

Cata-Fisheye Camera for Panoramic Imaging

Gurunandan Krishnan Shree K. Nayar
Department of Computer Science
Columbia University
{gkguru, nayar}@cs.columbia.edu

Abstract

We present a novel panoramic imaging system which uses a curved mirror as a simple optical attachment to a fisheye lens. When compared to existing panoramic cameras, our “cata-fisheye” camera has a simple, compact and inexpensive design, and yet yields high optical performance. It captures the desired panoramic field of view in two parts. The upper part is obtained directly by the fisheye lens and the lower part after reflection by the curved mirror. These two parts of the field of view have a small overlap that is used to stitch them into a single seamless panorama. The cata-fisheye concept allows us to design cameras with a wide range of fields of view by simply varying the parameters and position of the curved mirror. We provide an automatic method for the one-time calibration needed to stitch the two parts of the panoramic field of view. We have done a complete performance evaluation of our concept with respect to (i) the optical quality of the captured images, (ii) the working range of the camera over which the parallax is negligible, and (iii) the spatial resolution of the computed panorama. Finally, we have built a prototype cata-fisheye video camera with a spherical mirror that can capture high resolution panoramic images (3600x550 pixels) with a 360° (horizontal) \times 55° (vertical) field of view.

1. Introduction

Wide-angle and panoramic imaging have had significant impact on a variety of real-world applications. For example, being able to “see” in all directions provides situational awareness in surveillance and autonomous navigation tasks. Image-based immersive virtual reality [1, 2, 11, 13] enables realistic exploration of indoor and outdoor environments. Panoramic video conference [3] and telepresence devices aid in effective virtual collaboration. All these applications rely on capturing high-quality panoramic videos. In many of these applications, it is not required to capture the entire spherical field of view (FOV). It is sufficient to capture 360° panoramas with a reasonable vertical field of view that has more or less equal coverage above and beneath the horizon. For instance, in the case of video conferencing, we only need to capture the sitting participants and the table. The same is true for navigation, where the vehicle only needs to view the roads and the obstacles around it. In this paper, our focus is on capturing such panoramic fields of view.

Current techniques for capturing panoramic videos can be classified as either dioptric (systems that use only refractive optics) or catadioptric (systems that use both reflective and refractive optics). Dioptric systems include camera-clusters and wide-angle or fisheye lens based systems. Catadioptric systems include ones that use single or multiple cameras with one or more reflective surfaces.

Dioptric camera-clusters [14, 20, 28] compose high resolution panoramas using images captured simultaneously by multiple cameras looking in different directions. In this approach, due to their finite size, the cameras cannot be co-located. Hence, they do not share a single viewpoint and, in general, have significant parallax. To address this issue, Nalwa [22] uses a pyramidal (piecewise planar) mirror to co-locate the viewpoints of multiple cameras, thus eliminating the parallax (also see [17]). In all cases, the use of multiple cameras requires careful geometric and photometric calibration of the cameras as well as synchronization with respect to time. Moreover, the use of multiple cameras can make the imaging system bulky and expensive.

A popular dioptric solution is the use of fisheye lenses [21]. Since they have a hemispherical FOV, fisheye lenses can only be used to capture hemispherical panoramas [31]. Slater [26] proposed the use of two back-to-back fisheye lens cameras to obtain full spherical panoramas. Similar to camera-clusters, this approach suffers from large parallax. Also, if we need to capture a reasonable vertical coverage about the horizon, say $+30^\circ$ to -30° , then in terms of optics, image detectors and bandwidth, it is wasteful to capture the entire sphere.

Catadioptric systems [8, 9, 10, 12, 16, 25, 30] employ one or more reflective surfaces capable of projecting a wide view of the world onto a single sensor. Several of these designs [18, 22, 23, 24] satisfy the so-called single viewpoint (SVP) constraint. Baker and Nayar [4] derived the family of mirrors that satisfy the SVP constraint, when used with perspective lenses. Since in these designs perspective cameras view the world through highly curved reflective surfaces, the vertical resolution (number of pixels per unit elevation angle) of the computed panoramas is not uniform and is, in general, poor near the apex of the mirror [23].

Yagi and Kawato [29] and Lin and Bajcsy [18] have used conical mirrors in their catadioptric designs, with the latter satisfying the SVP constraint. In this case, we obtain the

best vertical resolution. The conical mirror acts as a planar mirror in the radial direction and hence the vertical resolution is the same as that of the camera. However, this restricts the maximum vertical FOV of the panorama to half the FOV of the camera. For instance, if we wish to capture a panorama with a 60 degree vertical field of view, the camera must have a 120 degree field of view. Even if such a system is used, the resolution of the panorama approaches zero at the apex of the cone. More importantly, conical mirrors are known to produce strong optical aberrations due to astigmatism that fundamentally limit the resolution of the captured panorama [27].

In this work, we present a novel panoramic imaging system that uses a curved mirror as a simple and compact attachment to a fisheye lens. We refer to this as a cata-fisheye system. Fisheye lenses usually have a hemispherical FOV. In our system (see Figure 1), the curved mirror is placed within the central part of the fisheye FOV. This enables us to capture light rays beyond the hemispherical FOV of the fish-eye lens. A similar system, proposed by Davis *et al.* [15], uses a flat mirror to capture panoramic images with vertical FOV below the horizon (less than 0° and greater than -45°). Unlike their design, in our system the desired panorama is captured in two parts - the upper part is captured directly through the fisheye lens, while the lower part is obtained through the mirror. These two parts are seamlessly stitched to produce the final panorama.

Our approach to panoramic imaging can be viewed as a hybrid of the dioptric and catadioptric approaches. As we shall show, it overcomes many of the limitations of previous systems. Although our design is in part catadioptric, it uses a very shallow mirror that produces minimal optical aberrations. Thus, we are able to capture high quality panoramas. Our system captures its two parts of the panoramic field of view from two compact viewpoints. In this respect, it too has an inherent parallax. However, the two viewpoints are very close together and hence the parallax is negligible except for very small distances from the system.

We have developed a prototype camera that captures high resolution panoramas with a $360^\circ \times 55^\circ$ field of view. Although our prototype has almost equal coverage above and below the horizon, cata-fisheye systems can be designed with a wide range of vertical fields of view.

2. Cata-Fisheye Image Formation

Figure 1 illustrates the design of a radially symmetric cata-fisheye system. It uses a mirror in conjunction with a fisheye lens in order to capture light rays from beyond the FOV of the fisheye lens. If the overall vertical field of view of the system extends from θ_{au} to θ_{al} , then the upper part of the FOV (from θ_{fu} to θ_{fl}) is captured directly by the fish-eye lens and the remainder (from θ_{mu} to θ_{ml}) through the mirror. A small overlap in the two views allows them to be seamlessly blended into the final panorama.

The properties of the cata-fisheye system shown in Figure 1, can be concisely expressed as:




Figure 1. The cata-fisheye system “looks” at the world partly (from θ_{fu} to θ_{fl}) through the fisheye lens and partly (from θ_{mu} to θ_{ml}) through the convex mirror. $\Delta\theta_o$ is a small overlap between the two fields of view. All angles corresponding to the incident rays are measured from the vertical optical axis.




Figure 2. (a) Illustration of image formed by a rotationally symmetric cata-fisheye system. (b) The final panorama is generated by stitching the mirror-panorama and the fisheye-panorama.

$$\theta_{fu} = \theta_{au}, \quad (1a)$$

$$\theta_{ml} \geq \theta_{al}, \quad (1b)$$

$$\Delta\theta_o > 0. \quad (1c)$$

where, $\Delta\theta_o = (\theta_{fl} - \theta_{mu})$ is the overlap between the two FOVs. Also, $\theta_{fl} = \Omega_f/2$, where Ω_f is the FOV of the fish-eye lens¹. The first two conditions, (1a) and (1b), state that the mirror and the lens should not obstruct the panoramic FOV and the condition (1c) ensures that no part of the desired panoramic FOV is missed. All angles are measured from the vertical optical axis of the fisheye lens.

Figure 2a illustrates the image as formed by a cata-fisheye system. Unlike the image obtained by a regular fish-eye lens, the cata-fisheye image is not continuous with respect to the FOV. The outer orange annulus represents the upper FOV imaged directly through the fisheye lens. The inner blue annulus represents the FOV captured through the mirror. The small overlap is depicted as thin gray rings at the outer boundaries of the two views. The final panorama is produced by stitching the two individual panoramas corresponding to the FOVs of the mirror (mirror-panorama) and the fisheye lens (fisheye-panorama). We discuss the generation of panoramas in detail in Sections 4 and 5.

¹The FOV of a fisheye lens is usually specified as twice its vertical coverage, *i.e.* 180° fisheye lens covers 90° in the vertical direction.




Figure 3. A cata-fisheye camera design with a spherical mirror with radius of curvature r_m and width d_m is placed at a height of h_m from the fisheye lens viewpoint O . The point P lies on the rim of the mirror and has a surface normal \hat{n} .

3. Designing a Cata-Fisheye Camera

The cata-fisheye concept allows us to design cameras with a wide range of FOVs by simply varying the parameters (shape and position) of the curved mirror and the FOV of the fisheye lens. In this section, we discuss determining the parameters of the mirror for a desired panoramic FOV.

The mirror used in a cata-fisheye system can have a wide range of shapes - spherical, paraboloidal, hyperboloidal, conical, or even a complex aspherical shape. To minimize optical aberrations (maximize image resolution) we have chosen to use spherical mirrors. Further, to keep the system rotationally symmetric, the mirror is assumed to be coaxial with the fisheye lens.

Another design parameter that needs to be set is the amount of overlap between the upper and lower views. Since the stitching of the two views takes place in the image, it is convenient to specify the overlap in terms of the number of pixels, p_o . The corresponding angular overlap in the two FOVs is given by,

$$\Delta\theta_o = \Omega_f/2 - f^{-1}[f(\Omega_f/2) - p_o]. \quad (2)$$

Here, $r = f(\theta)$ is the projection function of the fisheye lens, where r is the radial distance of a pixel from the image center and θ is the elevation angle of the incident ray with respect to the optical axis of the fisheye lens.

Figure 3 illustrates the two-dimensional cross-section of a typical cata-fisheye system with a spherical mirror. The fisheye lens is assumed to have a single viewpoint with its position roughly known to be at O .² h_l is the distance between O and the tip of the lens and d_l is the width of the lens (including its casing). The spherical mirror, whose radius of curvature r_m and width d_m are to be determined, is assumed to be positioned at a distance h_m from the lens.

Let P be a point on the rim of the mirror. According to conditions in Equation (1), the position of P and the surface normal \hat{n} at P (and hence the radius and center of curvature of the mirror) should be such that an incoming ray at angle θ_{mu} is reflected by P towards O (see Figure 3). To satisfy this condition, the directional angle of the normal at P should be:

²Since fisheye lenses do not exactly have a single viewpoint, its assumed position may compromise the design and therefore the field of view of the imaging system. One way to solve this problem is to add tolerances to the limits of the desired field of view.

Panorama FOV (Required)	Fisheye FOV (Chosen)	Design Solution (in mm)		
		r_m	d_m	h_m
+30° to -30° (60° to 120°)	170°	70.00	47.88	4.6
	180°	76.36	44.45	4.5
	190°	81.49	42.18	4.4
	200°	96.22	39.73	4.2
	210°	106.98	37.15	4.1

Figure 4. Given the projection function and field of view of the fisheye lens and the desired panoramic field of view, the radius of curvature, width and position of the spherical mirror are determined using the design equations in Section 3. This table shows a few examples of cata-fisheye camera designs. In all cases, $\Delta\theta_o = 5^\circ$, $h_l = 5\text{mm}$ and $d_l = 55\text{mm}$.

$$\theta_{\hat{n}} = \frac{1}{2}(\pi + \theta_{fu} + \theta_{mu}). \quad (3)$$

From all of the above constraints, the radius of curvature and the width of the mirror can be easily determined to be:

$$r_m = \frac{(h_m + h_l) \sin \theta_{fu}}{\sin(\theta_{fu} - \theta_{\hat{n}}) - \sin \theta_{fu}}, \quad (4)$$

$$d_m = \frac{2(h_m + h_l) \sin \theta_{\hat{n}} \sin \theta_{fu}}{\sin(\theta_{fu} - \theta_{\hat{n}}) - \sin \theta_{fu}}. \quad (5)$$

Note that the radius and the width are computed as a function of the position h_m of the mirror. While any arbitrary h_m and a mirror with the corresponding r_m and d_m guarantee the upper limit (condition (1a)) and the required overlap (condition (1c)), they need not ensure that the lens will not obstruct the panoramic field of view (condition (1b)). Therefore, we perform a linear search to find the smallest h_m for which $\theta_{ml} \geq \theta_{al}$.

Note that there can be several design solutions for a desired panoramic field of view, based on the FOV of the fisheye lens and the overlap between the upper and lower views. The table in Figure 4 shows a few examples of cata-fisheye camera designs. In practice, we select the best optical design by evaluating all the possible designs with respect to several performance measures, including the resolution of the panorama, the camera's working range with respect to parallax, and the optical quality of the image. We discuss these performance measures in detail in Section 7.

4. Calibration of Cata-fisheye Camera

In order to geometrically map the captured image to the upper and lower panoramas and then stitch them, the relationship between the light rays entering the camera and the pixels in the captured image has to be known. This is completely governed by the design parameters. Due to manufacturing tolerances, however, the actual position and orientation of the mirror need not adhere strictly to the design specifications. This, compounded with the possible inaccuracy in the assumed fisheye lens viewpoint, requires us to calibrate the system to determine (i) the fisheye projection function $r = f(\theta)$, and (ii) the actual center of curvature (x_m, y_m, z_m) and orientation (α_m, β_m) of the mirror. To calibrate the fisheye lens, either a projection function provided by the manufacturer can be used or any of the existing

calibration methods may be used [5, 6, 7]. Here, we focus on the second part of the calibration.

We pose the calibration as a process of searching for a solution $\tilde{M} = (x_m, y_m, z_m, \alpha_m, \beta_m)$ in an R^5 space (M) for which corresponding local features in the overlapping FOV of the two views are aligned in the panoramic space. The camera is first placed in an environment with a large number of distant features, particularly in the overlapping region. Nearby features should be avoided as they would have larger parallax and hence tend to introduce errors in the alignment process. With the solution initialized to the design values, $\tilde{M} = (0, 0, h_m + h_l, 0, 0)$, we generate the mirror-panorama and the fisheye-panorama for the overlapping region. Using SIFT[19], we then extract a corresponding set of features from the two panoramas. The extracted features are represented in spherical coordinates. Any spurious matches resulting from this procedure can be easily sieved out by using the constraint that matched features should have very similar spherical coordinates. In the pruned feature set, let the i th feature in the mirror-panorama be denoted by $F_m^{(i)} = [\theta_m^{(i)}, \phi_m^{(i)}]$ and the corresponding feature in the fisheye-panorama be denoted by $F_f^{(i)} = [\theta_f^{(i)}, \phi_f^{(i)}]$. Now the problem of finding the solution \tilde{M} can be setup as a minimization of an objective function that is the sum of the squares of the errors in the alignment of the corresponding features:

$$\arg\min_{\tilde{M}} \sum [F_f^{(i)} - F_m^{(i)}(\tilde{M})]^2. \quad (6)$$

Note that, during this optimization, it is not required to search for the solution in the entire R^5 space. We only need to search in a small interval ΔM centered at the ideal design parameter values. The size of the interval is chosen based on the manufacturing tolerances and the ambiguity in fisheye lens viewpoint.

5. Computing the Panorama

Once the system is calibrated, generating the panorama with the required resolution is straightforward. Each pixel in the panorama represents a spherical coordinate/direction, and for each pixel we need to find the corresponding location in the captured image. We consider three possibilities here. A panorama pixel (θ_p, ϕ_p) may lie in (i) the direct FOV of the fisheye lens, (ii) the FOV of the mirror and (iii) both the fisheye lens' and the mirror's FOVs. If (θ_p, ϕ_p) is in the fisheye lens' FOV, then its corresponding location in the captured image is given by $(f(\theta_p)\cos\phi_p, f(\theta_p)\sin\phi_p)$ where $f(\theta_p)$ is the fisheye projection function.

Alternatively, if the pixel is in the mirror's FOV, then its corresponding location in the captured image can be determined by solving for the outgoing ray (θ_c, ϕ_c) from the lens which would finally proceed in the pixel's direction (θ_p, ϕ_p) after being reflected by the mirror. This is a non-linear problem and can be solved as follows. Let $R(\theta, \phi)$ be the direction in which a ray (θ, ϕ) from the origin O would proceed after being reflected by the mirror. Then, the following optimization solves for the required (θ_c, ϕ_c) in a least squares sense:

$$\arg\min_{(\theta_c, \phi_c)} [(\theta_p, \phi_p) - R(\theta_c, \phi_c)]^2. \quad (7)$$

Subsequently, the required image location is given by $(f(\theta_c)\cos\phi_c, f(\theta_c)\sin\phi_c)$. Observe that, solving the above problem for each pixel in the mirror's FOV can be computationally expensive. Fortunately, we are not required to do so. Since the reflecting point on the spherical mirror varies smoothly with the incoming ray direction, we only need to solve for pixels at discrete intervals and interpolate the solution for in-between pixels.

Finally, if the pixel lies in the overlapping region, we employ linear feathering to smoothen the blending. The entire mapping between the panorama and the captured image, can be stored as a simple lookup-table. Using the lookup table, we can map the captured images to high resolution panoramas in real time. *Note that the calibration process and the mapping calculations need to be performed just once for any given cata-fisheye camera.*

6. Camera Prototype

Figure 5b shows our prototype cata-fisheye camera that has a vertical FOV of 55° ($+25^\circ$ to -30°). It was built using off-the-shelf components – a 7 frames per second Lumina Lw570 digital video camera with a 5 megapixel sensor³ and a Fujinon FE185C046HA-1 fisheye lens with a FOV of 185° ⁴. Usually, an overlap of 20 to 30 pixels is sufficient to perform blending. However, in our prototype, the image formed by the fisheye lens is much larger than the sensor (see Figure 5c), resulting in the image being chopped off at the top and bottom. Since the missing parts of the FOV are near the equator, they could appear as holes in the computed panorama. In order to avoid these holes, we set the overlap to be larger than the holes (20°). Using Equations 4 and 5, the radius and width of the mirror were determined to be 70.12mm and 50.26mm, respectively, and its position to be 2.6mm above the tip of the lens.

With respect to packaging, our camera design is extremely compact. The size of the complete system is just 10cmx7cmx5.5cm (see Figure 5b). Moreover, the packaging is such that the mirror can be easily detached from the lens, thereby converting the cata-fisheye camera into a regular hemispherical camera.

Figures 10 and 11 show two examples of panoramas captured using our cata-fisheye camera. The first example corresponds to an outdoor campus scene. Figure 10a shows an image formed on the detector. Figure 10b shows the corresponding SIFT features extracted from the mirror and fisheye panoramas using our automatic calibration algorithm. Figure 10c shows the alignment of the features after the calibration process. The average error in the alignment was found to be 0.087° . The computed panorama, which has a resolution of 3600x550 pixels, is shown in two rows (Figures 10d and 10e). The second example, shown in Figure 11, corresponds to a typical video conference scenario. Here, in

³Since we purchased this camera, other manufacturers have developed cameras using the same detector that can output 15-30 fps.

⁴We found the FOV to be actually 194° and used this value instead.




Figure 5. (a) A hemispherical panoramic imaging system with just a fisheye lens. (b) Our prototype cata-fisheye panoramic camera. The inset shows the compact mirror attachment. (c) Image captured by a fisheye lens. (d) Image captured by the prototype.

addition to the captured panorama, we also show two perspective views computed from it that convey the high quality of captured information.

7. Performance Evaluation

In this section we describe some of the performance measures used to evaluate our prototype.

Optical Quality: The point spread functions (PSFs – also known as spot diagram) of a camera are an indication of the quality of the captured image – the more compact the PSFs, the better the image quality. We obtain the PSFs for our prototype camera (operated at F/4 and focused at infinity) using the Zemax optical simulator. The PSFs are shown in Figure 6. The individual spots correspond to different angles of incidence for three different wavelengths (red, green and blue). Observe that, the PSFs corresponding to the fisheye lens' FOV are more compact than those corresponding to the mirror's FOV. This blurring is due to optical aberrations like coma, astigmatism and field curvature induced by the curved mirror. Even so, the aberrations induced by the mirror are much smaller than those found in most catadioptric cameras. This is because our system only requires the use of a very

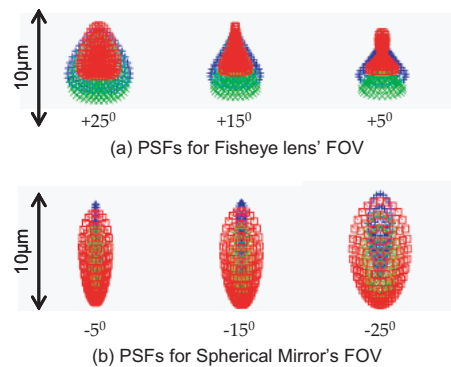


Figure 6. Point spread functions for various angles of incidence (indicated below the diagrams) for three wavelengths (red, green and blue) when the camera is operated at F/4.

F/#	2.0	2.8	4.0	5.6	8.0
Pixel Size (in microns)	3.85	2.7	2.5	3.3	4.9
Detector Res. (in pixels)	1227 x 1227	1749 x 1749	1890 x 1890	1432 x 1432	964 x 964

Figure 7. When operating at different F-numbers, the second row shows the minimum size of the pixel required to ensquare atleast 30% of the PSF energy. Based on the pixel size, the third row shows the corresponding resolution of detector that can be used.

shallow (low curvature) mirror.

When using a CCD or CMOS detector with a color mosaic (Bayer mosaic, for instance), a design is considered to be optically good if a detector pixel can ensquare at least 30% of the PSF energy [27]. Figure 7 shows a table containing the size of the pixel (second row) ensquaring 30% of the PSF energy, averaged over all incident angles for various F/#s. For instance, at F/4, the average PSF energy over all incident angles, is concentrated within a pixel of size 2.5 microns. Based on this pixel size and the size of the image formed by the fisheye lens, we can determine the detector resolution that would provide the highest quality images. In our prototype, we used an off-the-shelf detector with a pixel size of 2.2 microns, as it was closest in resolution and pixel size to the optimal resolution and pixel size (third column of Figure 7).

Parallax and Working Range: Since our system captures its two parts of the panoramic field of view from two different, albeit compact, viewpoints, it has an inherent parallax. In the presence of parallax, stitched panoramas often have ghosting artifacts. However, parallax is depth dependent, and becomes negligible after a certain scene depth (the minimum working distance). In other words, the camera can produce seamless panoramas without artifacts, only if the scene objects are farther than the minimum working distance. The range over which parallax is negligible is referred to as the working range of the camera.

Usually, parallax is defined in number of pixels and

Detector Res. (in pixels)	512x 512	1024x 1024	2048x 2048	2148x 2148
Working Range (in m)	0.65 – ∞	1.3 – ∞	2.56 – ∞	2.7 – ∞

Figure 8. The working range of the prototype camera for different detector resolutions, assuming the maximum acceptable parallax to be 2 pixels. The last column shows the working range of our current configuration.

hence varies with the resolution of the detector. For example, if we set the maximum acceptable error to be 2 pixels, then Figure 8 shows the working range of our prototype camera with different sensor resolutions. The working range of our prototype is 2.7m-∞ (fourth column in Figure 8).

Panorama Resolution: The cata-fisheye camera does not capture panoramas with uniform resolution. Here, we separately analyze the spatial variations in the vertical and horizontal resolution.

We know that, if the lens and the mirror are coaxial, then radial lines in the captured images map to vertical lines in the panorama. Hence, the vertical resolution (pixels per unit angle) of the panorama at a given θ is defined as $dr/d\theta$, where r is the radial distance of the corresponding pixel from the center of the captured image. Within the fisheye's FOV, the vertical resolution depends directly on the the fisheye projection function. Our prototype camera has a fisheye lens with a equidistant projection function (that is $r = k\theta$). Therefore, the vertical resolution of the fisheye panorama is uniform as $dr/d\theta = k$. On the other hand, the vertical resolution of the panorama in the mirror's FOV is determined by the mirror's geometry and position. We find a numerical approximation to the vertical resolution by computing the change in r for a small change in θ ($\Delta r/\Delta\theta$). Figure 9a shows the vertical resolution map for the panorama produced by our prototype camera. While the resolution is highest and uniform in the fisheye lens' FOV, it is lower in the mirror's FOV owing to its curvature.

The horizontal resolution is just the total number of pixels in the captured image that make up each scan line in the panorama. Figure 9b shows the horizontal resolution map. Note that, in contrast to the panoramas captured by other catadioptric cameras, the resolution of the cata-fisheye panorama is not smooth. However, in applications such as video conferencing, the lower part of the panorama is used to image closer objects (e.g. the table the camera sits on). In such cases, the resolution characteristics of our camera prove to be advantageous.

8. Discussion

While our current prototype uses a spherical mirror with a fisheye lens that has an equidistant projection model, convex mirrors with other profiles may be more suitable in the case of lens with different projection functions. We will explore the relationship between the shape of the mirror and the fisheye projection model in our future work.

Our approach has two main limitations that must be con-




Figure 9. (a) Vertical resolution map (in pixels per degree). (b) Horizontal resolution map (in pixels per scan line). See Section 7 for details.

sidered before it is used in any given application. First, it has an inherent parallax that makes it unsuitable for imaging objects that are very close to the camera. Second, it has a spatially varying resolution where the upper part of the panorama is captured at a much higher resolution than the lower part. This makes it unsuitable for applications that require the panoramas to be captured with more of less uniform resolution.

Acknowledgment

This research was conducted at the Computer Vision Laboratory at Columbia University. It was supported by a grant from the Office of Naval Research (No. N00014-05-1-0032).

References

- [1] <http://www.apple.com/quicktime/technologies/qtvr/>.
- [2] <http://maps.google.com/>.
- [3] <http://www.microsoft.com/uc/products/roundtable.msp>.
- [4] S. Baker and S. Nayar. A Theory of Single-Viewpoint Catadioptric Image Formation. *IJCV*, 35(2):175–196, 1999.
- [5] H. Bakstein and T. Pajdla. Panoramic mosaicing with a 180° field of view lens. In *Omnidirectional Vision*, pages 60–67, 2002.
- [6] A. Basu and S. Licardie. Alternative models for fish-eye lenses. *Pattern Recognition Letters*, 16(4):433–441, 1995.
- [7] S. S. Beauchemin, R. Bajcsy, and G. Givaty. A unified procedure for calibrating intrinsic parameters of fish-eye lenses. In *Vision Interface*, pages 272–279, 1999.
- [8] S. Bogner. Introduction to panoramic imaging. In *IEEE SMC Conference*, pages 3100–3106, 1995.
- [9] D. Buchele and W. Buchele. US Patent No. 2,638,033: Unitary catadioptric objective lens system, 1953.
- [10] J. Chahl and M. Srinivasan. Reflective surfaces for panoramic imaging. *Applied Optics*, 36(31):8275–8285, 1997.
- [11] D. Chapman and A. Deacon. Panoramic imaging and virtual reality: Filling the gaps between the lines. *ISPRS*, 53(6):311–319, 1998.
- [12] J. Charles, R. Reeves, and C. Schur. How to build and use an all-sky camera. *Astronomy Magazine*, 1987.
- [13] C. Cruz-Neira, D. Sandin, and T. DeFanti. Virtual reality: The design and implementation of the cave. In *SIGGRAPH*, pages 135–142, 1993.
- [14] R. Cutler, Y. Rui, A. Gupta, J. Cadiz, I. Tashev, L. He, A. Colburn, Z. Zhang, Z. Liu, and S. Silverberg. Distributed meetings: A meeting capture and broadcasting system. *ACM Multimedia*, 2002.
- [15] J. Davis, K. Castle, M. Todd, T. Stuhlinger, and M. Ruda. US Patent No. 5,841,589: Panoramic optical assembly having an initial flat reflective element, 1998.





Figure 10. (a) Image captured using our prototype camera operating at F/5.6. (b) SIFT feature detection and matching applied to the two panoramic strips generated for the overlapping regions in the views of the mirror and fisheye lens before calibration. For clarity, only 20% of the features are depicted in the figure. (c) Alignment of the features after the calibration process with an average alignment error of 0.087° . The position of the mirror's center of curvature is at $(-0.1492, -0.4463, 77.7345)$ (in mm) relative to the fisheye lens' viewpoint. The mirror's orientation is $(161.53^\circ, 0.822^\circ)$. (d,e) These images together represent the computed panorama whose resolution is 3600x550 pixels.

- [16] J. Hong, X. Tan, B. Pinette, R. Weiss, and E. Riseman. Image-based homing. *IEEE Control Systems Magazine*, 12(1):38–45, 1992.
- [17] U. Iwerks and S. Oaks. US Patent No. 3,118,340: Panoramic motion picture camera arrangement, 1964.
- [18] S. Lin and R. Bajcsy. True single view point cone mirror omni-directional catadioptric system. In *ICCV*, pages II: 102–107, 2001.
- [19] D. Lowe. Distinctive image features from scale-invariant keypoints. *IJCV*, 20:91–110, 2003.
- [20] D. McCutchen. US Patent No. 5,703,604: Immersive dodecahedral video viewing system, 1997.
- [21] K. Miyamoto. Fish eye lens. *JOSA*, pages 1060–1061, 1964.
- [22] V. Nalwa. A true omnidirectional viewer. Technical report, Bell Laboratories, Holmdel, NJ 07733, USA, 1996.
- [23] S. Nayar. Catadioptric omnidirectional camera. In *CVPR*,



(a)



(b)



(c)




Figure 11. (a) Captured image. (d, e) These images together represent the computed panorama whose resolution is 3600x550 pixels. (b, c) Two perspective views computed from the spherical panorama.

pages 482–488, 1997.

- [24] S. Nayar and V. Peri. Folded Catadioptric Cameras. In *Panoramic Vision*, pages 103–119. 2001.
- [25] D. Rees. US Patent No. 3,505,465: Panoramic television viewing system, 1970.
- [26] D. Slater. Panoramic photography with fisheye lenses. In *International Association of Panoramic Photographers*, 1996.
- [27] W. Smith. *Modern optical engineering : the design of optical systems*. McGraw-Hill, 2000.
- [28] R. Swaminathan and S. Nayar. Polycameras: Camera Clusters for Wide Angle Imaging. Technical report, 1999.

- [29] Y. Yagi and S. Kawato. Panorama scene analysis with conic projection. In *IEEE International Workshop on Intelligent Robots and Systems*, volume 1, pages 181–187, 1990.
- [30] K. Yamazawa, Y. Yagi, and M. Yachida. Omnidirectional imaging with hyperboloidal projection. In *International Conference on Intelligent Robots and Systems*, 1995.
- [31] S. D. Zimmermann. US Patent No. 5,185,667: Omniview motionless camera orientation system, 1993.