

Flexible Mirror Imaging

Sujit Kuthirummal and Shree K. Nayar
Columbia University
{sujit,nayar}@cs.columbia.edu

Abstract

The field of view of a traditional camera has a fixed shape. This severely restricts how scene elements can be composed into an image. We present a novel imaging system that uses a flexible mirror in conjunction with a camera to overcome this limitation. By deforming the mirror, our system can produce fields of view with a wide range of shapes and sizes. A captured image is typically a multi-perspective view of the scene with spatially varying resolution. As a result, scene objects appear distorted. To minimize these distortions, we have developed an efficient algorithm that maps a captured image to one with almost uniform resolution. To determine this mapping we need to know the shape of the mirror. For this, we have developed a simple calibration method that automatically estimates the mirror shape from its boundary, which is visible in the captured image. We present a number of examples that demonstrate that a flexible field of view imaging system can be used to compose scenes in ways that have not been possible before. This flexibility can be exploited in applications such as video surveillance and monitoring.

1. Cameras with Flexible Fields of View

Traditional cameras have fields of view of fixed shapes – rectangular for perspective lenses and circular for wide-angle lenses. This severely restricts how scene elements can be composed into an image. Therefore, an imaging system that provides control over the shape and size of the field of view (FOV) would be desirable. Such an imaging system would enable a photographer or an application to capture scenes in unconventional ways – including only the scene elements of interest and excluding all other elements. Thus, image pixels (a fixed resource) are devoted to capturing only what is desired. This is particularly important in the case of video, where the number of pixels is always limited by the fact that individual pixels must be large enough to collect sufficient photons within a short integration time. Also, for dynamic scenes, this flexibility would enable one to continuously vary the FOV as objects of interest move around. Such an ability is of value in video monitoring applications.

In order to have control over the shape of the FOV, the camera must include some form of flexible optics. In principle, it is possible to develop lenses with adjustable shapes. While such elements have been proposed for finer optical ad-

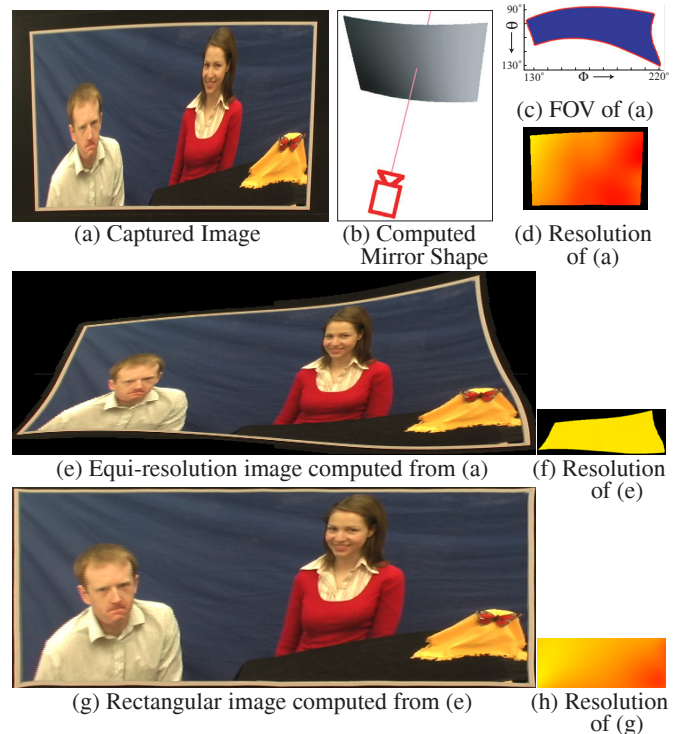


Figure 1. (a) Image captured by our flexible field of view imaging system. The FOV was shaped to include the butterfly (while excluding the scene region above it) and to include the person on the left. (b) Shape of the mirror estimated from its boundary in the captured image. (c) FOV captured in the image in (a), shown using spherical panorama coordinates – θ (elevation angle) and ϕ (azimuth angle). (d) Resolution of the image in (a), where resolution increases from red to yellow. (e) Equi-resolution image computed from the one in (a). (f) Resolution of the image in (e). (g) Rectangular image computed from the one in (e). (h) Resolution of the image in (g).

justments such as auto-focusing [16], they are hard to realize for our purpose of achieving a wide range of FOVs. We propose the use of a flexible mirror placed within the FOV of a perspective camera. In its normal state the mirror is planar and hence the camera’s FOV remains unchanged, except that it is rotated due to reflection by the mirror. By manually applying forces to the boundary of the mirror, we can generate a wide and continuous range of smoothly curved mirror shapes. Each shape results in a new FOV enabling us to capture the scene in a new way.

In order to analyze a captured image, we need to know the shape of the flexible mirror when the image was taken. Since the mirror shape can vary from one image to the next, it must be estimated from the captured image itself. Towards this end, we have developed a simple calibration method that estimates the 3D mirror shape from the 2D shape of its boundary, which is visible in a captured image. Once we know the mirror shape, we can determine different properties of the image, including its FOV and spatially varying resolution.

Figure 1(a) shows an image captured by our system. The bottom-right corner of the mirror was flexed to include the butterfly (while excluding undesirable scene elements above it) and the left side of the mirror was flexed to include the person on the left. Figure 1(b) shows the shape of the mirror estimated by our calibration method. The resulting FOV is shown in Figure 1(c) and can be seen to have a rather complex shape. Note that, in principle, one could have captured all the desired scene elements in an image using a wide-angle lens. However, in that case, since the shape of the FOV is fixed, the camera would capture scene elements that are not needed, i.e., image pixels would be ‘wasted’ on undesired elements. In contrast, in our system, we can adjust the shape of the FOV, so that image pixels are devoted to capturing only the scene regions of interest. Due to the non-planar shape of the mirror, the captured image is typically multi-perspective with varying resolution, as seen in Figure 1(d).

Since the images captured by our system are multi-perspective with varying resolution, they appear distorted (see Figure 1(a)). We propose that these distortions can be minimized by mapping the captured image to an equi-resolution image – an image in which all pixels subtend the same effective FOV (solid angle). We have developed a simple and efficient algorithm that uses the automatically determined mirror shape to compute this mapping. Our algorithm was applied to the captured image in Figure 1(a) to obtain the equi-resolution image in Figure 1(e). The resolution of this image is shown in Figure 1(f) and is seen to be uniform. As one can see, the equi-resolution image has a non-rectangular shape, which might be undesirable. A second mapping that employs thin-plate splines [11] is used to warp the equi-resolution image to a rectangular one without introducing significant resolution variation. The computed rectangular image and its resolution map are shown in Figures 1(g) and 1(h), respectively. This example demonstrates that our system enables us to capture complex scene compositions with minimal distortion. Note that since our algorithm digitally resamples a captured image, it does not improve on the captured image’s optical resolution.

2. Related Work

A wide variety of catadioptric systems (imaging systems that use lenses and mirrors) have been developed in fields as diverse as robotics, computer vision, computer graphics, and astronomy. Wide-angle imaging is a popular application and

some of these systems also capture equi-resolution images [2, 3, 9]. However, in all these systems the mirror shapes are fixed and hence their FOVs are not flexible. Recently, in [8, 13] it has been suggested that a flexible mirror can be emulated using a planar array of planar mirrors, where each planar mirror can be arbitrarily oriented. However, such a system is yet to be implemented as there are no mirror arrays available that provide the required level of control over mirror orientation. In astronomy, to remove the effects of atmospheric turbulence, some telescopes use adaptive optics [15] in the form of membrane mirrors that can change shape. Since such mirrors are used to adapt to phase changes in the incoming radiation, their shapes vary over very small ranges and cannot be used to vary FOV. Our system is similar in spirit to the flexible camera array of [14], in which many cameras are mounted on a flexible sheet and the sheet is flexed to capture a scene using an unconventional collective FOV. However, their objective is different. They combine images from all their cameras to create a collage of the scene, while we map an image captured by a single camera to a seamless undistorted composition of the scene.

Most techniques for calibrating catadioptric systems assume some knowledge of the shape of the mirror [5, 10, 12]. Since the mirror shape can vary in our system, we cannot make this assumption. In [17] the shape of a mirror is determined by analyzing how it reflects a known calibration target, while in [6] a method is presented to calibrate a generic imaging system by capturing a number of images of an active calibration target. Since in our system the mirror can assume different shapes in different images, such approaches would mean repeating the complete calibration process for each image – a task that is cumbersome for still images and impractical for videos with dynamic FOVs. To address these issues, we present a method that determines the 3D shape of the mirror from its 2D boundary in the captured image.

Since an image captured by our system is multi-perspective with varying resolution, it has distortions. Recently, a framework was proposed to characterize such distortions [21]. One approach to minimize these distortions is to use user-specified scene priors [19]. To handle arbitrarily complex and unknown scenes, we propose to minimize these distortions by mapping a captured image to an equi-resolution image. This mapping problem may be formulated as an equi-areal projection of a sphere onto a plane [18]. However, as we will show, for some imaging configurations this approach cannot be used due to ‘folding’ of the FOV. To address this problem, we have developed an alternative algorithm to compute equi-resolution images.

3. Capturing Flexible Fields of View

The prototype system we built to capture flexible FOVs is shown in Figure 2. It comprises of a Panasonic PV-GS180 camcorder that captures the scene reflected in a flexible planar mirror sheet. The camera captures 720×480 resolu-




Figure 2. Prototype system that captures flexible FOVs.

tion video. The mirror sheet¹ is 465×355 mm and made of acrylic. This sheet is mounted on another flexible plastic sheet whose center is attached to the metallic frame. To deform the mirror, we simply apply pressure on its edges or corners. We have used such a large mirror since optical aberrations due to the bending of the mirror decrease with mirror size and hence are less noticeable in the captured image. Also, it is a convenient size for the user to manually flex.

When we take a picture, the camera’s optical axis is usually parallel to the horizontal. To mimic this, when the mirror is not deformed (it is planar), we would like the effective viewpoint of our system to have a forward looking FOV. In our system, this is done by elevating the camera’s optical axis by 45° and tilting the plane of the mirror by 22.5° , as shown in Figure 2. The horizontal FOV of the camera is 28.22° and the mirror is 60 cm from the camera’s optical center.

As stated earlier, we estimate the 3D mirror shape from the shape of its boundary in a captured image. To ensure that we can detect the mirror boundary robustly, we pasted a trapezoidal border onto the mirror that is black with a thin white strip on the inside². This border can be seen in all the captured images. The mirror boundary, marked in Figure 3(a), is automatically detected by searching from the sides of the image inwards for a transition from black to gray. To represent the shape of this boundary we choose 32 equidistant points on each ‘horizontal’ side and 24 on each ‘vertical’ side. The coordinates of these 112 points are concatenated to form a 1D descriptor, \mathcal{D} , of the mirror boundary.

4. Determining Flexible Mirror Shape

To analyze a captured image, we need to know the shape of the mirror when the image was taken. Here, we describe how the mirror shape can be estimated from its boundary in a captured image.

4.1. Modeling Mirror Shape

The flexible mirror sheet can be deformed into various curved surfaces. These surfaces are assumed to be smooth and we represent them using a tensor product spline similar to how [7, 20] model a surface. Let $M(u, v)$ be the 3D location of the point on the mirror seen by image pixel (u, v) ,

¹Purchased on EBay for \$20.

²A trapezoidal border ensures that when the mirror is not deformed the reflective portion of the mirror appears as a rectangle in a captured image.




Figure 3. (a) One of the frames used for the off-line calibration of the mirror boundary to shape mapping. (b) Errors in mirror shape estimation. Please see text for details.

where (u, v) are normalized image coordinates. We represent the z-coordinate M_z of this 3D point as

$$M_z(u, v) = \sum_{i=1}^{K_f} \sum_{j=1}^{K_g} c_{ij} f_i(u) g_j(v). \quad (1)$$

Here, f_i and g_j are 1D spline basis functions, c_{ij} are the spline coefficients, and K_f and K_g are the number of spline basis functions used along the u and v dimensions of the captured image. The unknowns here are the spline coefficients c_{ij} which would be different for different deformations of the mirror. Once $M_z(u, v)$ is known, the 3D mirror point is given by $M(u, v) = [u \ v \ 1] M_z(u, v)$. In our experiments, f_i and g_j are quadratic spline basis functions and we use 6 such functions along each of the two spatial dimensions.

4.2. Off-line Calibration

It has been shown that for developable surfaces (like our planar flexible sheet) the 3D boundary can be used to determine the 3D surface [4]. However, in our case, only the 2D (image) shape of the boundary is known from the captured image and we need to determine the 3D mirror shape from it. More specifically, we need a mapping \mathcal{M} between the descriptor \mathcal{D} of the mirror boundary (described in Section 3) and the spline coefficients c_{ij} in Equation 1: $c_{ij} = \mathcal{M}(\mathcal{D})$.

Since it is not clear that a closed-form solution to the mapping, \mathcal{M} , can be found, we precompute it as a look-up table using the following calibration procedure. We paste a thin diffuse white sheet on the mirror surface and project a color-coded dot pattern on it using a projector. We then record videos of all mirror deformations that would be used to capture flexible FOVs. We repeat each deformation a number of times. In all, we collected 30,753 frames, one of which is shown in Figure 3(a). In each frame, we detect the mirror boundary and compute its descriptor \mathcal{D} . We also identify the projected dots (see Figure 3(a)) and triangulate them with the projector to get 3D points on the deformed mirror. These reconstructed points are then used to determine the spline coefficients c_{ij} in Equation 1. The RMS percentage errors for these fits were less than 0.05% per frame. The \mathcal{D} and c_{ij} for each frame are used to construct the look-up table. Note that this is done off-line and only once.

4.3. Estimating Mirror Shape for a Captured Image

For a captured image, the problem of finding the shape of the mirror is reduced to finding the look-up table entry whose mirror boundary descriptor is closest (in L2) to that




Figure 4. (a) Image captured by our system where the mirror was flexed to increase horizontal FOV. (b) The corresponding estimated mirror shape (gray surface) and viewpoint locus (blue surface).

of the given image. The corresponding spline coefficients are taken to determine the mirror shape.

4.4. Evaluation

To validate our assumption that the boundary can be used to determine the mirror shape we performed the following evaluation. We removed 1000 frames at random from the calibration set to form a test set. We know the mirror shapes for these frames. Using the look-up table we then ‘estimated’ the mirror shapes for these frames and compared the z-coordinates of the actual and estimated shapes. The RMS percentage errors over all frames was 0.0926%. We also computed the RMS, maximum, and median percentage errors for each frame. The maximum values of these are shown in Figure 3(b) and can be seen to be small.

5. Properties of Captured Images

Once the mirror shape is known for a captured image, we can compute different properties of the system that enable us to understand how the scene was mapped onto the image.

5.1. Viewpoint Locus

It has been shown that only rotationally symmetric conic mirrors placed at specific locations with respect to a perspective camera yield a single viewpoint [1]. Since in our system the shape of the mirror is not necessarily a rotationally symmetric conic, it generally does not possess a single viewpoint, but rather a locus of viewpoints – a caustic surface. Caustics of non-single viewpoint systems have been studied earlier [6] and we employ the same approach to determine caustics for our systems³. Figure 4(a) shows an image captured by our system where the mirror was flexed into a convex cylindrical shape to increase the horizontal FOV. In Figure 4(b), the computed mirror shape (gray surface) and the corresponding caustic surface (blue surface) are shown.

5.2. Field of View

Since we know the shape of the mirror, for each pixel (x, y) in the image, we can trace back the captured ray and determine its effective viewing direction $V_{x,y}$. This viewing direction can be represented as a point on the surface of the unit sphere. In this way, we can map each pixel of a captured image onto the unit sphere. We compute the FOV (effective

solid angle) imaged by a pixel (x, y) as the area of the spherical quadrilateral subtended on the unit sphere by the viewing directions of pixels in a 2×2 neighborhood whose top-left pixel is (x, y) . The FOV of a pixel can be easily shown to be

$$\omega_{x,y} = \sum_{i=0}^1 \sum_{j=0}^1 A(V_{x+i,y+j}, V_{x+1-i,y+j}, V_{x+i,y+1-j}), \quad (2)$$

where, $A(a, b, c) = -\sin^{-1}\left(\frac{(a \times b)}{|(a \times b)|} \cdot \frac{(a \times c)}{|(a \times c)|}\right)$. By adding the FOVs of all camera pixels that see the mirror, we get the FOV of the imaging system:

$$\Omega = \sum_{y=1}^{R-1} \sum_{x=1}^{C-1} P_{x,y} \omega_{x,y}, \quad (3)$$

where, the image is of size $C \times R$ and $P_{x,y} = 1$, if pixel (x, y) sees the mirror and 0 otherwise. Figures 1(c), 7(d), 8(d,j), and 9(c) show the FOVs captured in the images shown in Figures 1(a), 7(b), 8(b,h), and 9(a), respectively. Note that though our camera has a horizontal FOV of 28.22° and a vertical FOV of 20.89° , by using the flexible mirror, we are able to realize a wide range of effective FOVs – the horizontal and vertical FOVs become as large as 120° and 40° , respectively. The unconventional shapes of these FOVs enable us to compose scenes in novel ways. Observe that in some cases the FOV even folds over (see Figures 8(j) and 9(c)).

5.3. Resolution

The resolution of a pixel in a captured image is defined [1] as $\rho = \frac{A}{\omega}$, where A is the area of the pixel and ω is the solid angle of the scene imaged by that pixel (given by Equation 2). Since all pixels occupy the same area on the image plane, the relative resolution of a pixel is given by $\rho = \frac{1}{\omega}$. Figure 1(d) shows the solid angle resolution map corresponding to the captured image in Figure 1(a). An image captured by our system almost always has spatially varying resolution. The resolution map allows us to visualize where the scene has been compressed or stretched in the image.

In addition to the solid angle resolution map, we can compute three additional resolution maps that represent (a) horizontal resolution, (b) vertical resolution, and (c) angular resolution. To compute these quantities for a pixel (x, y) , consider the spherical triangle formed by the viewing directions of pixels (x, y) , $(x + 1, y)$, and $(x, y + 1)$, i.e., by $V_{x,y}$, $V_{x+1,y}$, and $V_{x,y+1}$, respectively. *Horizontal resolution* $HR_{x,y}$ is the length of the side of the spherical triangle joining vertices $V_{x,y}$ and $V_{x+1,y}$. Analogously, *vertical resolution* $VR_{x,y}$ is the length of the side joining $V_{x,y}$ and $V_{x,y+1}$. *Angular resolution* $AR_{x,y}$ is the angle of the spherical triangle at $V_{x,y}$. Together these three maps better capture how scene elements are compressed or stretched in an image than does the solid angle resolution map alone. We will use these maps to undistort the captured images. Note that given the horizontal, vertical and angular resolution maps, we can directly compute the solid angle of a pixel (x, y) as

³Once the mirror shape is known, we can determine the rays that are captured by the system, which we then use to compute the caustic as in [6].

$$\omega_{x,y} = \sum_{i=0}^1 2 \tan^{-1} \left(\frac{l_i \sin(AR_{x+i,y+i})}{1 + l_i \cos(AR_{x+i,y+i})} \right), \quad (4)$$

where,

$$l_i = \tan\left(\frac{HR_{x+i,y}}{2}\right) \tan\left(\frac{VR_{x,y+i}}{2}\right). \quad (5)$$

6. Undistorting Captured Images

We propose to minimize the distortions in a captured image (without having to know scene geometry) by mapping a captured image to an equi-resolution image – an image where all pixels correspond to equal solid angles.

6.1. Creating Equi-Resolution Images

One approach to create an equi-resolution image is to map all pixels in a captured image onto the surface of a unit sphere (as discussed in Section 5.2) and then project the sphere onto a plane using an equi-areal projection such as Mollweide, Hammer, etc. [18]. These are 1-1 bijective mappings from a sphere onto a plane. However, this approach cannot be used if the FOV folds over (see Figures 8(j) and 9(c)) i.e., if multiple points on the image plane map to the same point on the sphere. In such cases, since the mapping between the sphere and the captured image is not 1-1, a 1-1 mapping between the equi-resolution image and the captured image does not exist.

Proposed Algorithm

We propose an alternative algorithm to compute equi-resolution images from captured images that can be used even if the FOV folds over. Our algorithm has four steps:

1. Compute the four resolution maps – solid angle resolution, horizontal resolution, vertical resolution and angular resolution (Section 5.3).
2. Determine the smallest solid angle in the captured image from the solid angle resolution map and use that as the target solid angle – the solid angle that we want all pixels in the equi-resolution image to have. This is denoted by t .
3. Stretch the captured image, denoted by \mathcal{I}_C , horizontally (retain pixel y -coordinates), keeping the center column fixed, so that all pixels have the same horizontal resolution, denoted by a . The resulting image is denoted by \mathcal{I}_H .
4. Stretch \mathcal{I}_H vertically (retain pixel x -coordinates), keeping the center row fixed so that the vertical resolution along the center column is a and the solid angle at all pixels is t . The resulting equi-resolution image is denoted by \mathcal{I}_E .

We now explain steps 3 and 4 in greater detail. For step 3, the target horizontal resolution a has to be specified. A good estimate of a can be automatically computed from the target solid angle t , as described in the Appendix. This is the approach we have used for all examples in this paper. Alternatively, the value of a can be specified by the user. Given a , we linearly sample each row of \mathcal{I}_C starting from the center column and moving outwards, such that the horizontal




Figure 5. Illustrations of steps 3 and 4 in the computation of equi-resolution images. See text for details.

resolution at each pixel is a . For this, we use \mathcal{I}_C 's horizontal resolution map. This process is illustrated in Figure 5(a). The grid represents the pixel grid of \mathcal{I}_C , while the gray circles represent the sample points for the horizontally stretched image \mathcal{I}_H . Since this image will be used as the input for the next step, we need its vertical and angular resolution maps, which can be computed from \mathcal{I}_C 's resolution maps.

In step 4, we construct the equi-resolution image \mathcal{I}_E , row by row, starting from the center. Consider the pixel grid of \mathcal{I}_H shown in Figure 5(b). The gray circles d_i represent pixels on the center row of \mathcal{I}_H as well as \mathcal{I}_E , d_0 being the center pixel. We construct the row above the center row beginning with the sample e_0 in the center column. Since a is the vertical resolution along the center column of \mathcal{I}_E , the vertical resolution at e_0 should be a . Hence, we can use the vertical resolution map of \mathcal{I}_H to determine e_0 's y -coordinate. Also, the area of the spherical quadrilateral corresponding to e_0, d_0, d_1 , and e_1 should be t – the target solid angle. Given e_0 , the vertical resolution at e_1 that would satisfy this is given by

$$v_1 = 2 \tan^{-1} \left(\frac{\sin\left(\frac{(t-m_1)}{2}\right)}{\sin\left((H-G) - \frac{(t-m_1)}{2}\right) \tan\left(\frac{n_1}{2}\right)} \right), \quad (6)$$

where m_1 is the area of the spherical triangle formed by e_0, d_0 , and d_1 , H is the angular resolution at d_1 , G is the spherical angle between the sides corresponding to $e_0 d_1$ and $d_0 d_1$ (see Figure 5(b)), and n_1 is the arc length on the sphere corresponding to $e_0 d_1$. The value of v_1 can be used to determine the y -coordinate of e_1 . Similarly, the location of e_1 can be used to determine the location of e_2 . The same technique can be applied to compute all e_i 's. Due to numerical issues, the e_i sequence might not be smooth. To ensure smoothness, we fit a cubic approximating spline to the e_i values and obtain smoother estimates \hat{e}_i . These smoothed \hat{e}_i 's form a row⁴ of \mathcal{I}_E . All other rows of \mathcal{I}_E can be constructed similarly.

As discussed in Section 5.2, in some configurations the FOV folds over (see Figures 8(j) and 9(c)). In such cases, the boundary of the fold (where the FOV starts folding over – a band 5-10 pixels wide) has very small solid angles – as much as 200 times smaller than the solid angles at other points in the image. Consequently, these regions project onto very small areas in the equi-resolution image, i.e., the image looks ‘pinched’ in these regions. Therefore, we consider such regions as anomalies and substitute the horizontal, vertical and angular resolutions at such points by interpolating

⁴Note that to construct additional rows above this row, we need to determine the horizontal and angular resolution in \mathcal{I}_E at all \hat{e}_i . These can be computed from \mathcal{I}_H 's resolution maps.




Figure 6. (a,c) Renderings of a scene with black spheres reflected in a convex cylindrical mirror (a) and a concave cylindrical mirror (c). In both cases the curvature is along the horizontal. (b,d) Solid angle resolution maps for the images in (a,c). (e,g) Equi-resolution images computed from (a,c). (f,h) Solid angle resolution maps of the equi-resolution images in (e,g). (i,k) Rectangular images computed from the equi-resolution images in (e,g). (j,l) Solid angle resolution maps of the rectangular images in (i,k).

the resolutions at neighboring points.

Evaluation

To evaluate the ability of our technique to minimize distortions we performed the following simulation. We constructed a synthetic scene with a large number of black spheres arranged on a tessellated sphere obtained by recursively subdividing an icosahedron. We then rendered the scene as reflected in a convex cylindrical mirror and a concave cylindrical mirror⁵. The rendered images are respectively shown in Figures 6(a,c) in which the spheres are severely distorted and unevenly spaced. The solid angle resolution maps for these images are shown in Figures 6(b,d). As one can see, the resolution varies over the entire image. Figures 6(e,g) show the equi-resolution images computed from the rendered images. Note that the spheres are almost circular and evenly spaced, demonstrating that our mapping minimizes local distortions. Their corresponding solid angle resolution maps are shown in Figures 6(f,h). The RMS percentage errors are 0.04% and 0.03%, respectively.

6.2. Creating Rectangular Images

Equi-resolution images do not have rectangular boundaries and look odd since we are used to seeing rectangular images. We can correct this by applying an image warp to the equi-resolution image that maps its boundary to the boundary of a rectangle. In most cases, this additional warping does not introduce much resolution variation. To compute this warp, we need to establish correspondence between the boundaries of the equi-resolution image and the output rectangular image. This is done as follows. We discretize the ‘horizontal’ and ‘vertical’ sides of the boundary of the equi-resolution image so that they have p and q equidistant points, respectively⁶. The longer of the two hor-

⁵The mirrors were rotated about the horizontal axis to mimic the configuration in our physical setup.

⁶In our experiments, we used $p = 32$ and $q = 24$.

izontal sides and the two vertical sides determine the width and height of the rectangular image, respectively. Like the boundary of the equi-resolution image, the horizontal and vertical sides of the rectangle are also discretized to have p and q equidistant points, respectively, thereby establishing correspondences between the two boundaries. These correspondences are used to setup a thin-plate spline based image warp [11] that maps points (m, n) in the rectangular image to points (r, s) in the equi-resolution image. This warp consists of two maps $\mathcal{F}_r : (m, n) \rightarrow r$ and $\mathcal{F}_s : (m, n) \rightarrow s$, both of which minimize the energy

$$\int \int_{\sigma} [(\frac{\partial^2 \mathcal{F}}{\partial^2 m})^2 + 2(\frac{\partial^2 \mathcal{F}}{\partial m \partial n})^2 + (\frac{\partial^2 \mathcal{F}}{\partial^2 n})^2] dmdn, \quad (7)$$

where σ is the domain of the rectangular image. Figures 6(i,k) show the rectangular images computed from the equi-resolution images in Figures 6(e,g). Their corresponding solid angle resolution maps are shown in Figures 6(j,l). Note that the rectangular images have almost uniform resolution and low distortion.

7. Examples

We now present a number of scene compositions captured by our flexible FOV imaging system.

Birthday Snap: Often, when taking a picture of a group of people, like at a birthday (see Figure 7(a)), not everyone fits in the picture – we need a larger horizontal FOV. The image in Figure 7(a) was captured by our system when the mirror was not flexed. Figure 7(b) shows the picture captured by flexing the mirror’s two vertical edges and deforming it into a convex cylindrical shape. All 7 people now fit in the FOV. The estimated mirror shape and the resulting FOV are shown in Figures 7(c) and 7(d), respectively. Note that the horizontal FOV (121.72°) is much larger than the vertical FOV (20.25°) – the effective aspect ratio is very different from that of the image detector. At the same time, all the pixels are devoted to capturing only the desired scene elements. This is in contrast to what can be achieved if we use a wide-angle lens to capture all 7 people in an image. The spatially varying solid angle resolution of the image in Figure 7(b) is shown in Figure 7(e). The equi-resolution image computed from the captured image in Figure 7(b) is shown in Figure 7(f) and its solid angle resolution map is shown in Figure 7(h). Figure 7(g) shows the rectangular image computed from the image in Figure 7(f). Its resolution map is shown in Figure 7(i).

Panning Up: Figure 8(a) shows an image captured by our system when the mirror was not flexed. One can see only the feet of a person. Figure 8(b) shows the image captured on flexing the top edge of the mirror to include the complete person. The estimated mirror shape is shown in Figure 8(c), while Figure 8(d) shows the captured FOV. Figures 8(e) and 8(f) show the computed equi-resolution image and rectangular image, respectively.

Street Monitoring: Figure 8(g) shows a street captured

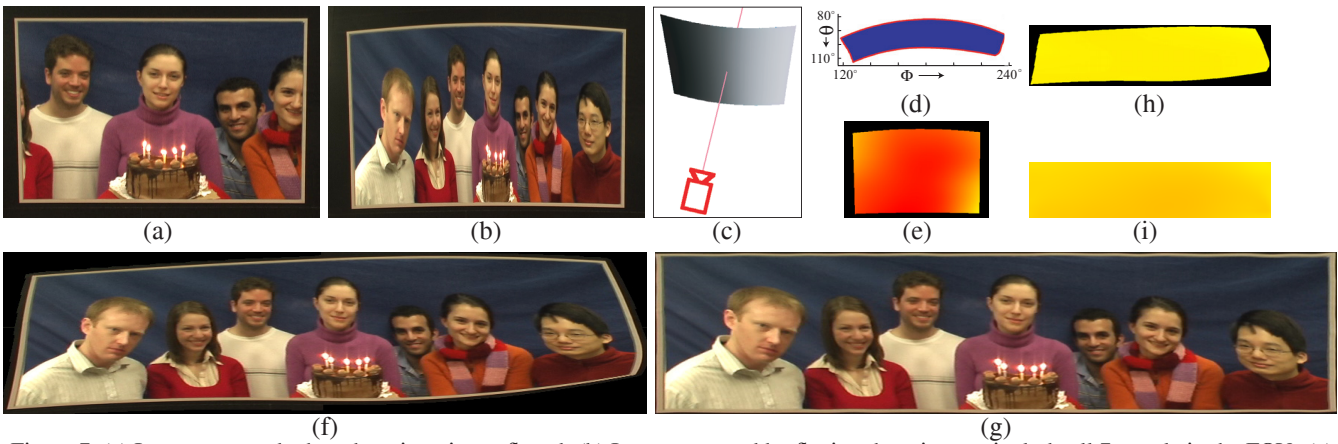


Figure 7. (a) Image captured when the mirror is not flexed. (b) Image captured by flexing the mirror to include all 7 people in the FOV. (c) Estimated mirror shape. (d) FOV and (e) solid angle resolution map for the image in (b). (f, h) Equi-resolution image computed from (b) and its resolution map. (g, i) Rectangular image computed from (f) and its resolution map. Please see the accompanying video.

when the mirror is not flexed. The ability to change the FOV enables our system to see the buildings on the sides of the street and the sky (top edge flexed) in Figure 8(h) and the left side of the street (left edge flexed) in Figure 9(a). Figures 8(i) and 9(b) show the corresponding estimated mirror shapes. The captured FOVs are shown in Figures 8(j) and 9(c). Figure 8(k) shows the equi-resolution image computed from the image in 8(h). The undistorted rectangular images computed from the captured images are shown in Figures 8(l) and 9(d).

These results demonstrate that the proposed imaging system enables us to compose scenes in ways not possible before – to include only the scene elements of interest and exclude all others. Hence, image pixels are devoted only to the desired scene elements. This not only opens up a new creative dimension in photography, but can also be useful for vision applications.

8. Limitations and Future Work

Though the proposed approach enables us to compose scenes in novel ways, it has certain limitations. Different fields of view are obtained by deforming the mirror to get various curved surfaces. Focusing on the reflections of a scene in a curved mirror is known to be difficult due to optical aberrations introduced by the curvature of the mirror. In our implementation, these effects are not noticeable as we have used a large mirror – for any given shape, a larger mirror has lower local curvatures and hence produces less blurring due to aberrations. We are currently exploring ways to make our system smaller by using a post-processing stage to reduce the blurring introduced by a smaller mirror. Since we know the shape of the mirror, we can determine its spatially varying point spread function and use it to deblur the captured image. Although this does not guarantee a perfect image, we believe it can significantly enhance the quality of the captured image. Finally, in its current form, our system can only be manually controlled to change the FOV. We are

exploring the use of a small number of servo-controlled actuators to enable computer-controlled flexing of the mirror. Such a capability would be of great value to surveillance and monitoring systems – they can ‘react’ to changes in the scene and keep only the scene elements of interest in view.

Appendix

In course of creating the equi-resolution image, to compute the target horizontal resolution a from the target solid angle t , we make the following assumptions. Since our algorithm maps the center row and column of the captured image \mathcal{I}_C onto the center row and column of the equi-resolution image \mathcal{I}_E , we assume that in \mathcal{I}_E , the horizontal resolution along the center row and the vertical resolution along the center column are equal (to a). We also assume that at the center of \mathcal{I}_E , the angular resolution is $\pi/2$. This implies that at the center of \mathcal{I}_E , we have a spherical triangle with area $t/2$, that has two equal (to a) sides and the included angle is $\pi/2$. Under these assumptions, a can be shown to be

$$a = 2 \tan^{-1}(\sqrt{\tan(t/4)}). \quad (8)$$

Acknowledgments

This work was conducted at the Computer Vision Laboratory at Columbia University. It was supported by a grant from the Office of Naval Research (No. N00014-1-06-1-0032).

References

- [1] S. Baker and S. K. Nayar. A Theory of Single-Viewpoint Catadioptric Image Formation. *IJCV*, pages 175–196, 1999.
- [2] J. Chahl and M. Srinivasan. Reflective surfaces for panoramic imaging. *Applied Optics*, pages 8275–8285, 1997.
- [3] T. Conroy and J. Moore. Resolution invariant surfaces for panoramic vision systems. *ICCV*, pages 392–397, 1999.
- [4] P. Decaudin, D. Julius, J. Wither, L. Boissieux, A. Sheffer, and M.-P. Cani. Virtual Garments: A Fully Geometric Approach for Clothing Design. *Computer Graphics Forum*, pages 625–634, 2006.
- [5] C. Geyer and K. Daniilidis. Catadioptric Camera Calibration. *ICCV*, pages 398–404, 1999.

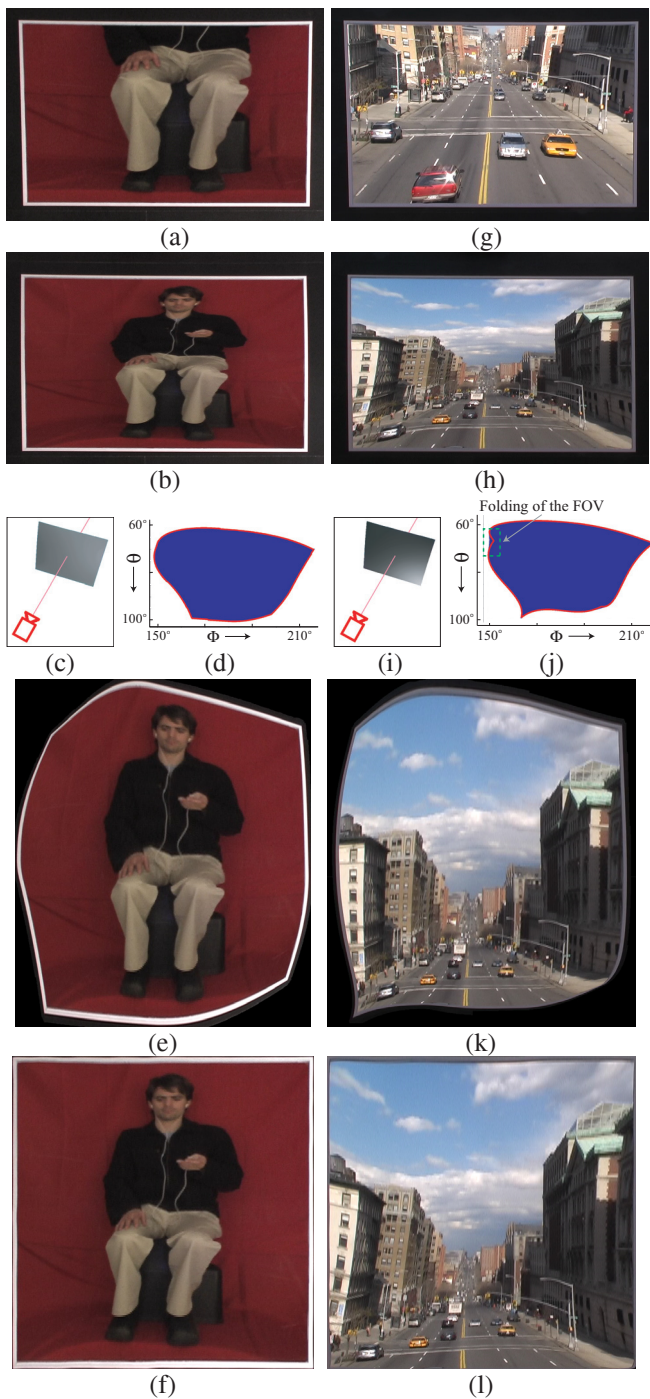


Figure 8. (a,g) Captured images when the mirror is not flexed. (b,h) Captured images when the mirror's top edge is flexed. (c,i) Estimated mirror shapes for the images in (b,h). (d,j) FOVs captured in images in (b,h). (e,k) Equi-resolution images computed from images in (b,h). (f,l) Rectangular images computed from images in (e,k). Please see the accompanying video.

- [6] M. Grossberg and S. K. Nayar. The Raxel Imaging Model and Ray-Based Calibration. *IJCV*, pages 119–137, 2005.
- [7] M. Halstead, B. Barsky, S. Klein, and R. Mandell. Reconstructing curved surfaces from specular reflection patterns us-

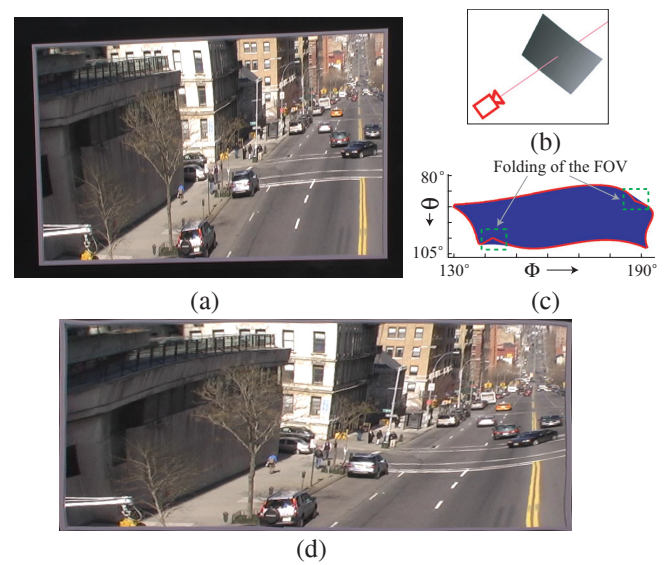


Figure 9. (a) Image captured when the mirror was flexed to include the left side of the street. (b) Shape of the mirror estimated for the image in (a). (c) FOV captured in (a). (d) Undistorted rectangular image computed from (a). Please see the accompanying video.

ing spline surface fitting of normals. *SIGGRAPH*, pages 335–342, 1996.

- [8] R. A. Hicks, V. T. Nasis, and T. P. Kurtzweg. Programmable imaging with two-axis micromirrors. *Optics Letters*, pages 1066–1068, 2007.
- [9] R. A. Hicks and R. K. Perline. Equireolution catadioptric sensors. *Applied Optics*, pages 6108–6114, 2005.
- [10] S. B. Kang. Catadioptric self-calibration. *CVPR*, pages 201–207, 2000.
- [11] S. Lee, K. Chwa, J. Hahn, and S. Shin. Image morphing using deformable surfaces. *Computer Animation*, pages 31–39, 1994.
- [12] B. Micusik and T. Pajdla. Autocalibration and 3D Reconstruction with Non-central Catadioptric Cameras. *CVPR*, 2004.
- [13] S. K. Nayar, V. Branzoi, and T. E. Boult. Programmable Imaging: Towards a Flexible Camera. *IJCV*, pages 7–22, 2006.
- [14] Y. Nomura, L. Zhang, and S. K. Nayar. Scene Collages and Flexible Camera Arrays. *EGSR*, pages 127–138, 2007.
- [15] F. Roddier, editor. *Adaptive Optics in Astronomy*. Cambridge University Press, 1999.
- [16] L. Saurei, G. Mathieu, and B. Berge. Design of an autofocus lens for VGA 1/4" CCD and CMOS Sensors. *SPIE*, 5249:288–296, 2004.
- [17] S. Savarese, M. Chen, and P. Perona. Local Shape from Mirror Reflections. *ICCV*, 64:31–67, 2005.
- [18] J. P. Snyder. Map Projections—A Working Manual. *U. S. Geological Survey Professional Paper 1395*, 1997.
- [19] R. Swaminathan, M. Grossberg, and S. K. Nayar. A Perspective on Distortions. *CVPR*, pages 594–601, 2003.
- [20] R. Swaminathan, M. Grossberg, and S. K. Nayar. Designing Mirrors for Catadioptric Systems that Minimize Image Errors. *Workshop on Omnidirectional Vision*, 2004.
- [21] J. Yu and L. McMillan. Modelling Reflections via Multiperspective Imaging. *CVPR*, pages 117–124, 2005.