Chapter 3

Catadioptric Systems

This chapter deals with catadioptric camera systems and starts by introducing the different types of mirrors that can be used, highlighting in particular which of the mirrors satisfy the single viewpoint constraint, and mentioning that one of the mirror types comes very close to achieving a full sphere of view. This is followed by a comparison between the omnidirectional images generated by the different mirror types. And finally the chapter concludes by reviewing some existing applications that use catadioptric camera systems.

3.1 Catadioptric Cameras

As mentioned in §2.3.4, **catadioptric cameras** are omnidirectional cameras that combine curved mirror(s) with conventional lens-based cameras to acquire wide fields-of-view (usually hemispherical or more). The camera captures what the mirror sees. The word **catadioptric** is coined from the words *catoptrics* and *dioptrics*. Catoptrics is the study of reflecting surfaces (mirrors), while dioptrics is the study of refracting elements (lenses).

The first known use of mirrors to capture a wide field-of-view are the *all-sky cameras* used in meteorology and astronomy, starting from the 1950s [NAYA97; CHAR87], most of which used arbitrarily-shaped convex mirrors. But it wasn't until 1970 that D.W. Rees introduced the first catadioptric camera designed specifically to have a single viewpoint. His system makes use of a *hyperboloidal* mirror and a camera equipped with a standard perspective lens [REES70]. Rees was also the first to use the camera system for generating perspective views from the omnidirectional images[1].

In the early 1990s, other different types of mirrors were used for computer vision applications, such as *conical* and *spherical* mirrors (mainly used in the area of robotics). Later in 1997, S. Baker and S. Nayar, performed geometric analysis of single-mirror catadioptric camera systems and developed a design using a *parabolic*-shaped mirror with an orthographic lens camera [BAKE98]. The late 1990s saw an increase in the use of catadioptric omnidirectional cameras in computer vision applications, with many publications on the aspects of omnidirectional systems, and in 2000, the start of the IEEE's workshop on omnidirectional vision.

3.2 Different Types of Catadioptric Systems

As briefly mentioned in the previous section, different mirror shapes have been used. Each mirror type has its own characteristics and different advantages and disadvantages. At first, the space of all possible mirror designs may appear limitless. But if the requirement for a single viewpoint is to be satisfied and if the systems are restricted to those with a single mirror and a single-lens camera, then the range of possible mirror shapes is much smaller. In fact, [BAKE99] have shown that the complete class of such mirrors consists only of those mirrors which have a 2D profile of a *conic section* – that is, the mirrors are 3D *swept conic sections* [2]. The series of Figures 3.1 to 3.6 below show catadioptric system based on different types of conic-shaped mirrors.

3.2.1 Planar Mirror

Not all conics produce useful omnidirectional cameras – the line conic is one such case: it generates a planar mirror (see Figure 3.1) that offers no increase in the field-of-view of the camera (2θ in this case). By positioning the camera so that its pinhole (indicated by ● in the figure) coincides with the single viewpoint (indicated by ⊙) of

[1] J.R. Charles in the late 1970s, developed a darkroom technique for converting the circular images of photographs taken with an all-sky cameras to rectangular prints. But the mirrors he used did not have a single viewpoint and the rectangular output format is not (geometrically correct) perspective.

[2] It's quite straightforward to see that the mirror must be symmetric about the optical axis of the camera lens, and therefore the 2D section profile of the mirror defines the 3D shape of the mirror.

the system (camera outline shown as dashed lines) and removing the mirror, we end up with the same field-of-view. The planar mirror is not a useful omnidirectional camera.

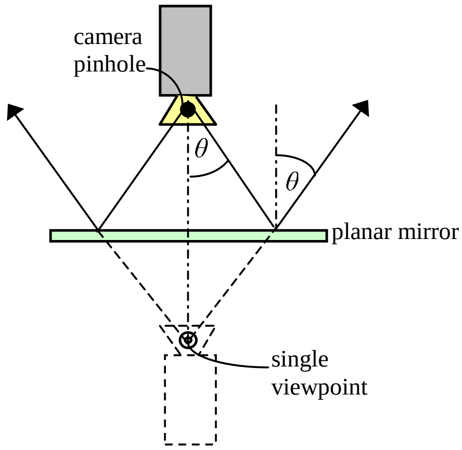

Figure 3.1: Planar mirror

3.2.2 Conical Mirror

For some other conics, the location of the single viewpoint in relation to the mirror and camera, does not allow practical catadioptric systems to be constructed. This is the case for the cone-shaped mirror (see Figure 3.2(a)), where the single viewpoint is located at the apex of the cone. It's not physically possible to place the camera's pinhole at the tip of the cone. In addition, with the pinhole at the tip of the cone, the only rays of light that can enter it are those that graze the mirror – rays ℓ_1 and ℓ_2 never reach the pinhole. This is a degenerate case [BAKE99].

But the design for conical mirrors can be adjusted slightly by moving the camera some distance away from the cone's apex, as indicated in Figure 3.2(b). The side effect is that the system now no longer has a single viewpoint – the viewpoint caustic is a circle (its position and radius depends on angle α of the mirror) [SWAM01]. This caustic circle shrinks in size and approaches the vertex as the camera is brought nearer to the cone's tip. Even though it no longer satisfies the single viewpoint constraint, the conical mirror is often used in robot-based applications [BRAS00].

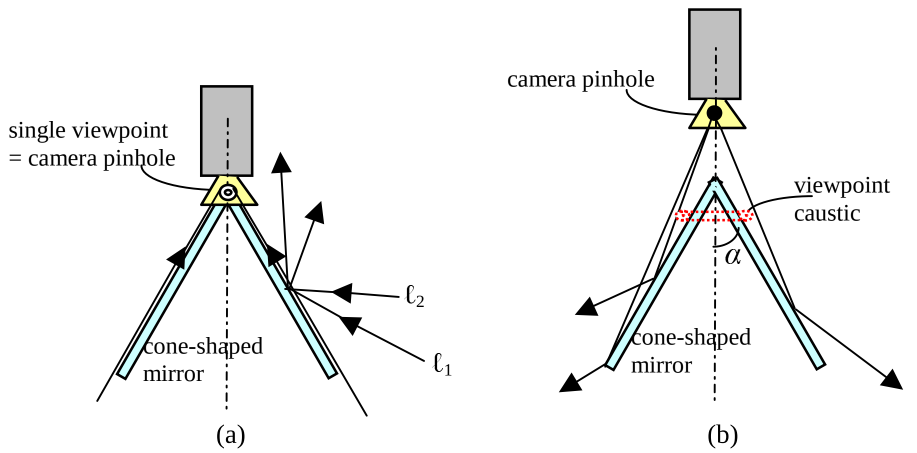

Figure 3.2: Conical mirror. (a) Degenerate case; (b) Non single viewpoint version

Advantages of using a conical mirror:

A1. Conical mirrors are quite easy to manufacture [ISHI98].

A2. Omnidirectional images acquired by conical mirrors have good resolution at the periphery, compared with the rest of the conic-shaped mirrors.

Disadvantages:

D1. Conical mirrors do not have a single viewpoint.

D2. Conical mirrors tend to suffer from high astigmatism [ISHI98].

3.2.3 Spherical Mirror

The spheroidal mirror, shown in Figure 3.3 below, is another degenerate case. The single viewpoint and camera pinhole position both coincide at the centre of the sphere and the camera system only sees itself – the ray of light ℓ, leaves the pinhole, reflects off the internal surface of the mirror, and passes back through the pinhole. A variation is to place the camera at some arbitrary position outside the sphere, at the expense of violating the single viewpoint constraint.

Advantages:

A1. Spherical mirrors are the easiest type of mirror to manufacture and they tend to be of high enough quality.

A2. Spherical mirrors have the widest field-of-view of all the conic-shaped mirrors [YAGI99].

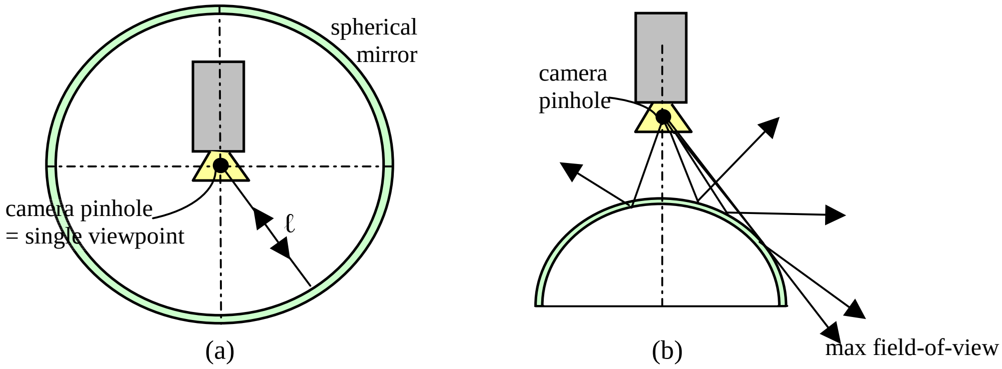

Figure 3.3: Spherical mirror. (a) Degenerate case; (b) Non single viewpoint version

Disadvantages:

D1. Spherical mirrors do not have a single viewpoint.

D2. Images acquired by these mirrors have good resolution in the centre of the image, but poor resolution at the periphery. In fact, of all the conic-shaped mirrors, the sphere creates the most distortion.

3.2.4 Ellipsoidal Mirror

The catadioptric camera based on a concave-shaped ellipsoidal mirror is the first system that offers an increase in the field-of-view while at the same time maintaining a single viewpoint. As Figure 3.4 shows, the camera's pinhole must be positioned to coincide with the ellipsoid's secondary (upper) focus, and the single viewpoint will then be at the ellipsoid's primary focus. The maximum possible field-of-view for the ellipsoidal mirror, a hemisphere of view, occurs when the mirror's edge lies in the same horizontal plane as the primary focus.

Advantages:

A1. The ellipsoidal mirror satisfies the single viewpoint constraint.

Disadvantages:

D1. The camera must be placed exactly at the ellipsoid's secondary focus, with very little tolerance for misalignment errors. This also complicates calibration of the catadioptric system.

D2. The maximum possible field-of-view is limited to a hemisphere.

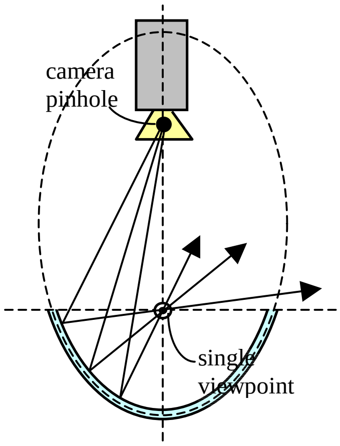

Figure 3.4: Ellipsoidal mirror

3.2.5 Hyperboloidal Mirror

The hyperboloidal mirror (Figure 3.5) also increases the field-of-view of the camera and has a single viewpoint. Compared to the ellipsoidal mirror, the hyperboloid can view much more than a hemisphere. The pinhole of the camera must be positioned at the secondary (upper) focus of the hyperboloid – this secondary focus is the focus of the other hyperbola (shown as a dashed curve in Figure 3.5)[3].

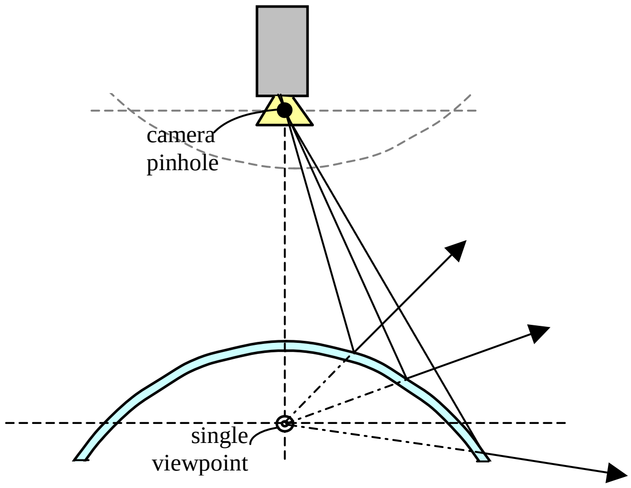

Figure 3.5: Hyperboloidal mirror

Advantages:

A1. Hyperboloidal mirrors satisfy the single viewpoint constraint.

---

[3] The two hyperbolic conic sections form a 3D non-ruled quadric called the *hyperboloid of two sheets* [HART00 pp.55-56]. The lower sheet is used for the shape of the mirror with its focus being the single viewpoint of the system. The upper sheet's focus serves as the position for the pinhole of the camera.

A2. Can achieve a very large field-of-view.

Disadvantages:

D1. It is very difficult to manufacture hyperboloidal mirrors, especially of good quality, and with high curvatures.

D2. Similar to the case of the ellipsoidal mirror, catadioptric systems based on hyperboloidal mirrors must have an exact alignment between the camera and the mirror, with very little tolerance for misalignment errors.

3.2.6 Paraboloidal Mirror

The use of a paraboloidal mirror in catadioptric cameras presents a special case. It is a well-known fact in optics that a parabola can be used to collect and focus light – this is called the *optical property of a parabola*. Rays of incoming light parallel to the axis of a parabolic mirror are reflected by the concave (inner) surface and pass through the focus of the parabola (shown in red in Figure 3.6 below).

By symmetry, rays of incoming light that pass through (or with a direction towards) the parabola's focus, will be reflected by the convex (outer) surface of the parabolic mirror in a direction parallel to the parabola's axis. Therefore, a paraboloidal mirror produces an *orthographic projection*, as can be seen in Figure 3.6, meaning that an orthographic camera is required. The common way of converting a normal perspective camera to an orthographic camera is to use a *telecentric lens* [WATA96][4].

The orthographic projection means that, contrary to the ellipsoidal and hyperboloidal mirrors, there are no restrictions on where the camera should be placed. The camera can be translated laterally and moved away or towards the mirror, as indicated in Figure 3.6, as long as the camera's orthographic axis is kept parallel to the mirror's axis. One use of this property is, for example, to allow the camera to dynamically 'zoom into' (move towards) the mirror to view a higher-resolution part of the omnidirectional image [BAKE99].

[4] Using a perspective camera with a paraboloidal mirror, instead of an orthographic camera, violates the single viewpoint constraint.

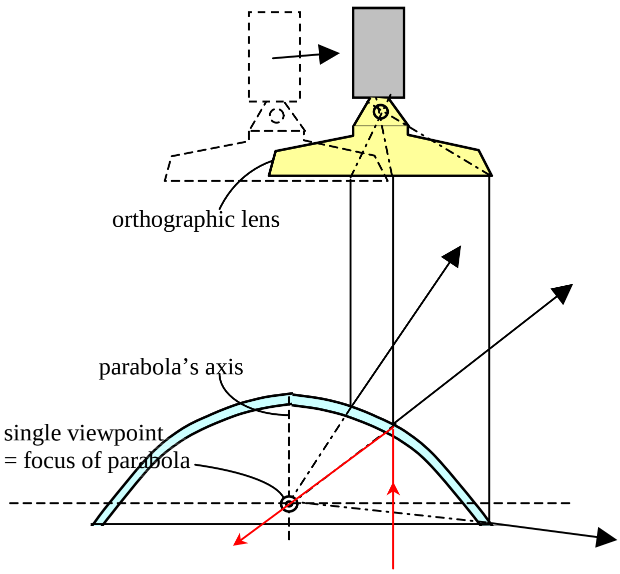

Figure 3.6: Paraboloidal mirror

Advantages:

A1. The paraboloidal mirror satisfies the single viewpoint constraint.

A2. These mirrors can achieve a very large field-of-view.

A3. No alignment between the mirror and the camera is required. This simplifies the task of calibrating the catadioptric system.

Disadvantages:

D1. Because of the orthographic projection, a telecentric lens (or equivalent) is required. These types of lenses tend to be expensive and large in size, so creating a bulky design.

3.3 Achieving a Full Sphere of View

Both the hyperboloidal and paraboloidal mirrors, and to a lesser extent the ellipsoidal mirror, can be used to build sensors with omnidirectional views that satisfy the single viewpoint. But it's not possible to get a truly full sphere of view with a single-mirror system as the mirror naturally hides part of the view.

In the case of the paraboloid, incoming light rays that make a 90° angle with the mirror's axis are reflected towards the camera when they strike the mirror in the horizontal plane containing the paraboloid's focus. Therefore, by 'cutting' the paraboloid exactly at the horizontal plane containing the focus, the field-of-view is exactly a hemisphere. More importantly, the single viewpoint is located exactly on the base of the mirror. Using two such mirrors, and aligning them back-to-back so that their individual single viewpoints coincide, one gets a catadioptric system that is truly omnidirectional – with a 360° by 360° [NAYA97]. This setup is shown in Figure 3.7.

One problem with all catadioptric systems is the blind spot caused by the camera. In addition, other obstructions of the field-of-view could arise from the support needed to hold the mirror and camera in place. One way of getting rid of this is to use a transparent support, for example, a glass tube to connect the mirror to the camera. Unfortunately, the glass surface causes *internal reflections* and the glass material can distort the view if not of high enough quality. Paraboloid-based systems have an advantage that they eliminate the internal reflection problem by their use of orthographic projection. For the other mirrors, [ISHI98] uses a novel idea of a black needle along the axis of the mirror and between the mirror and lens of the camera to eliminate internal reflections.

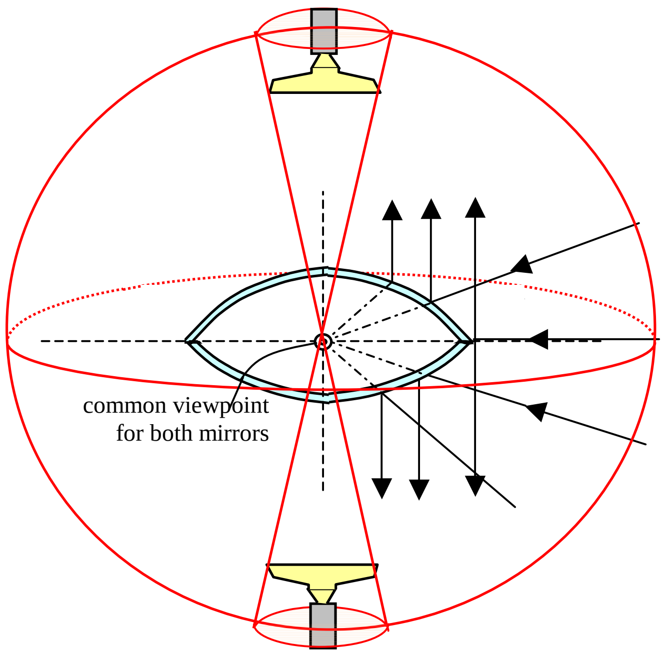

Figure 3.7: Achieving 360° by 360° with two paraboloidal mirrors

Another problem common to all catadioptric systems is that they tend to be quite large. Also the camera is usually placed at a certain distance from the mirror to minimise the blind spot caused by the camera. One solution to this problem is through the use of folded optics. *Optical folding* is a well-known method in optics, commonly used in astronomical telescopes and microwave devices. It involves using additional (secondary) mirrors to divert the path of light to achieve more compact designs [NAYA99]. A great advantage of optical folding is that it can be used to correct for the effect of *field curvature* (see §2.4.4, D3) of the primary mirror by designing the secondary mirror(s) with a negative curvature that cancels that of the primary mirror[5]. And can be used to conceal parts of the system, for example, the camera can be hidden behind the primary mirror, as in the original patent of D.W. Rees [REES70]. It has been shown in [NAYA99], that any folded system with two conic mirrors is geometrically equivalent to a single conic system.

Finally, one can mention the *panoramic annular lens* described in [ZHU99]. This consists of a glass block (the lens) with reflective coatings on two sides of the glass block acting as the 'mirrors'. Two versions are presented: one using a hyperboloid and ellipsoid; the second consists of an ellipsoid and a paraboloid giving an orthographic projection. The advantages of the PAL system are that there are no misalignment problems (mirrors' position fixed relative to each other and the lens), the system can be made very small, and the orthographic version has all the advantages of the single paraboloid mirror. The drawbacks are: a large obstruction and a limited vertical field-of-view (most of it along the horizon).

3.4 Simulations

In §3.2, the different types of useful catadioptric omnidirectional cameras have been explored. While working on this project, it was thought that it would be interesting to do a comparison of the omnidirectional images generated by each of these catadioptric cameras. In the absence of actual cameras to experiment with, the nearest thing that could be done was to generate synthetic views. It was decided to use ray-tracing techniques to simulate each of the cameras. This is far from ideal, but at least provides an indication of what type of output one can expect from each of the different mirrors – some indication of the mirror's (non-uniform) resolution, image distortion and focal blur.

[5] Optical folding can be implemented in a way to preserve the single viewpoint property of the system, that is, the final effective viewpoint of the whole system – the individual mirrors on their own may not have a single viewpoint.

An artificial room-like scene was created, consisting mostly of straight lines and planar objects. The simulated catadioptric systems were positioned in the centre of this room. The images were ray-traced using the *POV-ray* software package [POV02], chosen because of its support for mathematically-defined surfaces (for the mirrors), simulation of focal blur, and support for both perspective and orthographical camera projections. Two normal perspective views of the room are shown in Figure 3.8, taken from the top opposite corners.

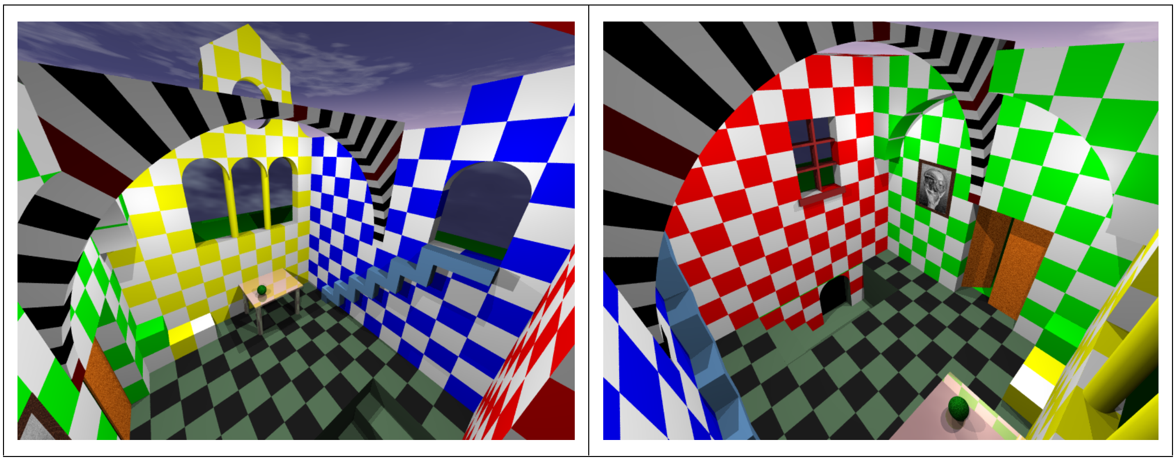

Figure 3.8: Simulations – the room

Figure 3.9 shows the scene as viewed by a spherical mirror and a cone-shaped mirror, both with a perspective camera. As discussed in §3.2, these catadioptric systems do not obey the single viewpoint constraint.

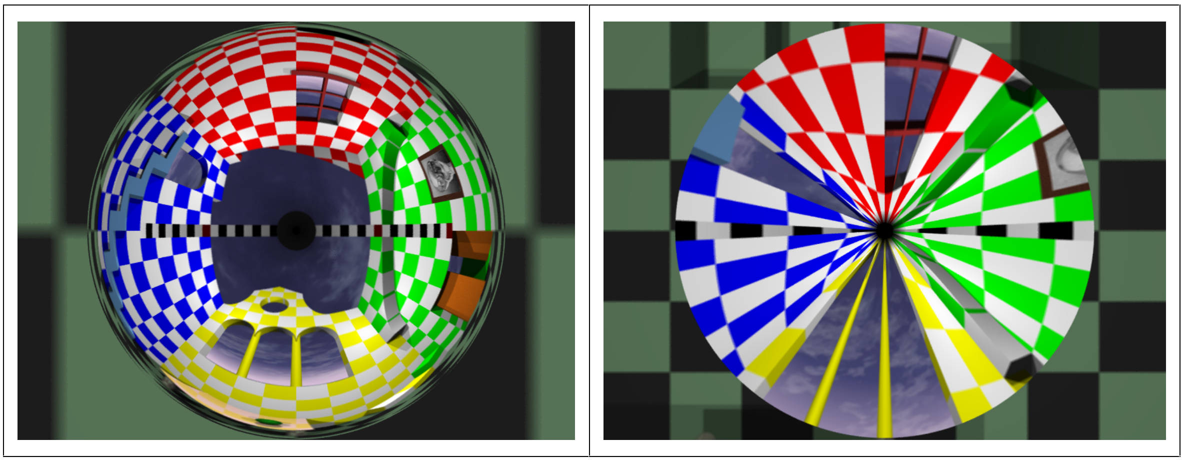

Figure 3.9: Simulations – Spherical and Conical mirrors

In this case, the spherical mirror has a vertical field-of-view of 340º – nearly a full circle! All the four walls, the sky and parts of the floor are visible. But the spherical mirror produces a lot of distortion, especially near the periphery of the image. The resolution in the central part of the image is very high, but then rapidly falls off the nearer to the edge.

The cone-shaped mirror used in this simulation has an internal cone angle of 120º – that is, the angle made by the two sides of the cone and the cone's apex. The vertical field-of-view of the cone is restricted to the horizon, starting from 0º at the horizon and goes up to an elevation of ~50º. The resolution is highest near the horizon (at the periphery of the image). The conical suffers from a lot of uneven distortion. The blurred black circle at the centre of the image is the reflection of the perspective camera.

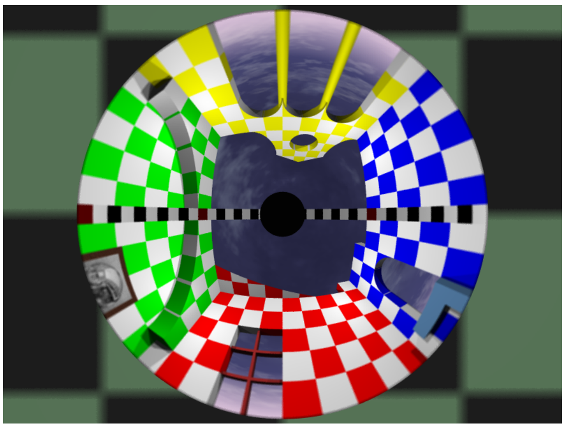

Figure 3.10: Simulations – Ellipsoidal mirror

Figure 3.10 above, shows the output of the ellipsoidal mirror. This has a vertical field-of-view of 180º (upper hemisphere), the maximum possible for this type of mirror. The resolution is non-uniform across the image, but it changes at a more gradual pace than that of the spherical mirror. Because this mirror is the only one having a concave shape, the image appears inverted, compared to the others.

The first image of Figure 3.11 shows the paraboloidal mirror, the orthographic camera (grey cylinder) and the support (a central black rod). Similar setups were used for the other simulations[6]. The second shows the image captured by this catadioptric camera.

[6] Except for the concave ellipsoid mirror, where a central support cannot be used as it would pass through the single viewpoint.

The paraboloidal mirror used in this simulation has a vertical field-of-view of 220º. Resolution is non-uniform and is highest near the edge of the image[7].

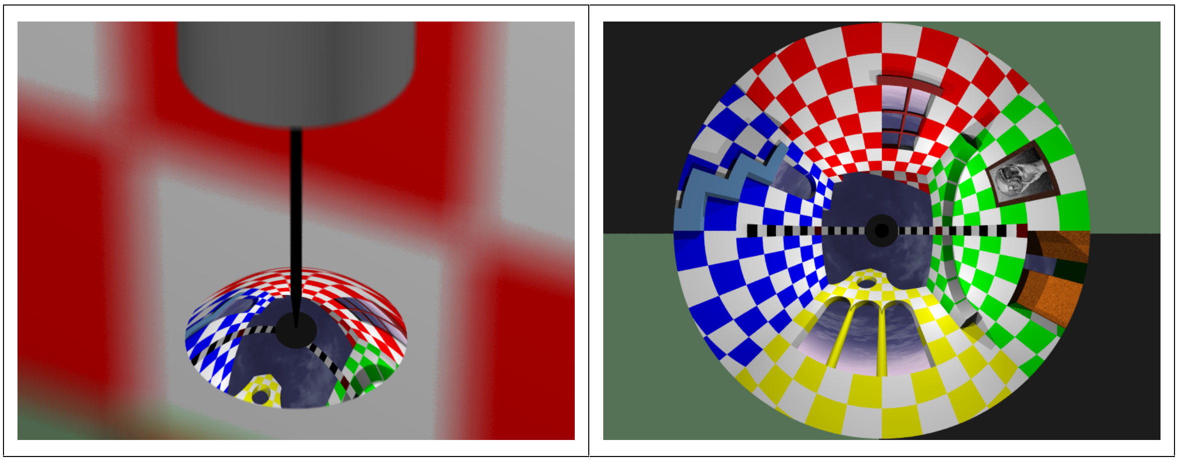

Figure 3.11: Simulations – Paraboloidal mirror

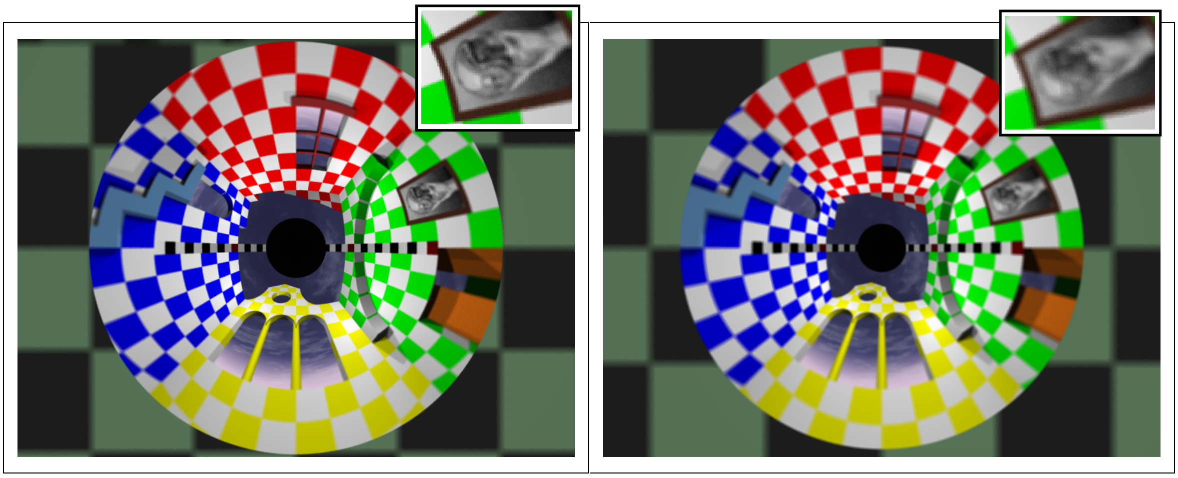

Figure 3.12: Simulations – Hyperboloidal mirror (without and with misalignment)

Figure 3.12 shows two views with a hyperboloidal mirror. The vertical field-of-view of the mirror used in this simulation is of 220º. There is a slight focal blur at the periphery in the first image due to the mirror's curvature. Hyperboloid and ellipsoid based catadioptric systems require the camera's pinhole to be placed exactly at the mirror's secondary focus. For the second image, the camera was moved slightly away

[7] The image generated by the ellipsoidal mirror suffers from some focal blur at the edges, due to the field curvature effect of the mirror. The paraboloid image does not show such blur. This is due to a limitation of POV-Ray – as of version 3.5, POV-Ray does not support focal blur for the orthographic projection (but only for the perspective projection).

from the hyperboloid's second focus. The second image shows that focal blur increases slightly.

The final simulation consisted of setting up two paraboloidal cameras back-to-back, as mentioned in §3.3, to achieve a 360º by 360º view. The two output images (upper and lower hemispheres) are shown in Figure 3.13. These omnidirectional images were then used to generate a 360º by ~320º panorama shown in Figure 3.14 below[8].

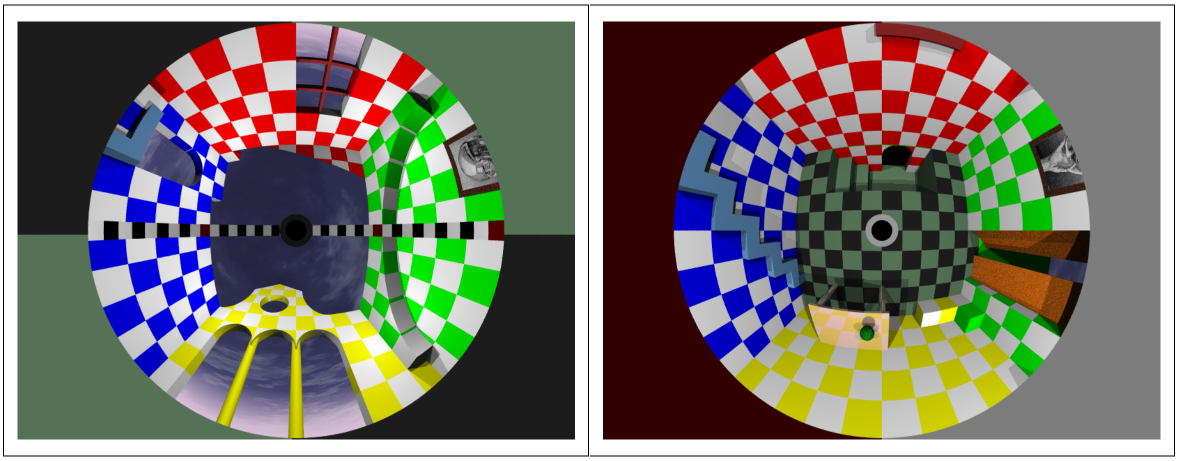

Figure 3.13: Two paraboloidal systems to achieve a full view

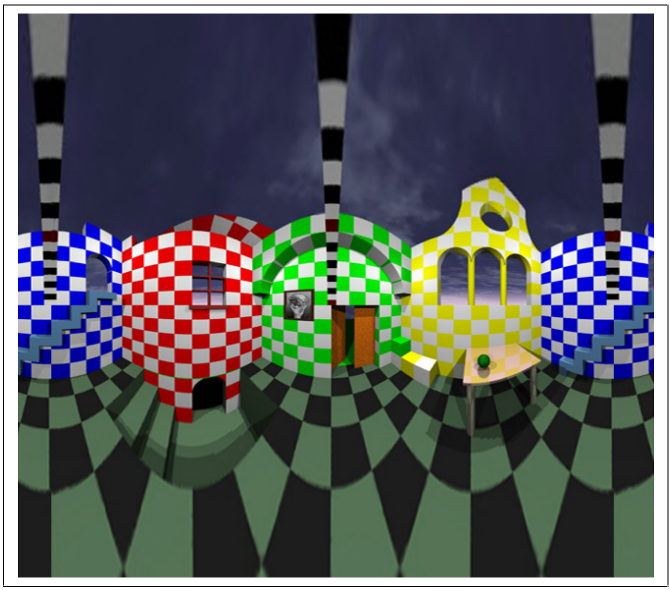

Figure 3.14: Panorama captured by two paraboloidal systems

[8] The re-projection of these two omnidirectional images to a cylindrical panorama was done with the OmniTracking application written for this project. The two images were processed separately, as in its current implementation, the program accepts only one omnidirectional video stream as input. The two panoramas (upper and lower hemispheres) were then placed on top of each other to get Figure 3.14. Also, the source images were converted into a short video stream by duplication of frames.

3.5 Applications using Catadioptric Cameras

This section gives a short list of some of the applications of catadioptric cameras as omnidirectional sensors in computer vision.

Table 3.1: Different mirror types and applications using them

mirror type	Applications
conical	*(SYCLOP sensor)* [BRAS00], *(COPIS sensor)* [YAGI95], [BALD99],
spherical	[THOM03], [WINT00], [ASOH01], *(spher. C-ODVS)* [ISHI98]
ellipsoid	
hyperboloid	*(HyperOmni sensor)* [MORI03; YAMA02], [SOGO00], [HUAN01], *(hyperb. C-ODVS)* [ISHI98]
paraboloid	*(ParaCamera)* [BOUL98b], [TAYL00], [RUI01], [STIE02], *(LOTS)* [BOUL99]
others	[ADOR02], *(pyramid of mirrors)* [BOUL98a], *(SVAVISCA sensor)* [GÄCH01], *(PAL lens)* [ZHU99].

3D Reconstruction and Image-based Rendering

- [TAYL00] implements a system, called *VideoPlus*, which uses a paraboloid omnidirectional camera for capturing the structure of environments and reconstructing a 3D model. The camera is mounted on mobile robot. Virtual walkthroughs can then be generated from the original video stream. Omnidirectional cameras offer the advantage that fewer images need to be taken of a particular scene.
- [ALIA01b] uses a paraboloidal mirror mounted on a mobile robot for implementing an image-based rendering application, based on the idea of the plenoptic function. The robot moves along an irregularly shaped grid and from this set of video streams, virtual walkthroughs can then be constructed (rendered) from any arbitrary point within the grid.

Tele-presence

- [BOUL98b] uses a parabolic camera, called the *paracamera*, in an application called *Remote Reality*. The video output of the camera is used to generate perspective views shown on a head-mounted display, with the possibility of many different observers concurrently looking at different directions. The system is also applied for the remote operation of a vehicle.

- [MORI03] uses a hyperbolic mirror (the *HyperOmni sensor*), for a real-time tele-presence application. The system runs over the Internet, where traditional cameras would fail due to the significant delays between a user altering his viewpoint and the remote camera responding to the command.
- [BALD99] uses a cone-shaped mirror for an immersive tele-presence application with a mobile robot in low-bandwidth environments.
- [THOM03] uses a spherical mirror (the *panospheric camera*) for a tele-viewing and tele-operation of a robot.

Robot localisation and ego-motion

- [BRAS00] uses a conical mirror for a sensor, called the *SYCLOP sensor*. This is used in an application for localisation of mobile robots. The non single viewpoint of the mirror is not a major drawback for robot-based applications since most (like this paper) are only interested in locating landmarks (vertical lines in this case) for navigation. They also implement a stereoscopic version of the application.
- [WINT00] uses a spherical mirror for robot navigation. Approximate models are used to partially correct the distortion of the spherical mirror – using certain assumptions about the ground-plane and its relation to the camera's orientation. Robot localisation and navigation are made simpler with omnidirectional images, because of the wider field-of-view.
- [ASOH01] also use a spherical mirror for a localisation and navigation application with a robot called *Jijo*. Edge features from a very low-resolution omnidirectional image are used as the navigation features. The robot is also equipped with another high-resolution camera for face recognition. This high-resolution camera is controlled based on the output of the omnidirectional camera.
- [ADOR02] uses a hybrid sensor, called *HOPS*, that combines a traditional camera to simulate high-resolution (foveal) imaging and an omnidirectional sensor for peripheral vision. This sensor is used for obstacle avoidance in robot navigation.

Meeting viewing

- [RUI01] uses a parabolic mirror to capture omnidirectional video of a meeting for later, offline viewing. The camera used in this application has a very high resolution (1300x1030). The main advantage of using an omnidirectional camera for meetings, is that all the participants are visible at the same time.

- [STIE02] implement an application for viewing meetings that performs automatic face detection. It also determines the focus of attention of the people present in the meeting by estimating the head pose. A parabolic mirror is used in this application.

Tracking and Surveillance

- [BOUL99] uses a paraboloidal mirror for a frame-rate surveillance and tracking application, named *LOTS*. This application works in complex outdoor environments where targets can be occluded, camouflaged, etc.
- [YAMA02] uses a hyperboloidal mirror, the *HyperOmni sensor*, for motion detection and target tracking in an indoor environment.
- [GÄCH01] implements a motion detection application, using the *SVAVISCA sensor*. This sensor has a log-polar pixel density so that it automatically 'dewarps' the omnidirectional image without the need for software.

Networked systems

- [SOGO00] uses 4 omnidirectional cameras to implement an *n-ocular stereo* application for real-time tracking of people within a room. The sensors used are the compact C-ODVS sensors developed by [ISHI98].
- [HUAN01] implements a similar application to the previous one, using a network of hyperboloidal-mirror cameras. It is used for detecting persons in a room, with the addition of face localisation and tracking. The application is called *NOVA*.
- [NG99] implements a *ubiquitous vision* system to monitor dynamically changing environments. Multiple omnidirectional sensors are used arranged in a grid-like fashion. The system can be used to generate virtual walkthroughs from any arbitrary position within the environment. This paper also implements a version of their system for tracking.
- [TRIV02] uses omnidirectional cameras in conjunction with normal cameras, microphones and range sensors in a multi-sensor network for incident detection in a road monitoring application. The omnidirectional camera can also be used to control the viewpoint of standard PTZ (pan-tilt-zoom) cameras.

Although some of the above papers have either developed their own catadioptric cameras, or the cameras used can be considered as prototypes, there are several off-the-shelf models available commercially.

The list presented in this section is far from being exhaustive. More sensors and applications can be found in [YAGI99; NAYA97]. Also, the following web site [OMNI] is dedicated to omnidirectional vision and contains references and links to several research groups and commercial omnidirectional cameras.

3.6 Conclusion

This chapter described the different types of mirrors that can be used in catadioptric camera systems. It was found that some of them do not obey the single viewpoint constraint. This chapter also illustrated the omnidirectional images acquired by the different types of mirrors by means of simulated views. And finally the chapter finished off by reviewing some existing applications that use catadioptric camera systems.

As a conclusion, the best type of catadioptric camera system is the one using the paraboloidal mirror and orthographic projection. This chapter mentioned that this system provides more flexibility in the camera design, since the camera sensor does not have to be placed at any particular fixed point with respect to the mirror, like in the case of the other mirrors. This chapter also showed that by using two paraboloidal systems, one can come very close to achieving a full sphere of view, and hence a true omnidirectional system.

The rest of this document will focus on the paraboloidal catadioptric system and this is the camera used for the application developed for this thesis.