Photometric Self-Calibration of a Projector-Camera System

Ray Juang
Department of Computer Science,
University of California, Irvine
rjuang@ics.uci.edu

Aditi Majumder
Department of Computer Science,
University of California, Irvine
majumder@ics.uci.edu

Abstract

In this paper, we present a method for photometric self-calibration of a projector-camera system. In addition to the input transfer functions (commonly called gamma functions), we also reconstruct the spatial intensity fall-off from the center to fringe (commonly called the vignetting effect) for both the projector and camera. Projector-camera systems are becoming more popular in a large number of applications like scene capture, 3D reconstruction, and calibrating multi-projector displays. Our method enables the use of photometrically uncalibrated projectors and cameras in all such applications.

1. Introduction

Projector-camera systems are commonly used in many applications like scene capture, 3D reconstruction, virtual reality, tiled displays and so on [17, 15, 19, 27]. The cameras and projectors used in these applications often require pre-calibration to assure accurate results. One particularly good example is that of multi-projector displays which allow users to move away from the rigidness of computer monitors and fixed displays [24, 19]. Cameras are now used regularly to calibrate such displays geometrically and photometrically [27, 2, 3, 6, 9, 11, 12, 13, 22, 21, 23, 26, 27, 10, 20, 18].

In this paper we present a self-calibration method that estimates the photometric parameters of an uncalibrated projector-camera system. The photometric calibration parameters of a projector/camera are its intensity transfer function and the spatial intensity variation function [8, 12]. The spatial variation, marked by a characteristic intensity fall-off from the center to fringe, is commonly called the vignetting effect. The vignetting effect need not be symmetric, especially for projectors, where it depends on the projector position, orientation, and the reflectance/tranmissive property of the screen. Our method estimates both the intensity transfer function and the spatial intensity variation for both camera and projector.

1.1. Related Work

Earlier work in photometric calibration of projection-based displays involved calibrating the projectors using either a precision optical instrument or a calibrated camera. [9, 26] find the projector intensity transfer function by using an expensive photometer or spectroradiometer. However, since photometers and radiometers can only measure one spatial location at a time, these methods cannot capture the spatial intensity variation of projectors. [20] uses a calibrated camera to estimate the projector intensity transfer function. First, the high dynamic range imaging technique described in [4, 16] is applied to estimate the camera's intensity transfer function. Once the camera is calibrated, the same high dynamic range technique is applied to recover the projector's intensity transfer function using this calibrated camera. The spatial intensity variation of the projector is then estimated by methods presented in [13, 11]. This entire method, however, assumes that the vignetting effect of the camera is negligible. This is only true for narrow apertures; hence, the camera is set to use a narrow aperture.

More recently, [25] presents the first method that estimates the intensity transfer functions of camera and projector by using isointensity curves in areas where a second projector overlaps the first. Since this method requires a second projector to compute the intensity transfer functions of a projector or camera, it cannot be applied to a single projector-camera system.

Achieving accurate photometric calibrations is an important issue even just for cameras. Several computer vision methods exist today to estimate the input transfer function of a camera [16, 5, 7, 14]. However, these use high dynamic range imaging in an outdoor setting where the user has relatively little control of the surrounding environment. Further, the problem of estimating the vignetting effect has been largely ignored. Not knowing the vignetting function forces applications to use their cameras at narrow apertures where the vignetting effect is negligible. Images taken in such settings have more noise than those taken at wider apertures. Thus, these applications have to address inaccuracies resulting from low signal to noise ratio. [1] presents

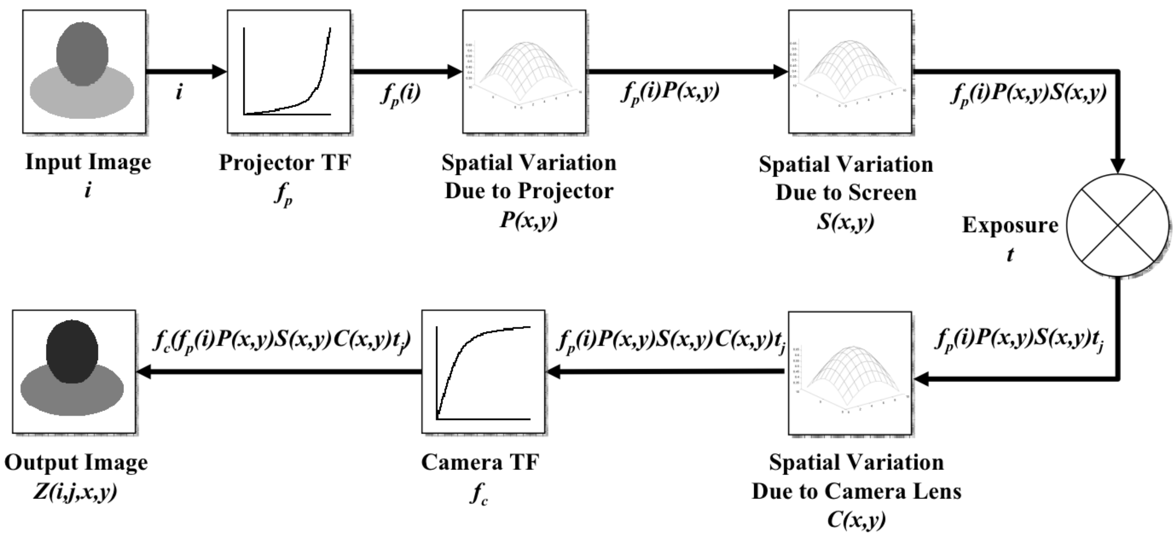

Figure 1. The transformation process of an image as it passes through a projector-camera system.

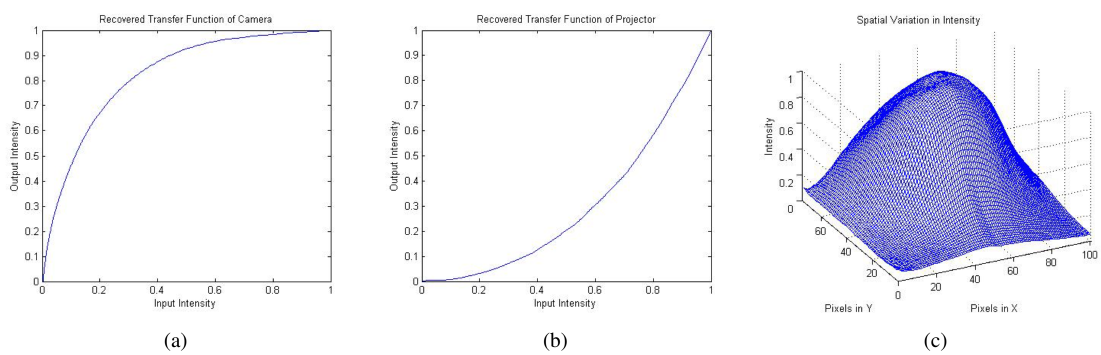

Figure 2. (a) The estimated camera input transfer function f_c. (b) The estimated projector input transfer function f_p. (c) The estimated spatial intensity variation due to projector, screen, and camera L.

an elaborate model that can estimate the vignetting effect for a camera whose input transfer function is already known or recovered. However, they require using a lens with high zoom capability.

1.2. Main Contributions

In this paper, we present a self-calibration technique for a projector-camera pair. To the best of our knowledge, this is the first work that estimates both the intensity transfer function and the vignetting effect of both projector and camera without using any other devices or physical props other than just a projector-camera pair. Any application that uses either a camera or a projector (or both) can thus benefit from this work. For example, one can now use photometrically uncalibrated cameras when using different photometric calibration techniques for single or multi-projector displays [13, 20]. Further, it can be also be used to photometrically calibrate cameras for any traditional computer vision applications like scene capture, 3D reconstruction, etc. Our method achieves a full photometric calibration of a camera by estimating both the input transfer function and the vignetting effect. This is achieved by using a projector, a device which is easily available anywhere today.

In the next section, we present the algorithm for estimating the photometric parameters of a projector-camera system. In Section 3, we present some example applications of how to use the calibrated devices. Finally, we conclude with future work in Section 4.

2. The Method

Our algorithm makes the following assumptions:

1. We assume a geometrically calibrated projector-camera system where a pixel (u, v) in the camera coordinate system is related to a pixel (x, y) in the projector coordinate system by a linear or non-linear warp $G(x, y) = (u, v)$. G can be determined by any standard geometric calibration method [27].

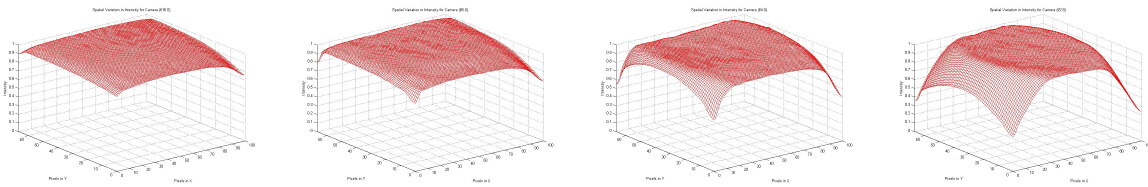

(a) (b) (c) (d)

Figure 3. The camera vignetting effect estimated after separation of parameters at (a) f/16, (b) f/8, (c) f/4 and (d) f/2.8. Note that the vignetting effect becomes more pronounced as the aperture size increases from (a) to (d).

2. Projectors and cameras are time-invariant devices whose photometric parameters do not change temporally.

3. The screen reflectance is time-invariant. It does not change when the power of light changes. Essentially, if the power of light increases, the radiance towards the camera increases proportionally.

Consider a spatially uniform grayscale input to the projector. Let the grayscale level be denoted by i. As per the model presented in [8], the uniform image is first transformed by a spatially invariant input transfer function of the projector, f_p, to create a spatially uniform output, $f_p(i)$. Next, the projector optics introduces a spatially dependent but input independent intensity variation $P(x, y)$. This results in a spatially varying image $f_p(i)P(x, y)$. This image is further modulated by the screen reflectance/transmissive function $S(x, y)$ to create another spatially varying image $f_p(i)P(x, y)S(x, y)$. The function is reflectance or transmissive depending on whether the system is a front or rear projection respectively. The light from the screen then reaches the camera. The amount of light accepted by the camera is scaled by its exposure time t_j, where j indexes different exposure times. The different exposures are instrumented by changing the shutter speed of the camera. This produces an image $f_p(i)P(x, y)S(x, y)t_j$ that passes through the camera optics which introduces another spatially dependent variation, $C'(u, v)$. The image thus generated is $f_p(i)P(x, y)S(x, y)t_jC'(u, v)$. To define the image in the projector coordinate space, we use $(u, v) = G(x, y)$ to define $C(x, y) = C'(G(x, y))$. The image in projector coordinate space is then given by $f_p(i)P(x, y)S(x, y)C(x, y)t_j$. Finally, this image is transformed by the spatially independent input transfer function of the camera, f_c, to generate the grayscale value recorded by the camera Z. Thus, Z is a function of the input i, the exposure time index j, and the spatial coordinates (x, y). This is illustrated in Figure 1. The final equation is

$$Z(i, j, x, y) = f_c(f_p(i)P(x, y)S(x, y)C(x, y)t_j). \quad (1)$$

In this equation, we first combine all of the spatially dependent terms into one term $L(x, y) = P(x, y)S(x, y)C(x, y)$. This represents the combined spatial variation introduced by the projector, screen, and camera optics in a closed form. Equation 1 thus becomes

$$Z(i, j, x, y) = f_c(f_p(i)L(x, y)t_j). \quad (2)$$

For cameras, the intensity transfer function is monotonic [4], and hence it is invertible. Note that the same is not true for projectors [12]. Assuming invertible f_c, the above equation becomes

$$f_c^{-1}(Z(i, j, x, y)) = f_p(i)L(x, y)t_j. \quad (3)$$

Taking the natural logarithm of both sides we get,

$$lnf_c^{-1}(Z(i, j, x, y)) = lnf_p(i)+ln(L(x, y))+ln(t_j). \quad (4)$$

To simplify the notation, we define $h_c = lnf_c^{-1}$ and $h_p = lnf_p$, The above equation then becomes

$$h_c(Z(i, j, x, y)) = h_p(i) + ln(L(x, y)) + ln(t_j) \quad (5)$$

where i ranges over the grayscale inputs, j ranges over the exposure times, and (x, y) ranges over the spatial coordinates of the projector. In this equation, Z and t_j are known while h_p, h_c and L are unknown. We want to recover h_p, h_c and L that best satisfy Equation 5 in a least-squares sense. Note that recovering h_p and h_c involves solving the functions for a finite number of samples in the complete range of input values. Varying i and t_j results in different values of Z for each pixel (x, y). We can use this to setup a system of linear equations. By solving this system we can recover h_p, h_c, and L as illustrated in Figure 2.

2.1. Separation of Spatial Parameters

The recovered $L(x, y)$, as shown in Figure 2, is the combined spatial intensity variation introduced by $P(x, y)$, $S(x, y)$, and $C(x, y)$. In this scenario, the spatial variation $C(x, y)$ is a function of the camera aperture setting. The vignetting effect becomes increasingly pronounced for larger aperture sizes. This happens since at wider apertures the camera deviates considerably from the pinhole model [1].

To assure a near uniform $C(x, y)$, scene capture methods operate with the camera set to a narrow aperture setting

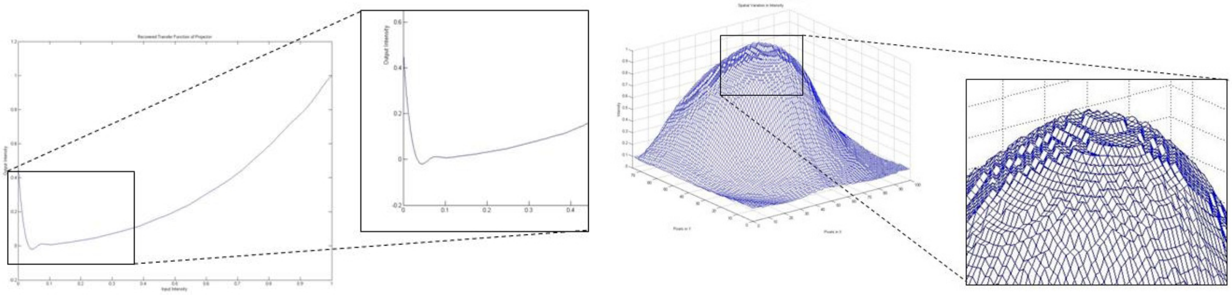

(a) (b)

Figure 4. (a) The estimated projector input transfer function f_p with a zoomed in portion to show the noise. (b) The estimated spatial intensity variation due to projector, screen, and camera L with a zoomed in portion to show noise.

[12, 4, 18]. As a result, the camera approaches the ideal pinhole model resulting in almost negligible spatial variation in $C(x, y)$. Though this leads to more noise in the acquired data, it is preferred over inaccuracies introduced by the presence of the vignetting effect at wider aperture settings [15, 18, 17].

We use this fact to separate $C(x, y)$ from $L(x, y)$. Let us assume that the camera offers different aperture settings, $a_1, a_2, \dots a_n$, where a_1 is the most narrow aperture setting. We first reconstruct $L(x, y)$ at different aperture settings a_k, $1 \leq k \leq n$ using the linear system of equations generated by Equation 5, denoted by $L_k(x, y)$. We assume that at a_1 aperture setting the camera vignetting is negligible. Hence, $C(x, y)$ is close to 1 and $L_1(x, y) = P(x, y)S(x, y)$. At wider aperture settings of a_k, $2 \leq k \leq n$, the vignetting effect of the camera $C_k(x, y)$ is then given by

$$C_k(x, y) = \frac{C_k(x, y)P(x, y)S(x, y)}{P(x, y)S(x, y)} = \frac{L_k(x, y)}{L_1(x, y)}. \quad (6)$$

Note that $C_k(x, y)$ is the camera vignetting effect represented in the projector coordinate space. $C'_k(u, v)$ is the same function in the camera's coordinate space and can be easily found using the geometric warp $G(x, y) = (u, v)$. Figure 3 shows the estimated spatial variation or vignetting effect at different apertures. As expected, an increase in aperture size (i.e decrease in f-stop) leads to more pronounced vignetting.

2.2. Performance

The system of linear equations achieved by Equation 5 is very large. Let D_i and D_Z each be the domain of h_p and h_c respectively, P be the set of pixels in the projector's coordinate space, and T be the set of camera exposures. This results in $|P||T||D_i|$ equations and $|P| + |D_i| + |D_Z|$ unknown variables for the system of equations defined by Equation 5. Typically $|D_i| = |D_Z| = 256$ and $|P| = 1024 \times 768 = 786432$ (assuming common XGA projector resolution). Since there are multiple exposures for each

(a) (b)

Figure 5. (a) The predicted image using the estimated parameters in Equation 5. (b) The image captured by the camera.

input in D_i, the size of the linear system is on the order of at least a few million equations. Solving such a huge system would make the method inefficient.

We address this inefficiency by first using a limited number of pixels in the projector space to solve for h_p and h_c. To ensure a sufficiently over-determined system, the criteria $|P||T||D_i| > |P| + |D_i| + |D_Z|$ should be satisfied. We use a subset of the projector's pixels and solve a linear system of equations of much smaller size. For $|D_i| = |D_Z| = 256$ and $|T| = 6$, a choice of 100 for $|P|$ is more than adequate. We first subsample L to a resolution of 10×10 pixels. We use this to setup a smaller linear system of equations and solve for h_p and h_c.

With the estimated h_p and h_c, we can substitute these into Equation 5 and quickly back-solve for $L(x, y)$ at the various projector coordinates, by rewriting Equation 5 as

$$ln(L(x, y)) = h_c(Z(i, j, x, y)) - h_p(i) - ln(t_j) \quad (7)$$

Ideally, any image that has an unsaturated Z at the spatial location (x, y) can be used to find $L(x, y)$. However, this will yield a noisy $L(x, y)$. To reduce this noise in $L(x, y)$, we can weigh values from multiple images, as detailed in the next section.

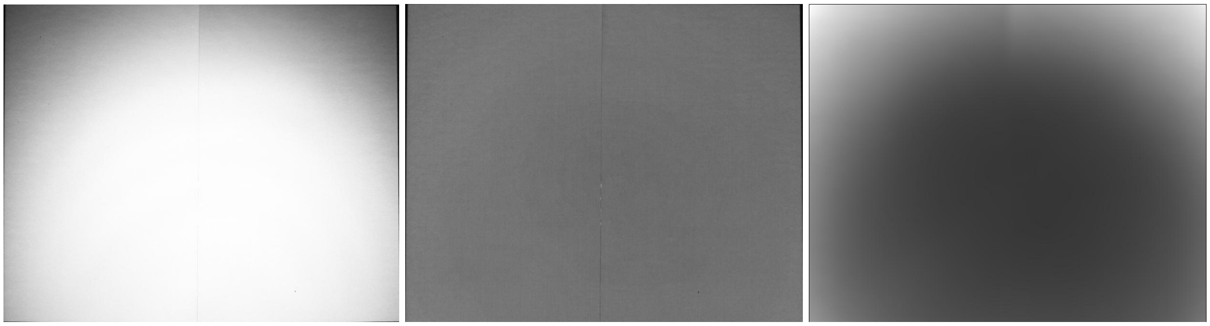

(a) (b) (c)

Figure 6. (a) The image captured by a camera when the projector displays a flat field. (b) The image captured by the camera when the projector displays the image in (c). (c) The corrected input image sent to the projector to create a visually flat field.

2.3. Accuracy

Noise is an important issue when solving for any large system of linear equations. The noise arises not only from the devices (camera and projector) but also from the screen. In particular, we use a rear projection screen with relatively high gain which has been shown to generate considerable noise[12]. If we do not take measures to address this, the recovered parameters can be very noisy as shown in Figure 4. To achieve a cleaner result, as in Figure 2, we can constrain the solution of our linear system to reduce the noise introduced in the estimated parameters.

To assure smooth h_p and h_c functions while solving the system of equations, we want to minimize the error function

$$E = \sum_{j \in T} \sum_{(x,y) \in P} \sum_{i \in D_i} [h_c(Z) - h_p(i) - ln(L(x,y)) - ln(t_j)]^2 \quad (8)$$

$$+\lambda \left(\sum_{Z \in D_Z} h_c''(Z)^2 + \sum_{i \in D_i} h_p''(i)^2 \right). \quad (9)$$

The first term assures that the solution arises from the set of equations given by Equation 5 in a least-squares sense. The second term is a smoothing constraint on the curvature of h_p and h_c, given by their second derivative. In the discrete domain, we use the Laplacian operator to find the curvature of h_p and h_c. For example, $h_p''(i) = h_p(i-1) - 2h_p(i) + h_p(i+1)$. The scale factor λ weighs the smoothness term relative to the data fitting term and should be chosen based on the amount of noise in Z.

Notice that the first term in Equation 9 gives equal weights to all recorded camera values Z, and the second term gives equal weights to all projector inputs i. However, the images with lower energy are much more likely to be affected by noise than signals with higher energy. In this scenario, this means that noise is high for lower values of i and Z. We want to weigh higher energy signals with greater confidence. To achieve this, we modify the first and second term of the error function in Equation 9 as

$$\lambda \left(\sum_{Z \in D_Z} w_c(Z) h_c''(Z)^2 + \sum_{i \in D_i} w_p(i) h_p''(i)^2 \right). \quad (10)$$

where w_p and w_c are the weighting functions corresponding to the projector input and the recorded camera values respectively. Since higher intensities have higher energy, we give the higher intensities greater confidence by using linear weighting functions, $w_c(Z) = Z$ and $w_p(i) = i$.

Another source of noise is when $L(x,y)$ is estimated using Equation 7. Once h_c and h_p are recovered from the sub-sampled projector space, they are then used to solve for $L(x,y)$ in the original projector space. Note that the value Z recorded from a spatial location (x,y) is different in images captured at different exposures for different projector inputs and can span the entire range of values in D_Z. Usually the image captured for input i at a particular exposure does not yield unsaturated outputs at all spatial pixel location (x,y). So, we need to use different images for reconstruction of L at different spatial locations. This yields a noisy $L(x,y)$ due to the presence of noise in both the projected and the captured images. An obvious way to reduce this noise is averaging. For each spatial location (x,y), we can average the multiple $L(x,y)$ values that we find by back-solving all the different images with an unsaturated value at (x,y) and thus reduce the effect of noise. Instead of averaging, we modify the averaging process by weighing the estimated $L(x,y)$ values according to a confidence measure that will suppress the impact of noise by assigning higher intensities greater confidence (since higher energy signals have greater signal to noise ratio). We weight $L(x,y)$ by a function $w_L(Z,i) = w_p(Z)w_p(i)$. The noise is most likely to affect images with low i and Z where $w_L(Z,i)$ is very low. Thus, the brighter and less noisy images are emphasized more than the noisy ones by using w_L. The back-solving is now done by modifying Equation 7

$$ln(L(x,y)) = \frac{\sum_{j \in T} \sum_{i \in D_i} w_L(Z,i)[h_c(Z) - h_p(i) - ln(t_j)]}{\sum_{j \in T} \sum_{i \in D_i} w_L(Z,i)}, \quad (11)$$

(a) (b) (c) (d)

Figure 7. (a,c) An image captured at aperture f/2.8. Note the darker corners where the vignetting effect is most apparent. (b,d) The estimated camera vignetting effect is used to correct images taken at aperture f/2.8. Note that the darkened corners are removed completely.

resulting in considerable noise reduction as illustrated in Figure 2.

2.4. Implementation

To test our methodology, we used a Kodak DCS ProSLR/n camera and a standard presentation projector, Epson 74c. We projected 32 flat grayscale fields with intensity levels uniformly sampled from 0 to 255. For each intensity level, 15 exposures were taken. This process was repeated for 5 different aperture settings: f/32, f/16, f/8, f/4 and f/2.8. The data collection took 45 minutes per aperture. To reduce the collection time, we tried reducing the number of exposures. 8 seems to be the minimum number of exposures needed before the effects of noise become visible. The exposures used, however, need to be well distributed amongst the range of available camera exposures. The program for recovering the projector-camera parameters was written in C++/Matlab and utilizes the OpenCV library. On a Pentium 4, 2.8 GHz PC, it takes 11-15 minutes to process and recover the projector-camera parameters, i.e. the camera and projector transfer functions, the camera vignetting effect at different apertures, and the spatial intensity variation of the projector and screen combined.

2.5. Verification

To verify the accuracy of the estimated parameters, we performed two experiments. In the first experiment, we took an arbitrary input image and applied the estimated parameters as per Equation 5 to generate a *predicted image*. We compared this with an image captured by the camera. Figure 5 shows that the predicted image is indistinguishable from the actual image captured by the camera thus verifying the accuracy of the parameters recovered by our method.

In the second experiment, we used our results to display a visually uniform gray field. Providing a flat gray input to the projector results in an image captured by the camera that appears to be non-uniform. This is due to the spatial intensity variation of the camera, projector, and screen $L(x, y)$. Note that to achieve a uniform gray field, we have to compensate for the variation in $P(x, y)S(x, y)$ achieved after separation of parameters (Section 2.1). If this separation is not performed, presence of $C(x, y)$ in $L(x, y)$ would lead to over-compression of the brightness near the center of the projector which in turn will cause significant loss in the dynamic range. To achieve a uniform image on the camera we divide the flat gray input image by $\min(P(x, y)S(x, y))$. We then project this result and capture it on camera. Since $P(x, y)S(x, y)$ is an accurate estimate, the captured result appears flat as illustrated Figure 6.

3. Using Calibrated Projectors and Cameras

We have only addressed photometric calibration for grayscale images. The analysis, however, extends directly for photometrically calibrating each color channel on a per-channel basis using already existing related work to describe how different color channels interact [18, 13]. We omit the details here because of the lack of space.

Projector-camera systems are used in a plethora of applications today, including scene capture, 3D reconstruction, and automated calibration of multi-projector displays [15, 17, 27]. In most of these applications, the camera is almost always set to a narrow aperture to avoid vignetting artifacts. This results in images having low signal to noise ratio which adversely affects important algorithmic components like feature matching or image blending. Increasing the aperture size decreases noise but also affects the accuracy of results by adding vignetting artifacts. Our projector-camera photometric calibration can be used to correct vignetting and thus, allow cameras to be used at wider apertures in such applications.

Figure 7 shows the result of using the estimated spatial intensity variation of the camera to correct for vignetting effects in images captured at wide apertures. Since the vignetting effect assumes a linear relationship between the input light and the captured camera values, the captured image is first linearized using the inverse of the estimated camera input transfer function. The reciprocal of the estimated vignetting effect is then multiplied with the linearized image to generate the corrected image in linear space. Finally, the input transfer function of the camera is applied to the corrected image to bring it back to the non-linear space.

As an example, we demonstrate using our method

(a) (b) (c) (d)

Figure 8. (a,c) Panorama generated from images taken at aperture f/2.8 without vignetting correction. Notice that there are dark vertical bands for the overcompensated blending regions. (b,d) The estimated camera vignetting effect is used to correct the images before they are stitched into a panorama. These contain no perceivable seams.

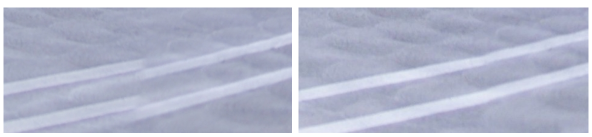

(a) (b)

Figure 9. (a,b) Zoomed in portion of the panorama in Figure 8(c) and (d) respectively. Due to poor feature detection, there is a mismatch in geometric matching in Figure 8(c). This is resolved when our method is used to generate Figure 8(d).

to improve feature matching and image blending in 2D panoramic image generation applications. We compare the panoramas generated from the same set of images in two ways. In the first approach, which is commonly adopted, the photometric parameters of the camera are unknown. The set of images taken by the camera are stitched together and then blended in the regions where adjacent images overlap. In the second approach, we first calibrate the camera using our self-calibration method. Thus, the photometric parameters of the camera are known. Then, the same set of images captured for the panorama is first linearized using the camera's inverse transfer function. They are then stitched together with overlap blending and finally brought back to the non-linear space by applying the inverse transfer function.

We find that the latter method provides better feature matching, enabling a better geometric match across images. Vignetting affects the fringes of each image which forms a big part of the overlap region with the adjacent image. Lower intensity in this region affects the quality of feature matching adversely. Further, the former panorama shows dark bands in the blending region which are eliminated completely in the latter one. This is because blending assumes linear input transfer function, which is not true in the former method, leading to over-compensation in the already darkened overlap regions. The latter method, on the other hand, performs the blending in a linear space eliminating the dark bands. Figure 8 and 9 illustrate the results.

Our projector-camera self-calibration method enables photometric calibration of multi-projector displays using an uncalibrated camera. Existing photometric calibration methods [13, 20, 11, 12] require a photometrically calibrated camera. Using our method, each projector can be independently calibrated using the same uncalibrated camera. The recovered projector parameters can then be used to photometrically calibrate the display using existing techniques that modify the parameters appropriately to achieve a seamless display [13].

4. Conclusion

In this paper, we have presented a method for complete photometric calibration of a projector-camera system. In addition to contributing to the state of the art in device calibration, our method enables the use of photometrically uncalibrated projectors/cameras in various applications ranging from multi-projector displays to 2D/3D scene capture.

This work has several future extensions. The described method in this paper estimates the photometric parameters at a fixed zoom setting. Vignetting effects change significantly with zoom settings. We would like to extend our work to capture these changes and find a model that will allow efficient storage and access of these parameters across such changes. A full color self-calibration of the projector-camera pair is another extension to consider. The commonly available white balance control on the projectors and cameras can have a significant role to play in this. One can imagine a closed-loop calibration technique that not only changes the exposure setting of the camera but also its white balance to instrument observations with changes in color parameters. This information can then be analyzed to achieve a full color calibration of the camera.

References

[1] N. Asada, A. Amano, and M. Baba. Photometric calibration of zoom lens systems. *Proceedings of the 13th International Conference on Pattern Recognition*, 1, 1996.

[2] H. Chen, R. Sukthankar, G. Wallace, and K. Li. Scalable alignment of large-format multi-projector displays using camera homography trees. *Proceedings of IEEE Visualization*, 2002.

[3] Y. Chen, D. W. Clark, A. Finkelstein, T. Housel, and K. Li. Automatic alignment of high-resolution multi-projector displays using an un-calibrated camera. *Proceedings of IEEE Visualization*, 2000.

[4] P. E. Debevec and J. Malik. Recovering high dynamic range radiance maps from photographs. *Proceedings of ACM Siggraph*, pages 369–378, 1997.

[5] M. Grossberg and S. Nayar. What is the space of camera response functions? *Proceedings of IEEE Conference on Computer Vision and Pattern Recognition*, 2003.

[6] M. Hereld, I. R. Judson, and R. Stevens. Dottytoto: A measurement engine for aligning multi-projector display systems. *Argonne National Laboratory preprint ANL/MCS-P958-0502*, 2002.

[7] S. Kim and M. Pollefeys. Radiometric self-alignment of image sequences. *Proceedings of IEEE Conference on Computer Vision and Pattern Recognition*, 2004.

[8] A. Majumder and M. Gopi. Modeling color properties of tiled displays. *Computer Graphics Forum*, June 2005.

[9] A. Majumder, Z. He, H. Towles, and G. Welch. Achieving color uniformity across multi-projector displays. *Proceedings of IEEE Visualization*, 2000.

[10] A. Majumder, D. Jones, M. McCrory, M. E. Papka, and R. Stevens. Using a camera to capture and correct spatial photometric variation in multi-projector displays. *IEEE International Workshop on Projector-Camera Systems*, 2003.

[11] A. Majumder and R. Stevens. LAM: Luminance attenuation map for photometric uniformity in projection based displays. *Proceedings of ACM Virtual Reality and Software Technology*, 2002.

[12] A. Majumder and R. Stevens. Color nonuniformity in projection-based displays: Analysis and solutions. *IEEE Transactions on Visualization and Computer Graphics*, 10(2), March–April 2003.

[13] A. Majumder and R. Stevens. Perceptual photometric seamlessness in tiled projection-based displays. *ACM Transactions on Graphics*, 24(1), January 2005.

[14] S. Mann. Comparametric equations with practical applications in quantigraphic processing. *IEEE Transactions on Image Processing*, 9(8):1389 – 1406, 2000.

[15] V. Masselus, P. Peers, P. Dutre, and Y. D. Willems. Relighting with 4d incident light fields. *Proceedings of ACM SIGGRAPH)*, 2003.

[16] T. Mitsunaga and S. Nayar. Radiometric self calibration. *Proceedings of IEEE Conference on Computer Vision and Pattern Recognition*, 1999.

[17] S. K. Nayar, G. Krishnan, M. D. Grossberg, and R. Raskar. Fast separation of direct and global components of a scene using high frequency illumination. *ACM Transactions on Graphics (SIGGRAPH)*, 25(3), 2003.

[18] S. K. Nayar, H. Peri, M. D. Grossberg, and P. N. Belhumeur. A projection system with radiometric compensation for screen imperfections. *Proceedings of IEEE International Workshop on Projector-Camera Systems*, 2003.

[19] C. Pinhanez, M. Podlaseck, R. Kjeldsen, A. Levas, G. Pingali, and N. Sukaviriya. Ubiquitous interactive displays in a retail environment. *Proceedings of SIGGRAPH Sketches*, 2003.

[20] A. Raij, G. Gill, A. Majumder, H. Towles, and H. Fuchs. PixelFlex2: A Comprehensive, Automatic, Casually-Aligned Multi-Projector Display. *IEEE International Workshop on Projector-Camera Systems*, 2003.

[21] R. Raskar. Immersive planar displays using roughly aligned projectors. In *Proceedings of IEEE Virtual Reality 2000*, 1999.

[22] R. Raskar, M. Brown, R. Yang, W. Chen, H. Towles, B. Seales, and H. Fuchs. Multi projector displays using camera based registration. *Proceedings of IEEE Visualization*, 1999.

[23] R. Raskar, J. van Baar, P. Beardsley, T. Willwacher, S. Rao, and C. Forlines. ilamps: Geometrically aware and self-configuring projectors. *ACM Transactions on Graphics*, 22(3), 2003.

[24] R. Raskar, G. Welch, M. Cutts, A. Lake, L. Stesin, and H. Fuchs. The office of the future: A unified approach to image based modeling and spatially immersive display. In *Proceedings of ACM Siggraph*, pages 168–176, 1998.

[25] P. Song and T. J. Cham. A theory for photometric self calibration of multiple overlapping projectors and cameras. *IEEE CVPR Workshop on Projector Camera Systems*, 2005.

[26] R. Yang, D. Gotz, J. Hensley, H. Towles, and M. S. Brown. Pixelflex: A reconfigurable multi-projector display system. *Proceedings of IEEE Visualization*, 2001.

[27] R. Yang, A. Majumder, and M. Brown. Camera based calibration techniques for seamless multi-projector displays. *IEEE Transactions on Visualization and Computer Graphics*, 11(2), March-April 2005.