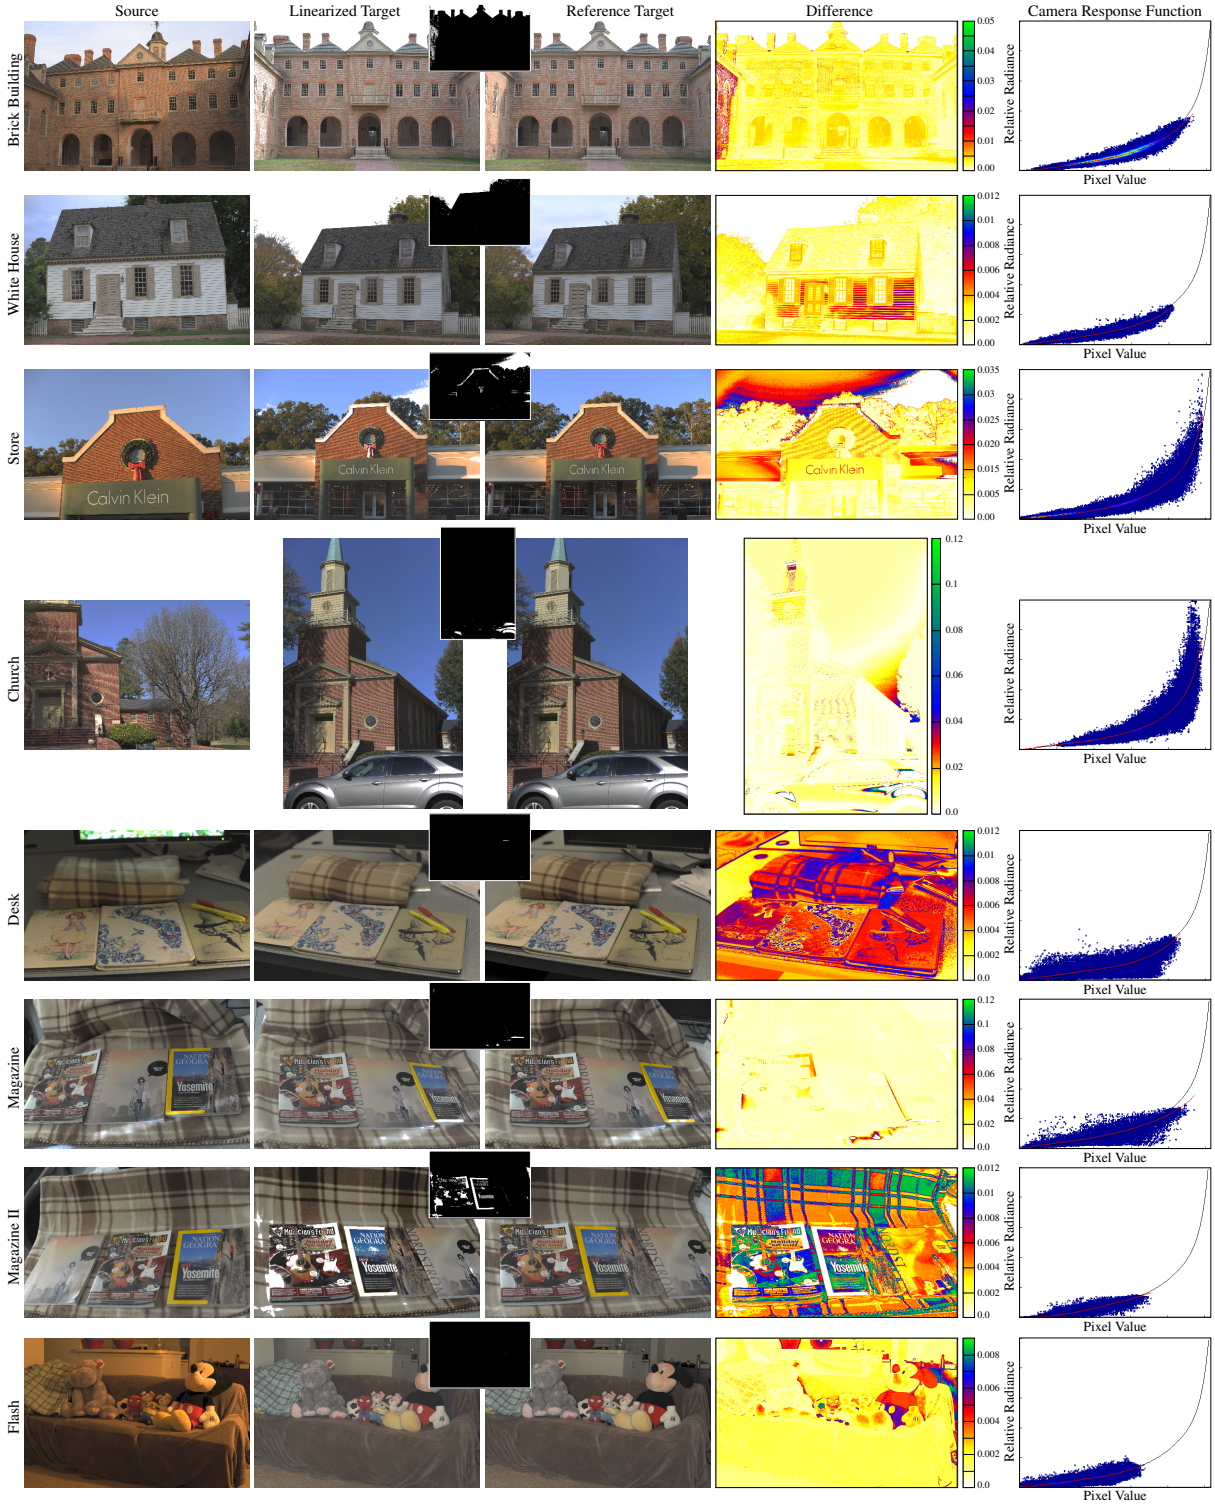

Figure 6: ***Radiometric Transfer Results.*** *Left to right: radiometrically calibrated source image, radiance transfer result (pixels outside the recoverable range are marked in white – also highlighted in the inset), reference linearized target image obtained by applying the (inverse) ground truth camera, false color difference, and camera response curve (recovered in red, ground truth in black). In addition, the recovered pixel radiance (times scale factor) for each pixel value is included in the right plot. Ideally this should form a tight 'cloud' around the ground truth camera response curve.*