

High Dynamic Range Panoramic Imaging

Manoj Aggarwal Narendra Ahuja
University of Illinois at Urbana-Champaign
405 N. Mathews Ave, Urbana, IL 61801, USA
{manoj,ahuja}@vision.ai.uiuc.edu

Abstract

Most imaging sensors have a limited dynamic range and hence can satisfactorily respond to only a part of illumination levels present in a scene. This is particularly disadvantageous for omnidirectional and panoramic cameras since larger fields of view have larger brightness ranges. We propose a simple modification to existing high resolution omnidirectional/panoramic cameras in which the process of increasing the dynamic range is coupled with the process of increasing the field of view. This is achieved by placing a graded transparency (mask) in front of the sensor which allows every scene point to be imaged under multiple exposure settings as the camera pans, a process anyway required to capture large fields of view at high resolution. The sequence of images are then mosaiced to construct a high resolution, high dynamic range panoramic/omnidirectional image. Our method is robust to alignment errors between the mask and the sensor grid and does not require the mask to be placed on the sensing surface. We have designed a panoramic camera with the proposed modifications and have discussed various theoretical and practical issues encountered in obtaining a robust design. We show with an example of high resolution, high dynamic range panoramic image obtained from the camera we designed.

1. Introduction

In many imaging applications including surveillance, navigation and photography, there is a need to capture high resolution images of a scene over large fields of view. In addition, it is desirable that the large brightness variation present in a real-world scene is captured without many areas being too dark (clipped) or too bright (saturated). The range of brightness levels that can be recorded by a sensor without clipping or saturation is often referred to as the dynamic range. A conventional digital camera, however, is limited in all three aspects, namely, resolution, field of view and dy-

amic range. The field of view is constrained by the sensor size and focal length of the lens (e.g. $30^\circ \times 22^\circ$ on a 16mm lens with $2/3''$ CCD sensor), the resolution is determined by the number of pixels on the sensor (640×480 for NTSC cameras) and the sensor and camera electronics determine the dynamic range. Most CCD sensors provide only 8-bits of brightness information, which are usually insufficient to capture the entire brightness variation in real scenes. The three limitations of a conventional imaging system have been examined individually and in pairs in various contexts in the literature. However, the problem of acquiring images with high resolution, large field of view and high dynamic range is still open and is the subject of this paper. We will first review the existing techniques which individually overcome the limitations on field of view and dynamic range, with or without high resolution, and then present the proposed approach.

Large field of view: There are two main classes of large field of view cameras: panoramic and omnidirectional. Omnidirectional cameras are capable of covering a field of view upto $360^\circ \times 360^\circ$, for instance over a sphere or a hemisphere. Most of the omnidirectional cameras proposed in literature use fish-eye lenses [16, 23, 26] or hyperbolic/parabolic mirrors [4, 8, 18, 24, 25] to image a large field of view onto a single sensor. These camera are capable of delivering images at video-rate, however, resolution of the images is constrained by that of the sensor. Recently, Coorg et. al. [5] and Nayar [19] proposed a high resolution omnidirectional camera. The higher resolution is obtained by sequentially acquiring multiple images covering different parts of the visual field and then stitching these images together, a process known as mosaicing. These designs trade the ability to get video-rate images for higher resolution images, and are better suited for stationary scenes.

Panoramic cameras cover upto a 360 degree cylindrical strip of the scene. The existing high resolution panoramic cameras just as the omnidirectional counterparts are based on mosaicing a sequence of images.

The images can be obtained either in parallel using multiple cameras [17] or sequentially by panning the camera [1, 10, 14, 21, 22]. The panoramic cameras based on multiple sensors can deliver images at video rate, however, they require precise arrangement of sensors and other optical elements. The panoramic cameras based on sequential capture require a single camera, but are not video-rate and hence are suitable for only stationary scenes.

Dynamic Range: Ordinarily, the dynamic range of any camera is given by that of the sensor used. This is particularly disadvantageous for omnidirectional and panoramic cameras since larger fields of view have larger brightness ranges. The dynamic range of a sensor can be improved through a number of techniques. The basic idea is to acquire multiple images using different exposure settings, thus capturing different sections of the illumination range. These multiple exposure images are then combined (intensity space mosaicing) to cover a larger brightness range which is a union of those covered by the individual images. There are two main types of high dynamic range cameras, depending on whether they have video-rate capability or not. The first type sequentially acquires multiple exposure images, each of which has a resolution identical to that of the sensor. These camera do not have video-rate capability and are suited for only stationary scenes [3, 6, 12, 13, 15]. The second type acquires images at video-rate, which can be achieved using several techniques, each having different tradeoffs. A number of special cameras have been disclosed in patents, which employ a single lens but multiple sensors such that the same scene gets simultaneously imaged on different sensors, preset to different exposure settings. Many of these designs typically require precise alignment of various optical elements and the sensors. In another approach, a special sensor is reported which has multiple sensing elements with different sensitivities for recording light for every pixel on the sensor. This approach uses a single sensor, but requires specialized and expensive hardware design. A special high dynamic range sensor has been proposed in [2], that instead of measuring the amount of charge accumulated in a pixel, measures the time it takes to reach saturation. The recorded times are used to convert them into a high dynamic range image. All the approaches discussed above generate high dynamic range images with resolution that of the sensor. A technique to trade off spatial resolution and quality for dynamic range is presented in [20]. N images with different exposures are captured on a single sensor at $(\frac{1}{N}th)$ the resolution of the sensor. Interpolation is used to construct a full-resolution image. The field of view of all the camera designs reviewed above is that of the individual images.

2. High dynamic range panoramic imaging

A straightforward approach to obtain a high resolution, panoramic image with high dynamic range would be to construct a high dynamic range image of the (narrow) field of view at each angular position of the camera using intensity space mosaicing, and then stitch together the individual high dynamic range images using spatial mosaicing. Alternatively, panoramic images generated under multiple exposure settings can be fused together by intensity space mosaicing to obtain a high dynamic range panoramic image. Either variant of this simple approach requires two controlling mechanisms, one for camera rotation and other for altering exposure settings.

We describe a simple modification to any existing panoramic/omnidirectional cameras in which the process of improving the dynamic range is coupled with the process of increasing the field of view. Our method requires only a single control mechanism for camera rotation, with no need for expensive hardware modifications or additional processing requirements. The basic idea is to capture every scene point under different exposure settings as the camera rotates. The desired change in exposure setting has been achieved using a filter with spatially varying transmittance (mask) as also done in [20]. However, the spatial variations in exposures suggested in this paper are different both in form and functionality from that in [20].

2.1. Proposed Approach

Consider a mask with spatial variation in transmittance as shown in Fig. 1(a), placed adjacent to the sensor of a high resolution panoramic camera based on rotation about lens center and mosaicing [14](Fig. 1(c)). The darker regions of the mask represent low transmittance or equivalently low exposure setting for the underlying portion of the sensor and brighter regions of the mask represent higher exposure. We will refer to the regions of common transmittance on the mask as bands. As the camera rotates, the image of a scene point falls at a different location on the sensor as shown in Fig. 2. Since the transmittance of the mask varies along the direction of camera rotation, every scene point visible during the camera scan, gets imaged under different exposure settings as the camera rotates. The sequence of images obtained for each scene point can then be combined to construct a high dynamic range panoramic image. The design of this camera parallels the design of the Non-frontal imaging camera (Nicam) [11]. In Nicam, the sensor is tilted which allows a scene point to be imaged under different focus settings as the




Figure 1. Spatially varying masks for large field of view high dynamic range image and their placement (a) Mask with transmittance varying along the rotation direction, suitable for panoramic cameras, (b) A circularly symmetric mask, better suited for omnidirectional high resolution imaging. (c) Top and side view of the proposed camera with the mask and the motor.

camera rotates. The resulting images can be used to construct omnifocus panoramic images. In our method, we instead use a graded filter to allow a scene point to be imaged under different exposure settings as the camera rotates.

The key idea in this approach is to exploit the rotation of the camera which is anyway necessary to obtain a large field of view (with high resolution), to obtain images of a scene point under different exposure settings as well. Depending on the field of view to be covered, a cylindrical strip (panoramic camera) or a hemisphere/sphere (omnidirectional cameras), the required variation in exposure can be realized by a suitable mask. For example, a omnidirectional camera in which the camera is panned and tilted to cover a sphere [5], or the high resolution omnidirectional camera proposed in [19], a circularly symmetric mask might be more suitable, such as one shown in Fig. 1(b).

3. Design of a high dynamic range panoramic camera

There are a number of practical and theoretical issues involved in the complete design of the camera, which we will discuss in the various subsections.




Figure 2. The rays from an object point O intersect the sensor at different locations for different angular positions ($P_1 - P_4$) of the camera. Since the mask adjacent to the sensor has spatially varying transmittance, any visible object point gets imaged under multiple exposure settings.

3.1. Preparation and placement of mask

A mask can be accurately implemented using several techniques summarized in [20]. These include placing a suitable mask adjacent to the detector array or using special sensors with different sensing elements having sensitivities corresponding to the mask. These techniques require very special handling and could be difficult to implement. For example, most solid state sensor have a protective glass cover, which is difficult to remove without destroying the sensor and hence a mask cannot be easily placed adjacent to the sensor. Our approach does not require placing the mask adjacent to the sensor. It is also robust to irregularities in the band boundaries, and misalignments of the bands with the sensor grid. A mask can be simply implemented by manually attaching a combination of neutral density thin-film filters on the sensor's protective surface. For example, if only a two band mask is required, with transmittance values 100% and 25%, a filter with 25% transmittance covering approximately half of the sensor's glass surface would suffice.

A consequence of such a simple implementation of the mask is that the effective spatial transmittance is not identical to that on the mask, but in fact a blurred version of the mask. The phenomenon is explained in Fig. 3. It shows cones of light from three scene points, undergoing refraction in the lens, passing through the mask and then converging on the sensor. The intercept of the cone with the mask varies with the angular position of the scene point. For example, for scene point O_1 , the area of intersection lies entirely in band A. The

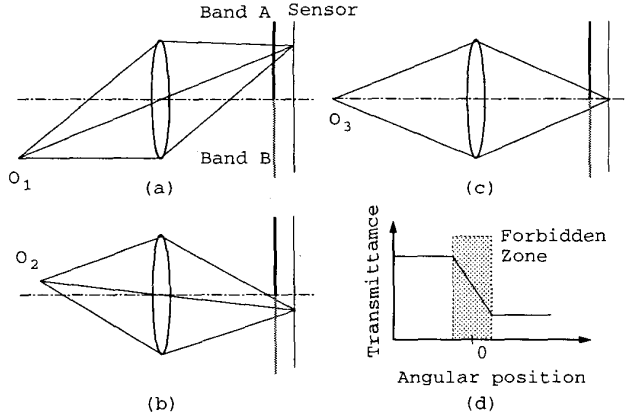


Figure 3. The effect of placing the mask at a distance from the sensor. (a)-(c) show how the light cone from the object point intercepts the mask after refraction. In (a) and (b), the intercept is confined to a single band, but in (c), it straddles both bands. (d) A plot of the resulting amount of light that reaches the sensor as a function of angular position of the object point.

cone of scene point O_2 barely misses intercepting the mask at band boundary. While, the cone corresponding to point O_3 intersects the mask on the boundary of two bands. This implies, the amount of light reaching the sensor is given by a weighted combination of the intercept area in band A and band B. The resulting amount of light that reaches the sensor as a function of the angular position of the scene point is plotted in Fig. 3(d). We will refer to the transition regions in the effective transmittance as forbidden zones. The size of the forbidden zones depends on the aperture size and the distance of the mask from the sensor. As we will show in Sec. 3.2, the angular step size of the motor can be suitably chosen to ensure each scene point gets imaged through each band entirely at least once, which is sufficient to create a high dynamic range mosaic.

Another practical issue in the construction of the mask is that the transition between bands may be quite rough on a micron scale (Typical pitch of pixels on a CCD sensor is 5-10 microns). This edge irregularity will affect the effective spatial transmittance near the band boundaries and may result in a slightly larger forbidden zone.

3.2. Angular step size of the motor

Consider a mask and camera configuration in which each forbidden zone spans at most k columns on the

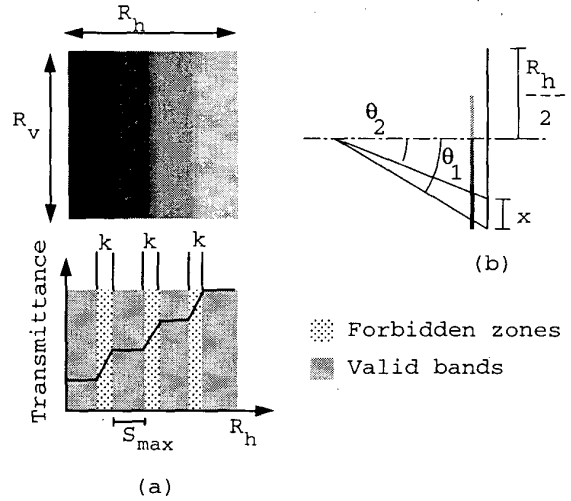


Figure 4. Forbidden and valid zones on the sensor and estimation of angular step for the camera

sensor. If we choose the angular step of the motor such that every scene point gets imaged at least once in non-forbidden zone of each band, then it is possible to reconstruct a high dynamic range image which is not affected by the irregularities or blur in the transmittance of the mask.

Let the horizontal resolution of the sensor be R_h and the number of exposure settings (bands) required to cover the desired dynamic range be n . It is clear from Fig. 4, that in order to ensure that every scene point gets imaged at least once for each of the exposures, the distance between the images of the same point for two successive positions of the camera must not exceed s_{max} , where

$$s_{max} = \left(\frac{R_h}{n} - k \right).$$

Since the image of a point shifts the maximum near the edge of the sensor, the maximum angular step size or the speed for the motor must be small enough to ensure an edge pixel does not shift by more than s_{max} units. This constraint can be used to easily determine the motor step size. From Fig. 4(b) we have

$$\tan \theta_1 = \frac{R_h}{2v} \quad \text{and} \quad \tan \theta_2 = \frac{R_h - 2x}{2v}.$$

Then, the difference $(\theta_1 - \theta_2)$ gives the maximum angular step for the motor. For all step sizes smaller than the maximum step, it is ensured that every scene point gets imaged at least once for each desired exposure (band).

3.3. Spatial variation in the sensitivity of the sensor

Photometric sensitivity refers to the ratio of image irradiance to radiance of a scene point towards the entrance pupil. Even in the absence of the graded mask, the sensitivity of a lens is spatially variant and decays in directions off the optical axis. In real lenses, various factors contribute to this decay including lens foreshortening, lens aberrations and vignetting, which are difficult to estimate theoretically without any knowledge of the internal structure of the imaging lens [9]. The impact of spatially varying sensitivity is that the irradiance at the sensor from the same scene point but at different angular positions of the camera are different. A correction factor which is an inverse of the sensitivity variation is thus necessary in order to normalize the recorded intensity values to a common scale. The sensitivity is further altered by the presence of the mask, and the final sensitivity is a product of the (blurred) transmittance of the mask and the inherent sensitivity of the lens and sensor combination. Since the dependence of irradiance from different directions is a complex function of lens design and imaging parameters, we chose to estimate the net variation in sensitivity empirically through a simple experiment described below.

We used a spatially uniform light source as the object and observed its image on the sensor equipped with a mask. The variation in the image gives the desired (net) variation in the sensitivity of the sensor. We used a light box (KLV7000, www.hakubausa.com) as the light source. The image of the light source was observed to change on repositioning the light box, which implied that the light source was not as uniform as we desired. We added light diffusing elements (plexi-glass and opal sheets) to the front of the lightbox, re-tested it for uniformity and used the combination as the light source.

Since, the mask creates a significant variation in the amount of the light reaching the sensor, a single image of the uniform light source may not be adequate to accurately estimate the net spatial variation in the sensitivity of the entire sensor. To overcome this problem, we acquire multiple images with different shutter speed settings (provided on a camera) such that each pixel on the sensor gets imaged within the reliable intensity range in at least one of the images. We define reliable range to be the range of intensity/luminance levels whose minimum is above the noise levels and maximum is below the saturation level minus the noise variance of the camera. We then acquire a image of a dark scene (with lens cap on) and label it as the "darkimage". For CCD cameras, this darkimage usually has non-zero mean and is a consequence of dark current present in CCD's [7]. We subtract the darkimage from each of the images and select the reliable portions. For each image,




Figure 5. The spherical grid and coordinate system

the intensities in the selected regions are scaled by dividing them by the shutter time for that image. Since each pixel falls in the reliable intensity range in at least one of the images, at the end of the procedure, we will have a number associated with each pixel on the sensor. This array of numbers after dividing by the maximum value among them will give the variation in the sensitivity across the sensor. The scaling factor required in the mosaicing process is obtained by taking an inverse of the sensitivity values.

The technique suggested above is applicable for linear cameras. In general, the relationship between the sensor irradiance and recorded intensity is nonlinear. A number of techniques have been suggested in literature to determine this nonlinearity [6, 15] and these can be adapted to estimate the scaling function required to account for the nonlinearity and the spatial variation in the photometric sensitivity.

3.4. Imaging geometry, calibration and image acquisition

The camera-lens system we used was perspective with slight distortion. The camera was mounted on motor with the rotation axis passing through the lens center, to ensure a common viewpoint at all angular positions of the motor. The camera rotates with a maximum angular step size as determined in Sec. 3.2 and the manner in which the camera covers the field of view is shown in Fig. 2.

We chose a spherical coordinate system (ϕ, θ) as shown in Fig. 5. Let the ray joining a object point to the lens center be denoted as $R_{\phi, \theta}$. In order to construct the mosaic, we need to find the intersection of the ray $R_{i, j}$ from object at (ϕ, θ) with the sensor for any given angular position (γ) of the camera. If $\gamma = 0$, the problem is identical to regular camera calibration which has been addressed by several researchers, even for lenses

with distortion. For $\gamma = 0$, let the relationship between the coordinates (x, y) of the image of an object point at angular position (ϕ, θ) be given by $x = f(\phi, \theta)$ and $y = g(\phi, \theta)$. Since, the camera is rotating about the lens center, the ray joining the object point to the lens center is invariant and rotating the camera amounts to rotating the coordinate system. Therefore, for non-zero gamma, $x = f(\phi - \gamma, \theta)$ and $y = g(\phi - \gamma, \theta)$. The relationship between the angular position γ in degrees to the motor count is known and documented in the manuals for the motor.

Using the above geometric relationship, we can compute the position of the object point in the image for different positions of the camera. This eliminates the need to track a scene point in an image sequence.

3.5. Constructing the high dynamic range panoramic mosaic

Let us label the sequence of images obtained at regular angular intervals (γ_k) of the camera by I_k . We construct a spherical mosaic surface aligned with the object space' coordinate system. We refer to the mosaic surface as \mathcal{M} , which is sampled uniformly along both ϕ and θ axis. The sampling rate is chosen based on the sensor resolution.

Let us assume that the sensor resolution in the horizontal and vertical directions is R_h and R_v , respectively, and the field of view of the individual images be $F_h \times F_v$. Then the sampling rate along the ϕ and θ axes are chosen to be $\frac{R_h}{F_h}$ and $\frac{R_v}{F_v}$, respectively. The range along ϕ axis is given by the total field of view covered by the rotating camera and can be as large as 360° . The range along θ axis is given by F_v . Consider a point on the spherical grid (ϕ_i, θ_j) and the ray $R_{i,j}$ joining that point to the lens center. The ray $R_{i,j}$ would intersect the sensor only in some of the images. Further, the intercept on the sensor varies as a function of the angular position of the camera (Fig. 2), which can be computed using the projection equations obtained by camera calibration as discussed in Sec. 3.4.

The intensity for point $\mathcal{M}_{i,j}$ is obtained as follows. We find the points of intersection of ray $R_{i,j}$ in all the images. Images in which $R_{i,j}$ does not intersect the sensor, or intersects within forbidden bands are disregarded. We note that the rotation step size for the camera motor has been chosen such that each scene gets imaged once for each exposure, without counting the situations when the point lands in the forbidden zones. The ray $R_{i,j}$ may not intersect the images at sensor grid points, therefore, the intensity value for $\mathcal{M}_{i,j}$ needs to be determined using interpolation. In any single image not all pixels in the neighborhood around a point of intercept may be within reliable intensity range. To

avoid using clipped or saturated neighborhood values to determine $\mathcal{M}_{i,j}$ we interpolate from only the reliable neighborhood pixels but from *all* selected images where $R_{i,j}$ intercepts, instead of any single best image.

Interpolation weights for reliable neighborhood pixels are computed as fractions representing the relative distance of the intersection point to the pixels. The intensities for the reliable pixels are scaled using the brightness correction functions determined in Sec. 3.3. An interpolation scheme, such as cubic or bilinear interpolation can then be used to determine the intensity value corresponding to grid point $\mathcal{M}_{i,j}$. This procedure is repeated for all grid points on the mosaic surface yielding the desired mosaic. It is worth noting that the interpolation is a necessary step for mosaicing a panoramic image on a spherical grid and is not a consequence of our method.

4. Experiments

We designed a grayscale panoramic camera with a mask as shown in Fig. 6. The mask consisted of a 25% transmittance filter (neutral density Kodak Wratten gelatin filter) covering left half of the sensor. The camera used was Pulnix TM720, with a 16mm Navitar lens (DO 1614) focused to infinity and aperture setting 5.6. There was one forbidden zone of width 100 columns (we chose to be conservative). The rotation step used for the experiments was approximately 3 degrees which was equivalent to a 100 pixel shift along the ϕ axis on the spherical grid.

We present a sample high dynamic panoramic images acquired with the camera. Fig. 7 displays results for an experiment on an indoor scene. The scene had significant light variation, bricks were generally darker, while the regions around the staircase and the poster boards much brighter. The first four images (a) - (d) display a subset of the images in the acquired sequence. The left halves of these images faithfully capture the brighter areas of the scene, while the right halves capture the darker areas. Fig 7(e) displays a panoramic of the same scene under identical camera settings, but without the mask. The high dynamic range panoramic image constructed from the sequence is shown in Fig. 7(f). The image has been histogram equalized to compress the dynamic range to 256 gray levels for the purpose of display. Figs. 7(g)-(l), show some sections of the high dynamic range panoramic image, which have been magnified and shown under different photometric scales to demonstrate the detail captured by the camera in both dark and bright regions.




Figure 6. (a) The mask used in the prototype camera. The darker band was a neutral density filter with 25% transmittance. the lighter portion represents absence of filter (b) The resulting transmittance variation as observed on the sensor.

References

- [1] R. Benosman, E. Deforas, and J. Devars. A new catadioptric sensor for the panoramic vision of mobile robots. In *Workshop on Omnidirectional Vision*, pages 112–116, 2000.
- [2] V. Brajovic and T. Kanade. A sorting image sensor: An example of massively parallel intensity-to-time processing for low latency computational sensors. In *IEEE Conference on Robotics and Automation*, pages 1638–1643, April 1996.
- [3] P. J. Burt and R. J. Kolczynski. Enhanced image capture through fusion. In *International Conference on Computer Vision*, pages 173–182, 1993.
- [4] J. S. Chahl and M. V. Srinivasan. A complete panoramic vision system, incorporating imaging, ranging, and three dimensional navigation. In *Workshop on Omnidirectional Vision*, pages 104–111, 2000.
- [5] S. Coorg, N. Master, and S. Teller. Acquisition of a large pose-mosaic dataset. In *Conference on Computer Vision and Pattern Recognition*, pages 872–878, 1998.
- [6] P. E. Debevec and J. Malik. Recovering high dynamic range radiance maps from photographs. In *Proceedings of ACM, SIGGRAPH*, pages 369–378, 1997.
- [7] G. E. Healey and R. Kondepudy. Radiometric CCD camera calibration and noise estimation. *IEEE Transactions on Pattern Analysis and Machine Intelligence*, 16(3):267–276, March 1994.
- [8] R. A. Hicks and R. Bajcsy. Catadioptric sensors that approximate wide-angle perspective projections. In *Workshop on Omnidirectional Vision*, pages 97–103, 2000.
- [9] C. Kolb, D. Mitchell, and P. Hanrahan. A realistic camera model for computer graphics. In *Proceedings of SIGGRAPH*, pages 317–324, 1995.
- [10] A. Krishnan and N. Ahuja. Panoramic image acquisition. In *Conference on Computer Vision and Pattern Recognition*, pages 379–384, 1996.
- [11] A. Krishnan and N. Ahuja. Range estimation from focus using a non-frontal imaging camera. *International Journal of Computer Vision*, 20(3):169–185, December 1996.
- [12] B. C. Madden. Extended intensity range imaging. Technical Report MS-CIS-93-96, Grasp Lab, University of Pennsylvania, 1996.
- [13] S. Mann and R. W. Picard. On being ‘undigital’ with digital cameras: Extending dynamic range by combining differently exposed pictures. In *Proceedings of IS&T 46th Annual Conference*, pages 422–428, May 1995.
- [14] L. McMillan and G. Bishop. Plenoptic modeling: An image based rendering system. In *Proceedings of SIGGRAPH*, pages 36–46, August 1995.
- [15] T. Mitsunaga and S. K. Nayar. Radiometric self calibration. In *Conference on Computer Vision and Pattern Recognition*, volume 1, pages 374–380, June 1999.
- [16] K. Miyamoto. Fish eye lens. *Journal of Optical Society of America*, 64:1060–1061, August 1964.
- [17] V. Nalwa. A true omnidirectional viewer. Technical report, Bell Laboratories, February 1996.
- [18] S. K. Nayar. Catadioptric omnidirectional camera. In *Conference on Computer Vision and Pattern Recognition*, pages 482–488, 1997.
- [19] S. K. Nayar and A. Karmarkar. 360 × 360 mosaics. In *Conference on Computer Vision and Pattern Recognition*, volume 2, pages 388–393, June 2000.
- [20] S. K. Nayar and T. Mitsunaga. High dynamic range imaging: Spatially varying pixel exposures. In *Conference on Computer Vision and Pattern Recognition*, volume 1, pages 472–479, June 2000.
- [21] S. Peleg. Panoramic mosaics by manifold projection. In *Conference on Computer Vision and Pattern Recognition*, pages 338–343, June 1997.
- [22] H. S. Sawhney and S. Ayer. Compact representations of videos through dominant and multiple motion estimation. *IEEE Transactions on Pattern Analysis and Machine Intelligence*, 18(8), August 1996.
- [23] Y. Xiong and K. Turkowski. Creating image-based VR using a self-calibrating fisheye lens. In *Conference on Computer Vision and Pattern Recognition*, pages 237–243, 1997.
- [24] Y. Yagi and S. Kawato. Panorama scene analysis with conic projection. In *International Workshop on Intelligent Robots and Systems*, pages 181–187, 1990.
- [25] K. Yamazawa, Y. Yagi, and M. Yachida. Omnidirectional imaging with hyperboloidal projection. In *International Conference on Intelligent Robots and Systems*, pages 1029–1034, July 1993.
- [26] S. Zimmerman and D. Kuban. A video pan/tilt/magnify/rotate system with no moving parts. In *IEEE/AIAA Digital Avionics Systems Conference*, pages 523–531, 1992.




Figure 7. Experimental results of high dynamic range panoramic imaging of an indoor scene. (a)-(d) Four images from the sequence acquired by a panoramic camera with a mask shown in Fig. 6. (e) A panoramic image constructed using a camera without the mask that shows many areas overexposed (f) The high dynamic range panoramic constructed using the sequence obtained from a camera with the mask. The image has been histogram equalized to show that the entire range of scene radiance was successfully captured. (g)-(l) Magnified results for some sections of resulting panoramic image. The brightness range within each individual section has been appropriately compressed between 0-255 graylevels to best display the intensity variations captured in that section.