Addressing Radiometric Nonidealities: A Unified Framework

Anatoly Litvinov and Yoav Y. Schechner

Dept. of Electrical Engineering
Technion - Israel Inst. Technology
Haifa 32000, ISRAEL
lanatol@tx.technion.ac.il , yoav@ee.technion.ac.il

Abstract

Cameras may have non-ideal radiometric aspects, including spatial non-uniformity, e.g., due to vignetting; a nonlinear radiometric response of the sensor; and temporal variations due to automatic gain control (AGC). Often, these characteristics exist simultaneously, and are typically unknown. They thus hinder consistent photometric measurements. In particular, they create annoying seams in image mosaics. Prior studies approached part of these problems while excluding others. We handle all these problems in a unified framework. We suggest an approach for simultaneously estimating the radiometric response, the spatial non-uniformity and the temporally varying gain. The approach does not rely on dedicated processes that intentionally vary exposure settings. Rather, it is based on an ordinary frame sequence acquired during camera motion. The estimated non-ideal characteristics are then compensated for. We state fundamental ambiguities associated with this recovery problem, while exposing a novel image invariance. The method is demonstrated in several experiments, where different frames are brought into mutual radiometric consistency. The accuracy achieved is sufficient for seamless mosaicing, with no need to resort to dedicated seam-feathering methods.

1. Introduction

Imaging systems are prone to radiometric problems [11, 27], such as nonuniform irradiance of the detector plane [1, 30, 31] and a nonlinear response of the sensor [4, 17, 22]. The spatial non-uniformity is caused by vignetting and foreshortening [1, 14, 30, 31], while the nonlinear response is usually due to the electronic characteristics of the sensor amplifier. Moreover, many cameras have an automatic gain control (AGC) mechanism. It temporally varies camera settings [3, 7, 17], such as exposure time or electronic amplification. As a result, the same scene point can correspond to different gray levels in different frames. Such inconsistencies are usually not negligible: even an inconsistency of 1% is noticeable in 8-bit sensors (and certainly when using higher sensitivity sensors). Moreover, humans can easily detect such minute radiometric errors, as demonstrated in Fig. 1.

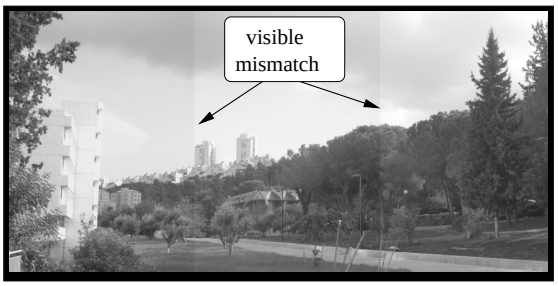

Figure 1. Illustrating human sensitivity to radiometric mismatch. Several consecutive image parts were biased by 3% with respect to each other. Even such a small mismatch creates clear visual artifacts.

These inconsistencies limit computer vision algorithms that are based on quantitative photometric measurements. In addition, image mosaicing [9, 10, 12, 25, 26, 27, 28] may be disrupted by them. Specifically, the inconsistencies may complicate image registration [3, 21], and create seams in the resulting mosaic [2, 5, 13]. It is important to alleviate such problems. This is especially true in novel mosaicing applications that make demanding use of all the acquired data , e.g., for high dynamic range imaging [27], multispectral sensing [26] and representations of dynamics [24]. Some methods were developed to eliminate the visual effect of seams [2, 5, 18, 28]. However, those pure image processing methods address perceptual effects, rather than addressing the root cause of the mismatch.

If the mentioned radiometric characteristics of the camera were perfectly quantified, they could be compensated for, thereby bringing all frames of a scene to mutual consistency. Apparently, these problems can be solved by careful camera pre-calibration. A standard calibration process, however, may not always be practical. For instance, zoom

lenses would require nonuniformity calibration at all possible zoom, focus and aperture settings. Moreover, some cameras (mainly consumer grade) do not allow for manual control of camera settings. Apparently, it is possible to solve the problem of spatial non-uniformity by precise optical design. However, this utopian line of thought may pose impractical tolerances on the optical engineering process. Hence, this problem is a fact of life [1, 30, 31]. In addition, the performance of cameras (particularly their electronics) may somewhat degrade over long periods of time, especially in harsh environments.

To counter these problems, we thoroughly examine the radiometric phenomena that cause inconsistencies. The question we ask is whether frames can be brought to mutual consistency, even when neither the spatial non-uniformity, the radiometric response nor the temporal gain variations are known? As we detail in Sec. 3, partial approaches to this problem had been suggested [6, 4, 16, 17, 20, 21, 22, 27]: some excluding the effect of unknown gain, others excluding the unknown nonlinearity of the radiometric response, while the rest ignore the unknown spatial non-uniformity.

In this work we propose a *unified* solution to all of these problems. The unknown spatial non-uniformity, radiometric nonlinearity and temporal gain variations are *blindly self-calibrated* based on an image sequence of an ordinary scene, as part of the mission the imaging system is designed to do. All of these aspects are estimated simultaneously. Moreover, our approach, does not need parametric models to estimate unknown functions.

We discuss fundamental ambiguities associated with our estimation problem. We show that in spite of the ambiguities, the recovered image is invariant, up to a simple transformation. Our method is able to compensate the raw frames, such that all of the resulting images are consistent. Hence, the output value of a scene point is identical for all frames (excluding random noise), whether it is imaged at the center of a frame, or at the periphery, irrespective of the temporal gain. Using the compensated images, we construct wide field of view (FOV) image mosaics. The compensation quality is sufficiently high to practically eliminate the visibility of seams in mosaics, *without resorting to dedicated seam-removal techniques*. These merits are demonstrated in experiments using real data.

2 Statement of the Problem

Consider the imaging system depicted in (Fig. 2). The detector plane is often non-uniformly irradiated even if the scene radiance is uniform. Typically this problem is optical: due to vignetting and foreshortening effects [1, 14, 30, 31], the sensor periphery is not illuminated as intensely as its center. This spatially varying optical response is denoted as $M(\vec{x})$, where $\vec{x} = (x, y) \in [1, N]^2$ is a location vector in the local coordinate system of the frame. Accounting for

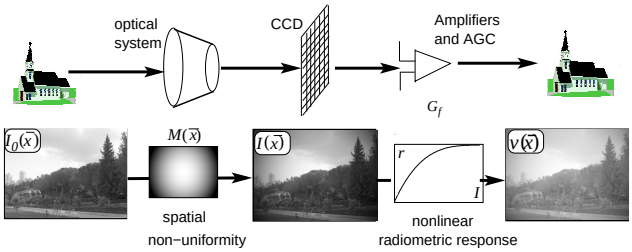

Figure 2. The imaging system model. The optical system induces spatial inhomogeneities on the image, which are characterized by $M(\vec{x})$. The camera electronics has an unknown radiometric response $r(I)$ and unknown frame dependent gain G_f.

this spatial variation, the irradiance falling on the camera detector is

$$\tilde{I}(\vec{x}) = M(\vec{x}) I(\vec{x}) \ . \tag{1}$$

Here I is the irradiance falling on the camera detector when $M = 1$.

The irradiance $\tilde{I}$ is converted into an electric signal via the camera amplifier. In some cameras the amplification is purely electronic. In others, the exposure time is extended to compensate for low light flux. Another mechanism involves electron multiplication by image intensifiers [7, 11]. We denote the effective amplifier gain in frame f as G_f. Some cameras have a built-in AGC mechanism. It temporally changes the amplification of the image signal as a function of the detector irradiance.

Finally, the electric signal may be nonlinearly related to the image irradiance. We denote this radiometric response as $r(I)$. The nonlinearity is usually caused by nonlinear electronic amplifiers. Alternatively, in image intensifiers it can be caused by the nonlinear response of the phosphors [7, 11] of the intensifying tube. The graylevel value v in pixel $\vec{x}$ at frame f is:

$$v(\vec{x}) = r\left[G_f \tilde{I}(\vec{x})\right] = r\left[G_f M(\vec{x}) I(\vec{x})\right] \ . \tag{2}$$

The range of gray level values in the detector is $v \in [0, v_{\max}]$, where $v_{\max} = 255$ for 8 bit digital cameras.

Suppose that the radiometric response function $r(I)$, optical non-uniformity $M(\vec{x})$ and the gain G_f are unknown. We aim to estimate an image which is equivalent to the scene radiance in a wide FOV. In this paper we

- Estimate all the mentioned radiometric aspects, based only on an image sequence taken for mosaicing.
- Derive a fundamental limitation of this estimation.
- Exploit this estimation to such quality that seamless image mosaics can be created without resorting to feathering techniques.

3 Prior Explorations

Several studies have addressed some of the mentioned radiometric aspects. Methods to estimate the radiometric response based on a single image are derived in [6, 17]. Other methods for estimating the response [4, 8, 22, 29] use a dedicated process, in which the exposure settings are intentionally varied in time. However, many existing cameras have no manual control of exposure. Hence, methods that are based on multiple exposures are problematic in such cases. Other methods assume that the radiometric response has a specific parametric model [3, 6, 20, 23]. Recent advances have included estimation of the temporal gain [16, 21] variations. In any case, those studies have not dealt with spatial non-uniformities.

Estimation of spatial non-uniformity[1] has been done in [27, 31]. However, those methods recover neither the unknown radiometric response, nor the temporal gain variations. In contrast, all the above mentioned radiometric characteristics are considered in [13]. Nevertheless, Ref. [13] does not recover any of the functions. Rather, it concentrates on image enhancement aimed at reducing seams in image mosaics.

4 The Fundamental Equation

The camera radiometric response function r is typically monotonic. Therefore, r is invertible, hence we can write

$$r^{-1}\left[v(\vec{x})\right] = G_f M(\vec{x}) I(\vec{x}) \quad . \tag{3}$$

Thus,

$$\log\left\{r^{-1}[v(\vec{x})]\right\} = \log\left[M(\vec{x})\right] + \log(G_f) + \log\left[I(\vec{x})\right] \quad . \tag{4}$$

Define the functions

$$h\left[v(\vec{x})\right] \equiv \log\left\{r^{-1}[v(\vec{x})]\right\} \quad , \quad m(\vec{x}) \equiv \log\left[M(\vec{x})\right] \ , \\ g_f \equiv \log(G_f) \ \text{ and } \ i \equiv \log(I) \quad . \tag{5}$$

Hence, from Eq. (4)

$$h\left[v(\vec{x}_p^{f_0})\right] - m(\vec{x}_p^{f_0}) - g_{f_0} = i \quad , \tag{6}$$

where $\vec{x}_p^{f_0}$ denotes the pixel in frame f_0 onto which a scene point is imaged. Another way to express Eq. (6) is

$$(1 \ -1 \ -1) \cdot \left\{h\left[v(\vec{x}_p^{f_0})\right] \ \ m(\vec{x}_p^{f_0}) \ \ g_{f_0}\right\}^T = i \ . \tag{7}$$

This is a basis for a matrix formulation of the problem. Acquiring scene points in different frames provides many linear equations in the form of Eq. (7). This set of equations can be written as

$$\mathbf{B}\mathbf{s} = \mathbf{i} \ , \tag{8}$$

where

$$\mathbf{s} = \begin{bmatrix} \mathbf{h} \\ \mathbf{m} \\ \mathbf{g} \end{bmatrix} \ . \tag{9}$$

is the sought vector, comprising

$$\mathbf{h} = \left[h(0), \ \ldots h(v_{\max})\right]^T \quad , \\ \mathbf{m} = \left[m(1), \ \ \ldots m(\tilde{N})\right]^T \ \text{ and } \ \mathbf{g} = \left[g_1, \ \ \ldots g_F\right]^T . \tag{10}$$

Here F is the number of frames. The matrix $\mathbf{B}$ is sparse, while each row has some entries of "1" and "-1" at column indices that faithfully express Eq. (7). The vector $\mathbf{m}$ is a column-stack representation of $m(\vec{x})$. For general non-uniformity we set $\tilde{N} = N^2$. For radial non-uniformity [13, 14] we can set $\tilde{N} \approx N/\sqrt{2}$. We aim to estimate $\mathbf{s}$, while $\mathbf{i}$ is unknown, and subsequently recover an estimate of I.

5 Fundamental Ambiguities

5.1 Scale Ambiguity

Before proposing a practical solution method, we need to understand fundamental limitations that are associated with the problem. Suppose that $\mathbf{s}$ and $\mathbf{i}$ are vectors that are consistent with Eq. (8). We can bias each of the system function by individual offsets: $\hat{\mathbf{h}} = \mathbf{h} + c_h$, $\hat{\mathbf{m}} = \mathbf{m} + c_m$ and $\hat{\mathbf{g}} = \mathbf{g} + c_g$, defining a new vector $\hat{\mathbf{s}}$ in analogy to Eq. (9). Suppose we bias the unknown vector $\mathbf{i}$ to $\hat{\mathbf{i}} = \mathbf{i} + c_i$, where $c_i = c_h - c_m - c_g$. It is easy to see from Eq. (7) that $\mathbf{B}\hat{\mathbf{s}} = \hat{\mathbf{i}}$, hence satisfying the fundamental equation. For this reason, we can only determine M, G and $r^{-1}(v)$ up to unknown scale factors.

This is a problem that is common to *all* methods for radiometric calibration, that do not use quantitative knowledge about the scene irradiance. It is usually unimportant, since it means that the irradiance of the scene I is known up to a scale $C_I = \exp(c_i)$. We handle the ambiguity by normalizing the estimates such that $\max[\hat{r}(I)] = 1$, $\max[\hat{M}(\vec{x})] = 1$ and $\min(\hat{G}_f) = 1$.

5.2 Exponential Ambiguity

More interestingly, the estimation is prone to an exponential ambiguity. Suppose that $\mathbf{s}$ and $\mathbf{i}$ are vectors that are consistent with Eq. (8). We may multiply both sides of Eq. (8) by an arbitrary scale factor γ

$$\mathbf{B}\gamma\mathbf{s} = \gamma\mathbf{i} \ . \tag{11}$$

Hence, also the vector $\hat{\mathbf{s}} = \gamma\mathbf{s}$ is a solution to the problem, if we change the unknown vector $\mathbf{i}$ by $\hat{\mathbf{i}} = \gamma\mathbf{i}$. Using Eq. (5),

[1] Interestingly, spatial non-uniformities can be of benefit in image mosaicing, since they can improve the dynamic range of the imaging system [25, 27].

the estimated radiometric functions of the system are

$$\hat{r}^{-1}(v) = e^{\hat{h}(v)} = e^{\gamma h(v)} = \left[r^{-1}(v)\right]^{\gamma} \ , \qquad (12)$$

$$\hat{M}(\vec{x}) = e^{\hat{m}(\vec{x})} = e^{\gamma m(\vec{x})} = \left[M(\vec{x})\right]^{\gamma} \ , \qquad (13)$$

and

$$\hat{G}_f = e^{\hat{g}_f} = e^{\gamma g_f} = \left[G_f\right]^{\gamma} \ . \qquad (14)$$

Since γ is arbitrary, the solution is ambiguous. An analogous ambiguity was found in [8] for methods that simultaneously estimate the radiometric response function and exposure ratios (not treating spatial non-uniformity).

6 An Image Invariance

In Sec. 7 we detail a method for estimating $\hat{M}(\vec{x}), \hat{G}_f$ and $\hat{r}^{-1}(v)$. The nonuniformity/nonlinearity functions can be compensated for, in order to estimate the *bottom line*: the image irradiance I. However, as explained in Sec. 5, any estimate of $\hat{M}(\vec{x}), \hat{G}_f$ and $\hat{r}^{-1}(v)$ is prone to an ambiguity, given by Eqs. (12,13,14). Hence, we now describe how this ambiguity influences the estimation of image irradiance.

We express the estimated irradiance value $\hat{I}$ using Eq. (3),

$$\hat{I} = \frac{\left[r^{-1}(v)\right]^{\gamma}}{\left[M(\vec{x})G_f\right]^{\gamma}} = I^{\gamma} \ . \qquad (15)$$

This means that we may err in estimating the true value of the image irradiance by an exponential ambiguity. Note Eq. (15) behaves *exactly* like a γ-correction operation. To appreciate the significance of this result, recall that our entire analysis does not assume any parametric form for the response function $r(I)$. Yet, thanks to our process, the problem becomes much simpler: it eventually boils down to a parameterized problem, having only a *single* parameter, γ.

Thus, no matter how complicated the original response function is, and regardless of the temporal gain and spatial nonuniformities, the image resulting from our method is γ-distorted, at most. This is a novel type of *image invariance*: different cameras with very different radiometric characteristics and nonidealities eventually may yield, by our method, images that are essentially identical, up to a simple transformation.

7 Estimation

Recall that Eq. (6) is based on a measurement of a scene point in frame f_0. Due to camera motion, the same scene point is measured in frame f_1 at image pixel $\vec{x}_p^{f_1}$. Assuming that the scene radiance is static,

$$h\left[v(\vec{x}_p^{f_1})\right] - m(\vec{x}_p^{f_1}) - g_{f_1} = i \ . \qquad (16)$$

Assume that the image registration process [3, 9, 10, 12, 27, 28] is successful, i.e. that corresponding image pixels are registered. Then Eqs. (6,16) yield

$$h\left[v(\vec{x}_p^{f_0})\right] - h\left[v(\vec{x}_p^{f_1})\right] - m(\vec{x}_p^{f_0}) + m(\vec{x}_p^{f_1}) = -g_{f_0} + g_{f_1} = 0. \qquad (17)$$

Tracking some of the scene points in several frames[2] provides many linear equations as Eq. (17). This set of equations can be written as $\mathbf{Rs} = \mathbf{0}$. We weighted each of these equations according to the signal-to-noise ratio of the measurements that each is based on. We followed the weighting proposed in [19].

To better condition the estimation, we impose smoothness on $M(\vec{x})$, $r(I)$ and G_f. This yields an over-constrained system of equations, whose least squares solution is

$$\hat{\mathbf{s}} = \arg\min_{\mathbf{s}}(\mathbf{s}^{\mathrm{t}}\mathbf{A}^{\mathrm{t}}\mathbf{A}\mathbf{s}) \ , \qquad (18)$$

where

$$\mathbf{A} = \begin{bmatrix} \mathbf{R} \\ \lambda_h \mathbf{D}_{\mathrm{h}} \\ \lambda_m \mathbf{D}_{\mathrm{m}} \\ \lambda_g \mathbf{D}_{\mathrm{g}} \end{bmatrix} \ . \qquad (19)$$

Here $\mathbf{D}_{\mathrm{h}}$, $\mathbf{D}_{\mathrm{m}}$ and $\mathbf{D}_{\mathrm{g}}$ express Laplacians operating exclusively either on h, m or g, respectively. We used Laplacian operators as described in [19]. The parameters λ_h, λ_m and λ_g weight the penalty for unsmooth solutions relative to the penalty for disagreement with the data. Eq. (18) is solved by singular value decomposition (SVD).

Interestingly, simultaneous constant functions $h(v) = const_1$, $m(\vec{x}) = const_2$ and $g_f = const_3$ satisfy Eq. (17). A constant $h(v)$ implies $r^{-1}(v) = I_0$, where I_0 is a constant. This is a non physical solution for the radiometric response function, because in this case $r(I)$ is not defined for the full range of I, but rather only for a single value. We avoid such simultaneous constant functions by excluding the SVD solutions that correspond to a trivial $r^{-1}(v)$.

Handling the γ Ambiguity

The γ-ambiguity (γ-correction) mentioned in Sec. 6 is insignificant in applications that do not require quantitative photometric measurements. Such applications include display for human interpretation. In these cases, recall that the user can tune the display γ-correction for visual plausibility, and therefore there is no need to determine γ a priori.

A simple way to disambiguate the solution is to define a criterion for image quality, and then optimize it as a function

[2] At this stage there is no need for accurate image registration. The reason is that for the recovery of M, G and r we may decide to include only points that reside on smooth image regions, discarding pixels in proximity to edges. We discarded saturated pixels from the estimation.

of γ. As the criterion, we elected maximization of the image entropy, and used it in our experiments. Alternatively, Ref. [6] presents a blind γ estimation method based on minimization of high order frequency correlations. Moreover, we can alleviate the ambiguity by using other minor *a-priori* knowledge about either the camera non-uniformity, the radiometric response function or camera gain. In particular, knowledge about $M(\vec{x})$ at two distinct pixels is sufficient to resolve it.

8 Consistency and Seamless Mosaics

Despite the above mentioned γ-ambiguity, all the frames $\hat{I}$ resulting from Eq. (15) are *mutually consistent*. Their gray-level values have one-to-one correspondence to the true scene radiance. Since the raw frames $v_f(\vec{x})$ are obtained by camera motion, their compensated version $\hat{I}_f(\vec{x})$ can be used for creating a wide FOV image mosaic. Since the images are consistent, the resulting mosaic is seamless.

The consistency better facilitates accurate image registration [3, 21, 27]. The registered frames $\hat{I}_f$ are now combined such that each pixel has a single value. There are various methods for performing this combination [4, 15, 22]. We followed the one described in [19, 27]. We did *not* use any seam-removal feathering method.

9 Experiments

To demonstrate our method we performed experiments with real data, as well as simulations. First, we describe the simulations and then the results of experiments.

9.1 Simulations

To create a sequence of simulated images, we took a wide FOV image and divided it into a sequence of small frames. To simulate a 1D vignetting effect, each frame was multiplied by a known nonuniform function, e.g., $M(\vec{x}) = \exp\left[-(x/200)^2\right]$. We then multiplied each frame by a known gain $G_f \in [1, \exp(2)]$, which varied in time. We used the resulting frames as inputs to a simulated radiometric response function, e.g., $r(\tilde{I}) = \tilde{I}^{0.45}$. Fig. 3 shows several images which were created by this process. The known functions of one of the simulations are plotted as solid lines in Fig. 4.

We used the synthetic image sequence to obtain the estimates $\hat{r}^{-1}(I)$, $\hat{M}(x)$ and $\hat{G}_f$, plotted[3] as dashed lines in Fig. 4. Based of that, we recovered the compensated images $\hat{I}_f(x)$. Then we fused these images into a mosaic, as shown in Fig. 5. Note that no seams appear in Fig. 5, although we did not perform any process dedicated to seam removal. We checked our algorithm using a variety of functions $M(x)$, $r(I)$ and G_f. The results all resemble Figs. 4 and 5.

[3]Only for the display of the graphs we operate an exponential transformation on $\hat{r}^{-1}(v)$, $\hat{M}(x)$ and $\hat{G}_f$ using a single fitting parameter γ. We have not done this when creating the mosaics.

Figure 3. Sample frames based on a simulated spatial filter, time-variant gain and a nonlinear radiometric response function. Scene features become darker towards the frame periphery and have global brightness variations.

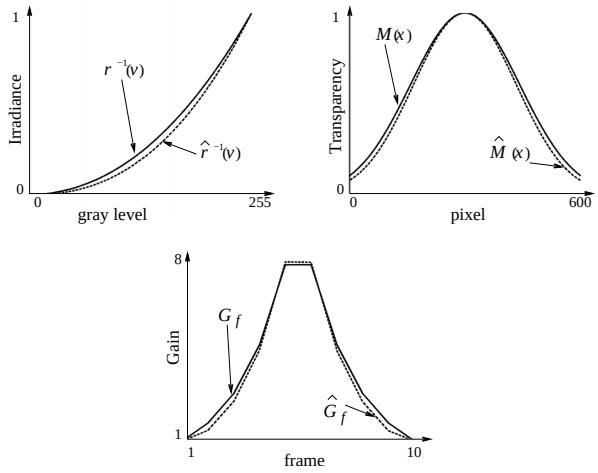

Figure 4. [Solid lines] The true inverted radiometric response $r^{-1}(v)$, true non-uniformity function $M(x)$ and true gain function G_f used in the image sequence, samples of which are shown in Fig. 3. [Dashed lines] The estimated functions.

9.2 Experiments using Real Images

Our experiments were performed using of Nikon D100 digital SLR camera and a PointGray Dragonfly firewire video camera with Sony ICX424 sensor. An advantage of the Nikon camera is its ability to acquire images in 12 bits using a linear radiometric response. This enables creation of the effect of arbitrary radiometric response functions in 8 bit images, in post-processing. Thus, the known radiometric response is helpful for assessing the effectiveness of our calibration method. To test the effect of amplifier AGC, we

Figure 5. A mosaic of the recovered images. No seams appear, although unknown vignetting and gain exist in the raw sequence, samples of which are shown in Fig. 3. Moreover, the recovered functions and images erred by an exponential function. No dedicated seam-removal method was applied.

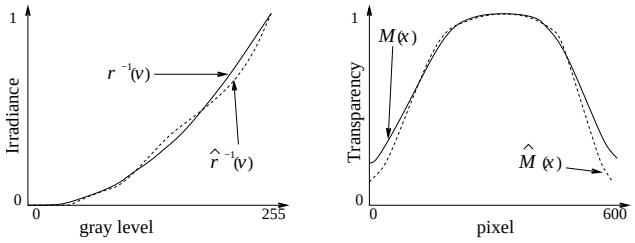

Figure 6. An experiment using real data. [Solid lines] The true inverted radiometric response and the true nonuniform function. [Dashed lines] The estimated functions.

complemented the SLR by the PointGray camera, which has AGC and an inherent native mode of gamma-correction.

To create a strong vignetting effect, we attached a spatially varying optical filter to the camera lens.[4] The filter attenuation of the irradiance varied along the x axis. In our experiments, we assumed that the effect of this 1D filter is much more significant than the intrinsic 2D vignetting of the lens.

Let us first consider the SLR-based experiment. The solid lines in Fig. 6 plot $r^{-1}(v)$ and $M(x)$ for this system. The function $M(x)$ was calibrated in a separate dedicated process, to enable future validation. Fig. 7 shows some sample frames from an acquired sequence. After acquisition, the frames were registered. We then created the matrix $\mathbf{A}$

[4]We placed the filter a few centimeters ahead of the lens. Had it been placed right next to the lens, it would have affected the aperture properties with little spatial variations in the image.

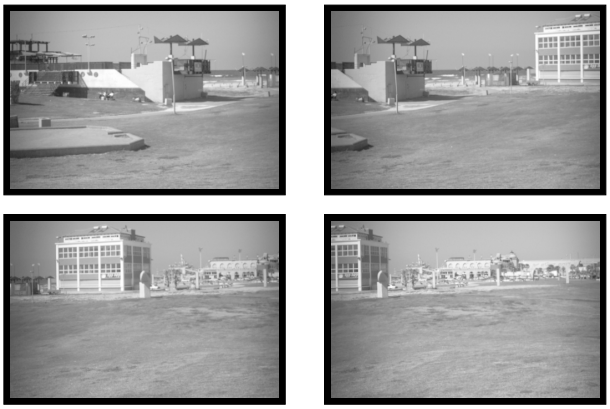

Figure 7. Image frames sampled from a sequence acquired by a Nikon D100 digital camera. A spatially varying optical filter had been attached to the camera lens. Scene features become darker towards the periphery of each frame.

according to Eq. (18). For this purpose, we selected 6000 random pairs of corresponding pixels in different frames. Using this data, we estimated $\hat{r}^{-1}(v)$ and $\hat{M}(x)$, which are plotted by dashed lines in Fig. 6.

Fig. 8 shows a mosaic resulting from the image combination process described in [19]. The non-uniformity was corrected, making all frames mutually consistent. We stress again, that we obtained seamless mosaicing without using any dedicated seam-removal methods like feathering etc. In contrast, Fig. 9 shows the image mosaic obtained without considering $M(x)$ and $r(I)$. Here, inconsistencies in the raw data create strong seams, whose removal would require a dedicated feathering process.

Analogous data and results based on the PointGray video camera are shown in Fig. 10. Here, the AGC was turned on. For this reason, the gain variations were estimated as part of the process and then compensated for. The estimation and mosaicing were based on 17 frames. Across most the resulting mosaic in Fig. 10, seams are unnoticeable.

10 Discussion

The presented method is simple and practical. It simultaneously leads to an estimate of the optical non-uniformity, the radiometric response function and the temporally varying gain. This estimation is prone to a simple ambiguity that can be resolved with very little prior knowledge. We demonstrated the capability of the estimation to yield seamless image mosaics.

We do *not* claim that our method makes standard methods for seam removal [2, 5, 18, 28] totally unnecessary. Those methods are of general scope, and can deal with cases in which seams are not caused by camera non-uniformity, but rather by scene and illumination dynamics. Neverthe-

Figure 8. An image mosaic based on analysis of a real experimental sequence, samples of which are shown in Fig. 7. No seams are apparent, although the original images suffered from an unknown vignetting, and the recovered vignetting and radiometric response erred by an exponential function. No dedicated seam-removal method was applied.

Figure 9. A mosaic based on a raw sequence, samples of which are shown in Fig. 7, without compensating for the estimated radiometric response function $\hat{r}(I)$ and optical non-uniformity $\hat{M}(\vec{x})$.

less, it is demonstrated that for static scenes, seams are removed simply by addressing their root cause.

There is still room for extensions to our framework. Some cameras have an auto-iris mechanism, which is analogous to AGC in the reaction to detector irradiance. An interesting example of this is the human eye, which has an auto-iris, a nonlinear radiometric response, and spatial non-uniformity of retina irradiance and sensitivity. We believe that in such a single-lens system, analysis as in this paper would be effective. However, most camera lenses consist of multiple components and stops. In these cases, the implications of an auto-iris may go beyond gain, since a smaller iris increases the depth of field. This may induce a slight sharpening of the vignetting distribution $M(\vec{x})$, modifying it temporally.

Acknowledgements

Yoav Schechner is a Landau Fellow-supported by the Taub Foundation, and an Alon Fellow. This work is supported by the German-Israeli Foundation (GIF) and the Ollendorff Minerva Center in the Dept. of Elect. Eng. at the Technion. Minerva is funded through the BMBF.

References

[1] M. Aggarwal and N. Ahuja, "A Pupil-Centric Model of Image Formation," Int. J. Computer Vision **48**, pp. 195-214 (2002).

[2] P. J. Burt and E. H. Adelson, "A multiresolution spline with application to image mosaics," ACM TOG **2**, pp. 217–236 (1983).

[3] F. M. Candocia, "On the featureless registration of differently exposed images," Proc. of Int. Conf. Imag. Sci. Syst. & Tech., Vol. I, pp. 163–169 (2003)

[4] P. E. Debevec and J. Malik, "Recovering high dynamic range radiance maps from photographs," Proc. SIGGRAPH, pp. 369–378 (1998).

[5] M. L. Duplaquet, "Building large image mosaics with invisible seam lines," Visual Information Processing VII, Proc. SPIE **3387**, pp. 369–377 (1998).

[6] H. Farid, "Blind inverse gamma correction," IEEE Trans. IP **10**, pp. 1428–1433 (2001).

[7] K. R. Fowler "Automatic gain control for image-intensified camera", IEEE Trans. on Instrumentation and Measurement **53**, pp. 1057–1064 (2004).

[8] M. D. Grossberg and S. K. Nayar, "Determining the camera response from images: What is knowable?", IEEE Trans. on PAMI **25**, 1455–1467 (2003).

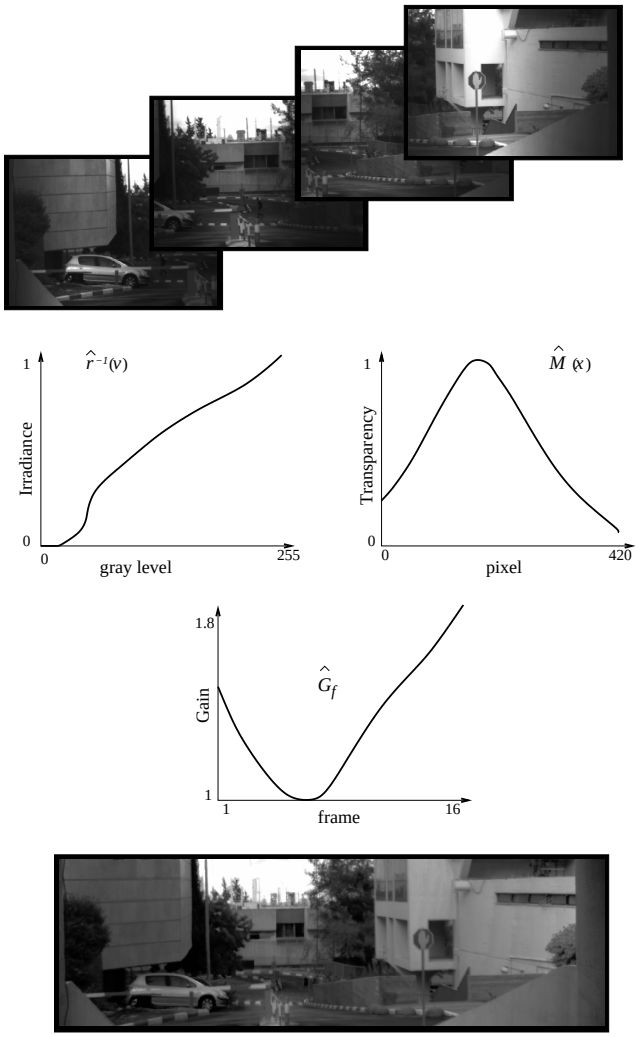

Figure 10. [Top] Image frames sampled from a sequence acquired by a PointGray camera with AGC. A spatially varying optical filter had been attached to the camera lens. Scene features become darker towards the periphery of each frame. [Middle] The estimated functions, including temporal gain variations. [Bottom] An image mosaic based on our analysis.

[9] M. Hansen, P. Anandan, K. Dana, G. van der Wal and P. Burt, "Real-time scene stabilization and mosaic construction," in Proc. Workshop on Appl. of Compt. Vis., pp. 54–62 (1994).

[10] S. Hsu, H. S. Sawhney and R. Kumar, "Automated mosaics via topology inference," IEEE Compt. Graph. and Appl. **22**, pp. 44–54 (2002).

[11] S. Inoue, *Video Microscopy,* Plenum Press , pp. 208–225 (1986).

[12] M. Irani, P. Anandan, J. Bergen, R. Kumar and S. Hsu, "Efficient representations of video sequences and their application," Signal Processing: Image Communications **8**, pp. 327–351 (1996).

[13] J. Jia and C. K. Tang, "Image registration with global and local luminance alignment," Proc. IEEE CVPR, Vol. I. pp. 156–163 (2003).

[14] S. B. Kang and R. Weiss, "Can we calibrate a camera using an image of a flat, textureless Lambertian surface?," Proc. ECCV Part. 2, pp. 640-653 (2000).

[15] S. B. Kang, M. Uyttendaele, S. Winder and R. Szeliski, "High Dynamic Range Video," Proc. SIGGRAPH, pp. 319-325 (2003).

[16] S. J. Kim and M. Pollefeys, "Radiometric alignment of image sequences", Proc. IEEE CVPR, Part. 1, pp. 645–652 (2004).

[17] S. Lin, J. Gu, S. Yamazaki and H. Shum, "Radiometric Calibaration from a Single Image," Proc. IEEE CVPR, Part. 2, pp. 938–946 (2004).

[18] A. Levin, A. Zomet, S. Peleg and Y. Weiss, "Seamless image stitching in the gradient domain," Proc. ECCV, Part IV, pp. 377–390 (2004).

[19] A. Litvinov and Y. Y. Schechner, "Radiometric framework for image mosaicking," JOSA A, **22** (2005). To be published.

[20] S. Mann, "Comparametric equations with practical applications in quantigraphic image processing," IEEE Trans. IP **9**, pp. 1389–1406 (2000).

[21] S. Mann and R. Mann, "Quantigraphic imaging: Estimating the camera response and exposures from differently exposed images," Proc. IEEE CVPR, Vol. 1, pp. 842–849 (2001).

[22] T. Mitsunaga and S. K. Nayar, "Radiometric self calibration," Proc. IEEE CVPR Vol. I, pp. 374-380 (1999).

[23] C. Pal, R. Szeliski, M. Uyttendaele and N. Jojic "Probability models for High Dynamic Range Imaging," Proc. IEEE CVPR (2004).

[24] A. Rav-Acha and S. Peleg, "Dynamosaics: Dynamic mosaics with non-chronological time," Proc. IEEE CVPR (2005).

[25] Y. Y. Schechner and S. K. Nayar, "Generalized mosaicing: Polarization panorama," IEEE Trans. PAMI **27**, pp. 631–636 (2005).

[26] Y. Y. Schechner and S. K. Nayar, "Generalized mosaicing: Wide field of view multispectral imaging," IEEE Trans. PAMI **24**, pp. 1334–1348 (2002).

[27] Y. Y. Schechner and S. K. Nayar, "Generalized mosaicing: High dynamic range in a wide field of view," Int. J. Computer Vision **53**, pp. 245-267 (2003).

[28] H. Y. Shum and R. Szeliski, "Systems and experiment paper: Construction of panoramic image mosaics with global and local alignment," Int. J. of Computer Vision **36**, pp. 101–130 (2000).

[29] Y. Tsin, V. Ramesh and T. Kanade, "Statistical calibration of CCD imaging process," Proc. IEEE ICCV, Vol. 1, pp. 480-487 (2001).

[30] R. G. Willson and S. A. Shafer, "What is the center of the image?," JOSA A, **11**, pp. 2946–2955 (1994).

[31] W. Yu, "Practical anti-vignetting methods for digital cameras," IEEE Trans. on Cons. Elect., **50**, pp. 975–983 (2004).