HDR Recovery under Rolling Shutter Distortions

Sheetal B Gupta, A N Rajagopalan
Department of Electrical Engineering
Indian Institute of Technology Madras, Chennai, India
{ee13s063,raju}@ee.iitm.ac.in

Gunasekaran Seetharaman
Information Directorate
AFRL/RIEA, Rome NY, USA
gunasekaran.seetharaman@us.af.mil

Abstract

Preserving the high dynamic irradiance of a scene is essential for many computer vision algorithms. In this paper, we develop a technique for high dynamic range (HDR) reconstruction from differently exposed frames captured with CMOS cameras which use a rolling shutter (RS) to good effect for reducing power consumption. However, because these sensors are exposed to the scene row-wise, any unintentional handshake poses a challenge for HDR reconstruction since each row experiences a different motion. We account for this motion in the irradiance domain by picking the correct warp for each row within a predefined search space. The RS effect is rectified and a clean frame is propagated from one exposure to another until we obtain rectified irradiance corresponding to all the exposures. The rectified irradiances are finally fused to yield an HDR map that is free from RS distortions.

1. Introduction

Real world scenes possess a far greater dynamic range than the images that capture them. The magnitude varies in the order of 4 on a unit log scale (to the base 10) which (obviously) cannot be accommodated in our digital frames. This is due to shortcomings of the camera sensors. When we capture an image, the camera lens refocuses all the scene information onto the sensor. This undergoes processing (both linear and non-linear) followed by quantization into 8 bit pixel values. Depending on the exposure time, the out-of-range irradiance values result in either underexposed or saturated intensities. Recovering this lost information (also referred to as HDR reconstruction) has increasingly drawn the attention of researchers [2, 3, 9, 11, 13] over the past decade. A prevalent approach is to extract scene irradiance by optimally combining information from multiple frames captured with different exposure times. This also involves estimating the camera response function (CRF) as part of the procedure. Another scheme is image fusion [10] that directly blends differently exposed intensity images. Based on quality measures such as saturation and contrast, it finds a weighted average of the input to yield the final output.

Though the above works are a good step forward, they make the unreasonable assumption that a pixel depicts a scene point at the same location in every frame. While a camera can be held still for low exposure frames, camera motion leading to alignment distortions is inevitable while capturing a series of differently exposed frames. This issue has been examined by a few researchers for cameras with CCD sensors. As these sensors acquire data all at once, a global motion model is commonly adopted. Lu et al. [8] discuss the problem of obtaining HDR images from motion blurred observations and jointly estimate the CRF, blur kernels and latent scene irradiance. Vijay et al. [16] solve for simultaneous deblurring and HDR imaging of a scene using a set of differently exposed and non-uniformly blurred images.

With CCD giving way to CMOS sensors in modern imaging devices, the focus of our work is on HDR for CMOS cameras. A major distinction between the two is in the shutter mechanism: while CCD sensors employ a global shutter, CMOS sensors systematically use an electronic rolling shutter (RS) where the rows of the sensor are scanned and read out at different times. Every row uses a common read-out circuit which reduces power consumption to less than 1/2 as compared to a CCD sensor of the same size. However, unlike the CCD, every row of a CMOS sensor experiences a different motion when there is relative motion between the camera and the scene. Depending on the camera path, this can cause RS induced distortions i.e., straight lines can appear bent. This renders motion estimation difficult since pixels belonging only to the same row have similar motion. For higher exposures, this may be accompanied by blur too.

The RS effect is a recent phenomenon and has already been addressed by some researchers. In [1] motion is estimated using L1 regularization by considering row-wise motion as high frequency and image motion as low frequency. In [4],[15] motions for key rows are estimated and linearly interpolated for the rest of the rows. Pichaikuppan et al.

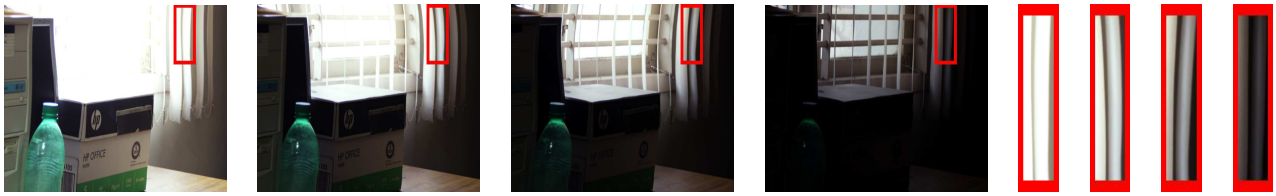

(a) (b) (c) (d) (e) (f) (g) (h)

Figure 1. Multiple exposure frames affected by rolling shutter effect.

[12] model rolling shutter along with motion blur considering each row of the distorted frame as a weighted average of warped versions of clean rows. They solve for change detection in the presence of rolling shutter and motion blur. It must be mentioned that these works assume identical exposure for all the images. Although there are no works that perform HDR imaging for RS affected images, there exists few schemes such as [5] that work at the sensor level. They control the read-out timing and the exposure length for each row by altering the logic of the control unit. In this paper, we address the problem of HDR image construction for differently exposed RS affected images. We solve for RS distortion frame-by-frame in the irradiance domain and simultaneously perform rectification. We use an existing algorithm to determine the CRF and use the inverse CRF to map intensity images to their respective irradiances. Starting with the lowest exposure, we map the first pair of irradiances to a common range and evaluate the motion between the two frames. We assume that the low exposure frames are free from RS distortion since the displacement of the camera in this short duration will be negligible. The computed motion is used to rectify the distorted irradiance so as to build a clean reference for the next higher exposure frame. This propagation of reference from one exposure to another yields rectified irradiances which are later fused together to construct the final HDR map of the scene.

We assume the scene to be distant enough to be considered as planar. Because the camera motion is small in HDR scenario, it can be well approximated by inplane translations with each row exhibiting a unique translation. Using a high aperture, we can use low exposure times which in turn results in negligible motion blur in the scene. A set of differently exposed frames captured using an RS camera is shown in Fig 1. Note that in the first image (a) and the last image (d) the region of well-exposed pixels is almost mutually exclusive. zoomed-in images (e-h) clearly depict the rolling shutter effect across different frames with different exposure.

The main contribution of this work is that *it is the first attempt of its kind for recovering an HDR image from rolling shutter affected images captured with CMOS cameras.* Also solving for RS distortions between two frames belonging to different exposures has not been attempted in the literature.

2. Rolling Shutter: Background

In this section, we explain in brief the working principle of the rolling shutter mechanism. In RS cameras, exposure for each row starts sequentially with a certain delay T_d. Each row is exposed for time T_e which is constant for all the rows but spans different intervals of the total exposure time. Thus each row is read at different times using the same read-out circuit.

Consider an image with M rows and N columns. If acquisition starts at time $t = 0$, then the i^{th} row is exposed during the time interval $[(i-1)T_d, (i-1)T_d + T_e]$. The total exposure of the frame is $(M-1)T_d + T_e$. The time delay T_d is usually fixed for a camera. Hence, varying the exposure time changes the duration for which each row is exposed. If T_e is very small, there will be no RS distortion in the images. This is because the displacement of the camera in such a short duration can be treated as negligible. This value of T_e can provide us with distortion-free low exposure frames. When T_e is increased (say to capture information from low-lit regions of the scene), the camera motion will introduce rolling shutter distortions in the image.

Let $\mathbf{f}$ be the irradiance corresponding to the clean frame (without motion) and $\mathbf{g}$ be the irradiance obtained while camera is in motion. In the above mentioned situation, the i^{th} row of the distorted irradiance $\mathbf{g}$ will correspond to the i^{th} row of the warped version of clean irradiance $\mathbf{f}$. If we consider the path followed by the camera as $u(t)$, then we can express each row of the distorted irradiance as

$$\mathbf{g}^i = \mathbf{f}^{(i)}_{\mathbf{u}((i-1)T_d)} \qquad \text{for } i = 1, 2, ..., M \tag{1}$$

where $\mathbf{f}^{(i)}_{\mathbf{u}((i-1)T_d)}$ is the i^{th} row of the warped version of $\mathbf{f}$ corresponding to the warp at time instant $(i-1)T_d$. Starting from $t = 0$, this is the i^{th} sample of continuous camera path u(t) considering T_d as sampling period. As the scene is distant and camera motion small, we can approximate the warp during the exposure $[(i-1)T_d, (i-1)T_d + T_e]$ by a single warp u($(i-1)T_d$)). If T_e is made very high, then the rolling shutter effect will also be accompanied with motion blur. In this paper we consider exposure times such that motion blur is minimal.

3. RS-HDR Imaging

If different exposure images affected by RS are directly subjected to HDR reconstruction, it will result in an HDR map with local artifacts due to unaccounted RS distortion. Hence, it becomes imperative to solve for the RS effect across the set of images. Starting with the lowest exposure, we consider consecutive pairs of images and evaluate camera motion for each row. This is then used to rectify the distorted frames which are subsequently used for final HDR reconstruction. We are motivated to perform all operations in the irradiance domain since the camera pipeline does not preserve linearity in the intensity domain. We employ [2] to estimate the CRF from a carefully chosen set of aligned images with different exposure settings. This enables us to determine the irradiance values corresponding to the intensities in an image using the inverse CRF. However, there are inherent difficulties in using these irradiance values when estimating motion between frames. This is because images corresponding to different exposures need not yield the same irradiance value at a given location. In fact, the values are preserved only when they fall in an admissible range as discussed next.

3.1. Irradiance Remapping

When sensors are exposed to a scene, exposure X which is the product of the irradiance E and exposure time Δt is collected on the film. Depending upon the exposure time, different ranges of irradiances are selected by the CRF and mapped to the intensity domain. Consider a scene with irradiance having double precision values ranging from 0 to R. This when multiplied by exposure time Δt acts as input to the CRF. Note from Fig. 2 that the CRF maps a definite range $[r_1, r_2]$ of X between 0 and 255. All values $\leq r_1$ are mapped to 0 while values $\geq r_2$ are mapped to 255. If we consider two frames captured with exposure times Δt_1 and Δt_2, the first frame when mapped to irradiance domain (using inverse CRF) will have values ranging from $\frac{r_1}{\Delta t_1}$ to $\frac{r_2}{\Delta t_1}$ while it will be $\frac{r_1}{\Delta t_2}$ to $\frac{r_2}{\Delta t_2}$ for the second frame. As the range corresponding to both the frames is different, motion estimation in rolling shutter affected images becomes increasingly difficult. One solution to this issue is to individually pick pixels belonging to the intersection in both frames and mask the others as zero. But this may lead to insufficient number of common irradiance values. To perform correct motion estimation, we set a function $\mathbf{m}$ which considers values that belong to the intersection of the two sets and maps all values of the frame within this range.

$$m(x) = \begin{cases} \mathbf{x} & k_1 \leq x \leq k_2, \\ k_1 & x < k_1, \\ k_2 & x > k_2. \end{cases} \tag{2}$$

Here k_1 refers to the $\max\left(\frac{r_1}{\Delta t_1}, \frac{r_1}{\Delta t_2}\right)$ and k_2 refers to $\min\left(\frac{r_2}{\Delta t_1}, \frac{r_2}{\Delta t_2}\right)$. Mapping all the values to a common range $[k_1, k_2]$ allows us to use these pixels for motion estimation. In the above scheme, the only pixels we neglect are the saturated pixels. Notice that in the common range $[k_1, k_2]$ these pixels are present only in the distorted frame i.e., the frames with higher exposures. We discard these values using a partial Tukey window. We create a 512 point Tukey window [16] and consider the second half as our mask to select the useful pixels in the distorted frame while giving low weights to the saturated pixels. To evaluate the warp, the values in $\mathbf{g}$ must be compared with the warped versions of $\mathbf{f}$ row-wise. Since $\mathbf{f}$ does not possess saturated pixels post-mapping, $\mathbf{m}(x)$ helps in arriving at a correct estimate of the motion between $\mathbf{f}$ and $\mathbf{g}$. The pixels of $\mathbf{g}$ which possess a higher weight when passed through the partial Tukey window are termed as 'useful pixels'.

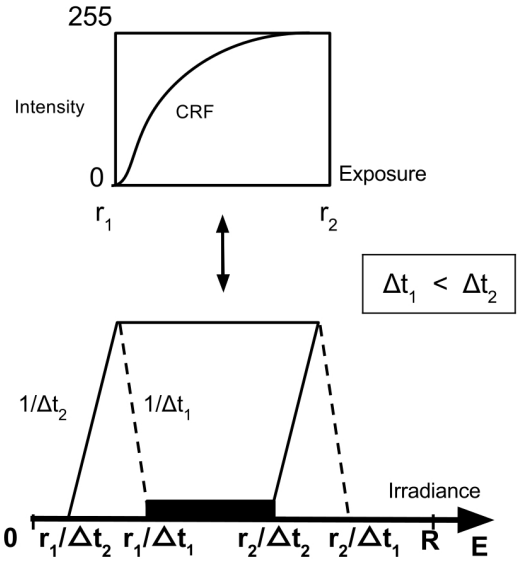

Figure 2. Irradiance to image mapping for different exposures.

3.2. Warp Estimation

We consider a block of rows around each row to ensure that an adequate number of useful pixels are employed to estimate the warp. Let this block be represented as B_i. Let $g_b^i \in \mathbb{R}^{P_i X 1}$ be the useful pixels in B_i (stacked as a column vector) where P_i is the number of useful pixels in B_i. We consider a search space S and construct a matrix $\mathbf{H_b^i}$ where each column of $\mathbf{H_b^i}$ contains the corresponding pixels in B_i of the warped version of the clean irradiance $\mathbf{f}$ for a pose $\tau_k \in S$. S is the discrete pose space that we define, and k ranges from 1 to $|S|$ where $|S|$ represents the number of poses. Thus, Eq. 1 can be represented as

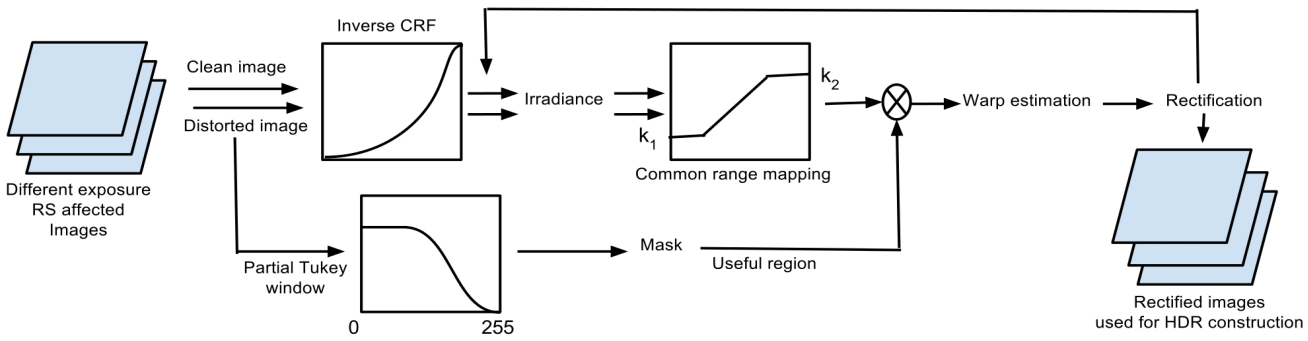

Figure 3. Irradiance estimation for different exposures of RS affected images.

$$\mathbf{g_b}^{(i)} = \mathbf{H_b}^{(i)}\boldsymbol{\omega}^{\boldsymbol{i}} \text{ such that } \|\boldsymbol{\omega}^{(i)}\|_0 = 1. \tag{3}$$

Here, $\boldsymbol{\omega}^{(i)}$ is a vector with a single non-zero entry that selects the correct warp in the search space defined by $\mathbf{H_b}^{(i)}$. To pick the warp in the search space with least error, we perform

$$\arg\min_k \ \|\mathbf{g}^{(i)} - \mathbf{f}_k^{(i)}\|_2^2 \tag{4}$$

Here $\mathbf{f}_k^{(i)}$ is the k^{th} column of $\mathbf{H_b}^{(i)}$ with k ranging from 1 to $|S|$ representing all the poses in the search space.

For the blocks in the distorted frame which possess the saturated region, very few useful pixels will be available for estimation of the warp. The pose chosen in this case by Eq. 3 may not be correct. We set a threshold to determine such blocks and ensure correct estimation by resorting to the following.

$$\widetilde{\boldsymbol{\omega}}^{(i)} = \arg\min_{\boldsymbol{\omega}^{(i)}} \left\{ \|\mathbf{g_b}^{(i)} - \mathbf{H_b}^{(i)}\boldsymbol{\omega}^{(i)}\|_2^2 + \lambda_1 \|\boldsymbol{\omega}^{(i)}\|_1 \right\} \quad \text{subject to } \boldsymbol{\omega}^{(i)} \succeq 0 \tag{5}$$

Instead of picking a single warp we allow ω^i to select multiple warps. This is done using least squares which forms the data term in Eq. 5. In addition to this, we impose an L1 prior on ω to ensure sparsity in the search space. The regularization parameter corresponding to this prior is chosen high enough so that it picks a single warp or utmost few warps to yield a centroid which is equivalent to the correct warp. When it selects multiple warps, the centroid pose is calculated as

$$\boldsymbol{\tau_c}^i = \frac{\sum \boldsymbol{\omega}_{\boldsymbol{\tau_k}}^{(i)} \boldsymbol{\tau_k}^i}{\sum \boldsymbol{\omega}_{\boldsymbol{\tau_k}}^{(i)}} \tag{6}$$

In order to ensure continuity of motion through the rows, we build a search space for each block considering a neighborhood around the estimated pose of the previous block. For every pair of frames we start with a block around the middle row since the first and last rows may yield wrong estimates due to loss of information at the boundaries. Consider $\mathbf{N_b} = \{\tau_c + ns : \tau_c - l < \tau_c + ns < \tau_c + l, n \in \mathbb{Z}\}$ where τ_c is the estimated pose from the previous block, l is the bound considered around the pose, and s is the step size. We initialize a large pose space for the block around the middle row and consider neighborhood $\mathbf{N_b}$ for the rest of the blocks.

3.3. Irradiance Rectification

Once the motion for each row has been estimated, we rectify the distorted frame using scattered interpolation in conjunction with the estimated camera motion. We apply inverse motion on every pixel of the RS image. As forward mapping from RS image can leave holes, we define a triangular mesh over this grid and resample over the regular grid using bicubic interpolation. The rectified irradiances are converted into images using the known CRF and given as input for irradiance estimation. Using Debevec and Malik's method [2] the different exposure rectified images are optimally combined to extract the scene irradiance. This is then remapped to match the low dynamic range of the display by tone mapping [14].

Summary of our algorithm: We begin by setting a threshold T_u for the exposure time. All frames with exposure duration less than or equal to T_u are considered to be free from rolling shutter distortion. Any frame from this set can be considered as clean reference for the next frame in the set. We consider frame 1 as clean frame and frame 2 as distorted frame. Both frames are converted to their respective irradiances using the known CRF. These irradiances are then mapped to a common range using Eq. 2 to yield new irradiance maps (say) f_r and g_r, respectively. In order to determine the useful region of the distorted frame g_r we consider overlapping blocks containing b (predefined

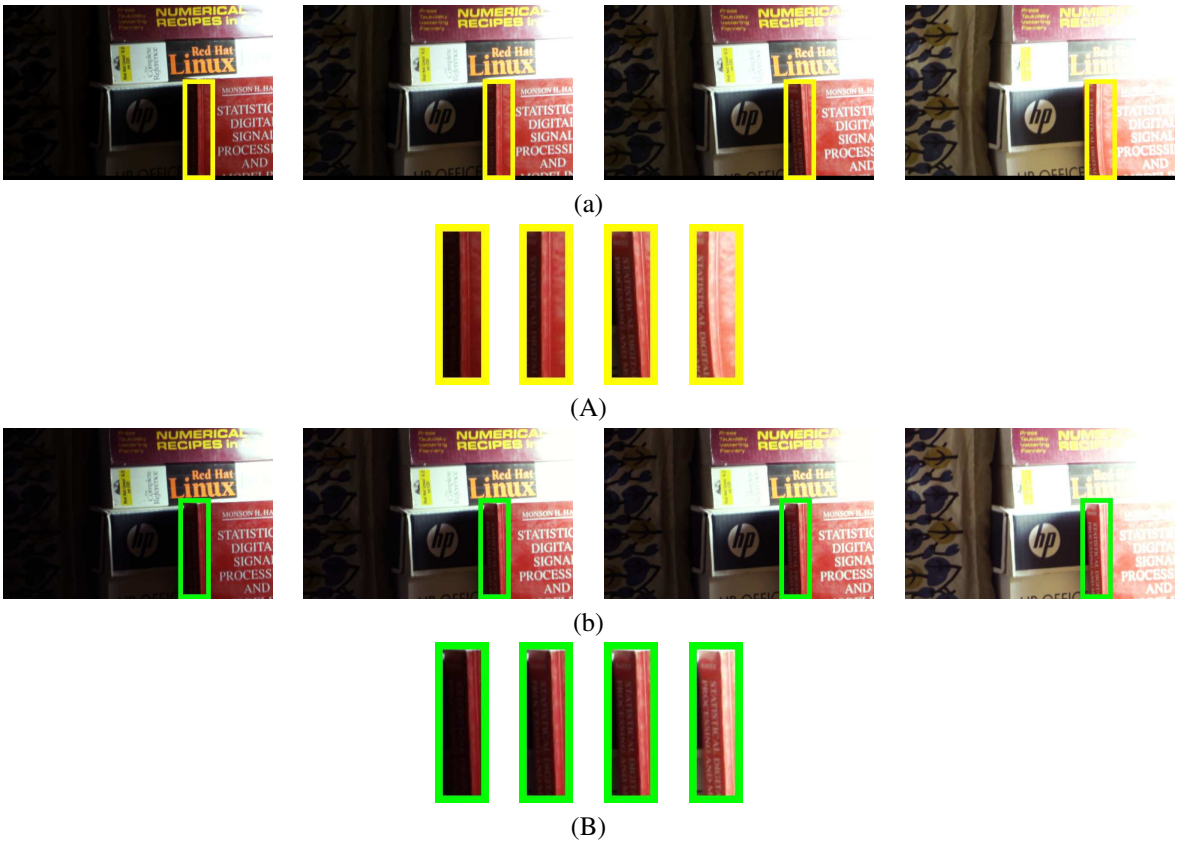

Figure 4. (a) Images captured with exposure times 0.05s, 0.08s, 0.1s and 0.2s. (A) Magnified portions of images in (a). (b) Rectified images corresponding to images in (a). (B) Close-up of corresponding regions from images in Fig.4(b).

number) rows around each row and pass it through a partial Tukey window. Starting with the block closest to the middle row, we consider a search space and find the corresponding block of warped versions of clean frame for every pose in the search space. The correct pose is estimated using Eq. 4. We use Eq. 5 when number of useful pixels is insufficient in the block. We set a threshold and evaluate warp for a block only if number of useful pixels in the block is greater than the threshold. Warps for the rest of the blocks are interpolated using the warps evaluated for these blocks. Once the warps are estimated for each row, the distorted frame is rectified using inverse warping to yield the reference for the successive frame in the set.

4. Experiments

To validate our algorithm we show results on both synthetic and real data. For the synthetic case, we begin by assuming a still camera to capture different exposure images without distortions. We then simulate the RS effect on the higher exposure frames for a given camera path. In the case of real experiments we use a Google Nexus 4 mobile camera to capture the desired images of different exposures affected by RS distortions.

4.1. Synthetic Experiments

We capture frames with different exposure times using Canon E60. We set up a scene containing a stack of books illuminated on the side using a table lamp. The CRF for this camera is derived using the code of [2]. The same camera settings were used in our experiments too. This also provides us with a clean irradiance of the scene. Using different camera paths for different images, we simulate the rolling shutter effect on the irradiance of the scene. We generate a camera path u(t) and fix a time delay that is used in order to yield samples u($(i-1)t_d$) for $i = 1..N$ (N corresponds to the number of rows in the image). The distorted frame is generated by considering the i^{th} row of the irradiance warped in accordance with the pose u(it_d) following Eq. 1. This is multiplied with the exposure time and converted into an image using the known CRF. Fig. 4(a) depicts the synthetically generated RS affected frames for different exposure times. We consider exposure times varying from 0.02 to 0.2 seconds. We set threshold as 0.02 seconds below which all frames are clean with no visible RS distortions. Fig. 4(A) contains the magnified portions of the images in

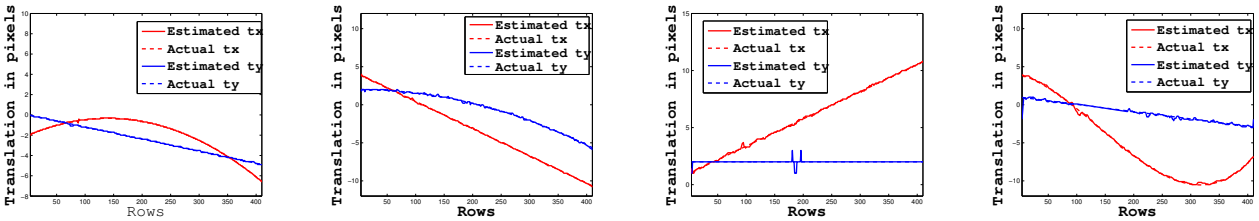

(b2) (c2) (d2) (e2)

Figure 5. (b2-e2) These plots display the actual and estimated camera trajectories of frames in Fig.4(b) with respect to clean images in Fig.4(a).

(a) (b) (c) (d) (e) (f)

Figure 6. (a) Irradiance obtained without RS rectification. (b) Result of [6]. (c) Irradiance obtained using our method. (d-f) Maginified regions of Fig.6(a-c) respectively.

Fig. 4(a) to reveal the presence of RS effect in the observations.

We set the size of block as 3 rows and initialise the search space for the middle block in each image. We then input these images to our algorithm that estimates the correct camera pose for each row and each image. The estimated camera trajectories are compared with the actual paths as shown in Fig.5. These warps are used to rectify the distorted frames to get clean images with different exposures in Fig. 4(b). We can clearly observe from Fig.4(B) that the curved lines have been correctly rectified (the edge of the book appears the same and straight in all the zoomed-in patches). The rectified images are used for irradiance estimation and displayed after tone mapping.

In contrast, if we give RS affected images as input for irradiance reconstruction using [2] we get the result in Fig. 6(a). Notice in Fig. 6(d) that the RS effect results in multiple curved edges or artifacts around a single edge. Fig. 6(b) shows results obtained with [6] but without applying the ghosting effect since our scene is static. This method aligns the images using SIFT and constructs HDR using a generalized weighted filtering technique. Even though SIFT is used, note that it is unable to remove the RS distortion completely. This is understandable as [6] is not designed to solve for RS distortions. Our result is displayed in Fig.6(c) after tone mapping while corresponding zoomed-in patch is shown in Fig. 6(f) which clearly reveals the strength of our method in rectifying the RS effect. This serves to underline the importance of our work.

4.2. Real Experiments

Fig.7 shows results on two different real data captured using Google nexus 4 with exposure values ranging from -2 to 2 with -2 being the lowest exposure (without any distortion). The RS distortion can be clearly seen in the close-up of patches in Fig. 7 (A) and Fig. 7(C) for the first and second example, respectively. The lower exposure value frames contain information which appear saturated in the higher exposures. On the other hand, high exposures contain information not visible in the low exposure frames. These images were given as input to our algorithm for HDR reconstruction. The intermediate results which are the rectified images are displayed in Fig. 7(b) and Fig. 7(d). The close-up patches in Fig. 7(B) and Fig. 7(D) display the rectification results obtained for each exposure frame. We can clearly observe in these images that all the edges are correctly aligned. These rectified images are then used to obtain the irradiance of the scene which is displayed in Fig. 8(c1) and Fig. 8(c2). In both cases, we observe the result is devoid of saturation effects and the intensities are proportional to the actual irradiance values in the scene.

We also compare our results with the methods in [2] and [6]. Using [2] we obtain the results shown in Fig. 8(a1) and Fig. 8(a2). These images show significant artifacts. In Fig. 8(b1) and Fig. 8(b2) we display the results obtained using [6]. When we compare these outputs, we observe that our scheme consistently outperforms these competing methods by delivering rectified irradiances that are free from RS distortions.

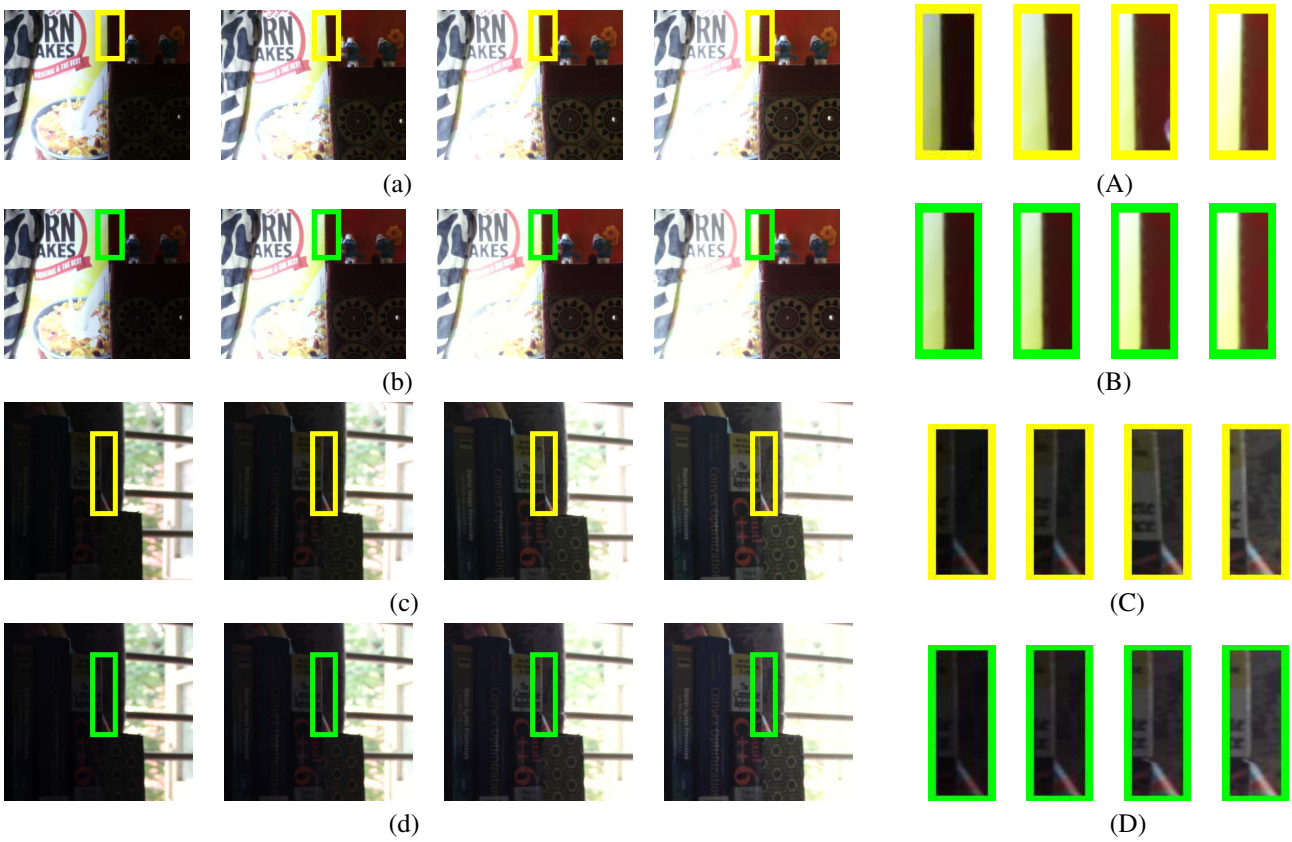

Figure 7. (a) Input for the first real example. (b) Rectified images for (A). (A) and (B) Close-up of patches in (a) and (b), respectively. (c) Input for second real example. (d) Rectified images for (C). (C) and (D) zoomed-in patches of (c) and (d), respectively.

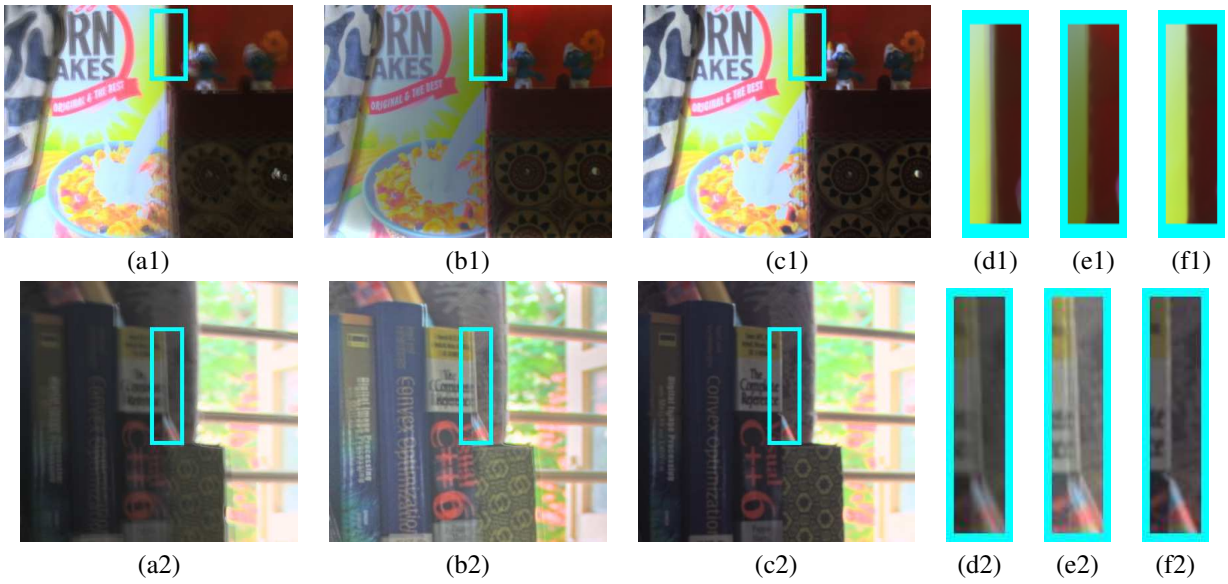

Figure 8. First real example: (a1) Irradiance obtained without RS rectification (b1) Result of [6]. (c1) Irradiance obtained using our algorithm. (d1-f1) Maginified regions of (a-c), respectively. Second real example:(a2) Irradiance obtained without RS rectification. (b2) Result of [6]. (c2) Irradiance obtained using our method. (d2-f2) Maginified regions of (a2-c2) respectively.

4.3. Algorithm complexity and runtime

For each block of rows in the image we get the correct warp either using Eq. 4 or Eq. 5 dependng upon the number of useful pixels present. We pick the correct warp in a single iteration for Eq. 4 whereas we use a gradient projection based approach to solve the L1-minimisation problem in Eq. 5 using SLEP [7]. It requires a sparse matrix-vector multiplication with order less than $O(P_i|S|)$ and projection into a subspace with dimension P_i in each iteration. Here P_i is the number of useful pixels whereas $|S|$ is the number of poses. Running time for (single image) using an unoptimised MATLAB code without any parallel programming on a 3.3GHz PC with 16GB RAM is 932 seconds of which evaluating motion requires 725 seconds whereas rectifying the RS affected image requires 207 seconds.

5. Conclusions

In this paper, we developed an HDR technique that combines information from different exposure frames captured using CMOS cameras. Due to row-wise acquisition, the higher exposure frames tend to exhibit RS distortion due to incidental camera motion. We solved for RS distortions on a frame-by-frame basis by bringing different exposure frames within the same region of scene irradiance and estimating the motion for each row of a distorted frame with respect to a previously derived clean frame. This involved rectification and propagation of clean frames from one exposure to another. The rectified irradiances were fused to get an RS-free HDR image. The method was validated on real as well as synthetic examples. Scope exists to extend this work to include the effect of motion blur too.

Acknowledgments

A part of this work was supported by a grant from the Asian Ofce of Aerospace Research and Development, AOARD/AFOSR. The support is gratefully acknowledged. The results and interpretations presented in this paper are that of the authors, and do not necessarily reect the views or priorities of the sponsor, or the US Air Force Research Laboratory.

References

[1] S. Baker, E. Bennett, S. B. Kang, and R. Szeliski. Removing rolling shutter wobble. In *Computer Vision and Pattern Recognition (CVPR), 2010 IEEE Conference on*, pages 2392–2399. IEEE, 2010. 1

[2] P. E. Debevec and J. Malik. Recovering high dynamic range radiance maps from photographs. In *ACM SIGGRAPH 2008 classes*, page 31. ACM, 2008. 1, 3, 4, 5, 6

[3] R. Fattal, D. Lischinski, and M. Werman. Gradient domain high dynamic range compression. In *ACM Transactions on Graphics (TOG)*, volume 21, pages 249–256. ACM, 2002. 1

[4] P.-E. Forssén and E. Ringaby. Rectifying rolling shutter video from hand-held devices. In *Computer Vision and Pattern Recognition (CVPR), 2010 IEEE Conference on*, pages 507–514. IEEE, 2010. 1

[5] J. Gu, Y. Hitomi, T. Mitsunaga, and S. Nayar. Coded rolling shutter photography: Flexible space-time sampling. In *Computational Photography (ICCP), 2010 IEEE International Conference on*, pages 1–8. IEEE, 2010. 2

[6] Y. S. Heo, K. M. Lee, S. U. Lee, Y. Moon, and J. Cha. Ghost-free high dynamic range imaging. In *Computer Vision–ACCV 2010*, pages 486–500. Springer, 2011. 6, 7

[7] J. Liu, S. Ji, and J. Ye. *SLEP: Sparse Learning with Efficient Projections*. Arizona State University, 2009. 8

[8] P.-Y. Lu, T.-H. Huang, M.-S. Wu, Y.-T. Cheng, and Y.-Y. Chuang. High dynamic range image reconstruction from hand-held cameras. In *Computer Vision and Pattern Recognition, 2009. CVPR 2009. IEEE Conference on*, pages 509–516. IEEE, 2009. 1

[9] S. Mann and R. Picard. On being undigitalwith digital cameras: Extending dynamic range by combining differently exposed pictures, 7 pages. 1

[10] T. Mertens, J. Kautz, and F. Van Reeth. Exposure fusion. In *Computer Graphics and Applications, 2007. PG'07. 15th Pacific Conference on*, pages 382–390. IEEE, 2007. 1

[11] T. Mitsunaga and S. K. Nayar. Radiometric self calibration. In *Computer Vision and Pattern Recognition, 1999. IEEE Computer Society Conference on.*, volume 1. IEEE, 1999. 1

[12] V. R. A. Pichaikuppan, R. A. Narayanan, and A. Rangarajan. Change detection in the presence of motion blur and rolling shutter effect. In *Computer Vision–ECCV 2014*, pages 123–137. Springer, 2014. 2

[13] E. Reinhard, W. Heidrich, P. Debevec, S. Pattanaik, G. Ward, and K. Myszkowski. *High dynamic range imaging: acquisition, display, and image-based lighting*. Morgan Kaufmann, 2010. 1

[14] E. Reinhard, M. Stark, P. Shirley, and J. Ferwerda. Photographic tone reproduction for digital images. In *ACM Transactions on Graphics (TOG)*, volume 21, pages 267–276. ACM, 2002. 4

[15] E. Ringaby and P.-E. Forssén. Efficient video rectification and stabilisation for cell-phones. *International Journal of Computer Vision*, 96(3):335–352, 2012. 1

[16] C. S. Vijay, P. Chandramouli, and R. Ambasamudram. Hdr imaging under non-uniform blurring. In *Computer Vision–ECCV 2012. Workshops and Demonstrations*, pages 451–460. Springer, 2012. 1, 3