10

Omnivision systems for mobile robots

In this chapter we describe the omnidirectional vision system developed for our RoboCup small-size and mid-size mobile robots. Our system consists of a combination of a video camera with, either an elliptic concave mirror, or a convex mirror with a special semi-conical shape. The concave mirror is a remarkable low-cost solution. The convex mirror was specifically designed to provide good resolution for objects less than 3 meters away from the robot, and exponentially lower resolution for objects situated between 3 and 12 meters. This chapter shows how to design such mirrors.

10.1 Local vision

The kind of robots that we want to consider now are mobile robots with their own integrated camera, i.e., so-called "local vision" systems. Ideally, such robots could be used in an office or factory environment. We would like them to automatically adapt to the peculiarities of buildings made for humans, so that we do not have to modify the environment just to fit the robot's capabilities.

Autonomous robots move around: usually they roll on the floor using two or more wheels. The robots we consider here carry their own camera. Therefore, the camera should be portable and small, yet powerful and capable of as much temporal and spatial image resolution as possible.

Some of the main issues for local vision systems are:

- the size, weight, and quality of the cameras,
- the processing speed of the local electronics,
- fast color discrimination,
- localizing the robot in the world using visual cues
- identifying obstacles

We discuss these issues in the following sections.

In the RoboCup mid-size league, the robots have local vision and communicate among themselves. The environment consists of a 12 by 8 meters long rectangular green carpet with two goal boxes on each side, one colored yellow, the other colored blue. There are four colored poles, one in each corner. The poles to the side of the yellow goal have three stripes colored alternatively yellow, blue, yellow. The other two poles are colored blue, yellow, blue. The main problem, therefore, is to recognize four main colors: blue, yellow, green, and the orange ball. Also, if we want to recognize the lines on the field, white must be distinguished from other colors.

The environment for this robotic soccer task has been standardized. However, the colors have not. "Yellow", "blue", and "green" is defined in a loose way, and the organizers of a competition can use any kind of shade of those colors. Therefore, the local vision system must be as flexible and adaptable as possible.

10.2 Omnidirectional vision

In most cases, it is advantageous to have a camera which is not just facing forward, but which can perceive the surroundings of the robot, that is, a camera which can get a full 360 degrees view of the space around the machine. Such vision systems are called *omnidirectional*, and there exist many variations of them. The simplest approach, probably, is to have a camera oriented upwards or downwards, with a wide-angle capable of providing a 180 degree field of vision. There are some security cameras based on this principle. However, the lenses for such wide angulars tend to be expensive and heavy.

The second alternative is to use several cameras mounted in such a way that they cover 360 degrees of vision. Four or five small cameras, mounted on the sides of a pentagonal or hexagonal structure can cover the required 360 degrees. The picture in Fig. 10.2 shows the "Flycam", an array of video cameras that can record a complete view of their surroundings. The image from the multiple cameras can be stitched into a single picture in real time.

The third alternative is to combine a video camera and a mirror to provide a 360 degree view of the robot surroundings. Arrangements in which a mirror is used to reflect the image, and a camera is used to capture the reflection, are called *catadioptric systems*. They are popular in mobile robotics because small mirrors and cameras can be combined to produce wide angle panoramic effects. A catadioptric mirror-camera system is the second-best alternative to our spherical retina which is able, with a small lens, to capture an almost a 180 degrees view of our surroundings. Until spherical imaging chips become available, we will have to combine standard cameras with mirrors.

Consider a sphere or a conic mirror hanging above a camera oriented upwards: both mirrors reflect the entourage into the camera and one can

Fig. 10.1. An omnidirectional system with multiple cameras

get the desired omnidirectional view. However, the resolution provided by a spherical mirror is usually very different for nearby and for distant objects. The latter receive a smaller portion of the camera chip real state. This can make the identification of distant objects difficult.

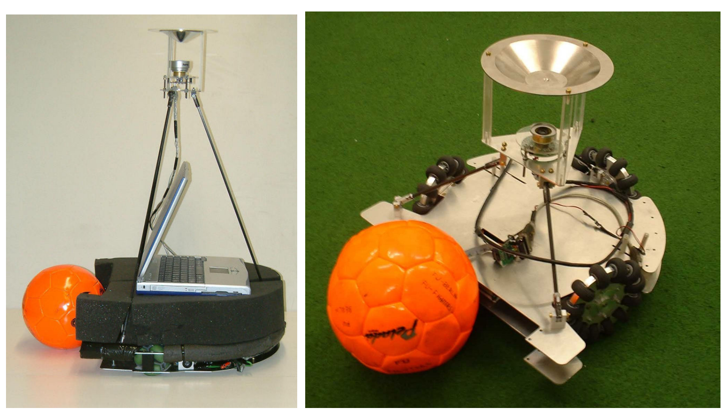

Fig. 10.2. The middle-size FU Fighters robots 2002 (left) and 2003 (right)

There is a need to find a compromise between the following conflicting objectives:

- We would like to get a long range view of the surroundings.
- Resolution for nearby and distant objects should be allocated according to the needs of the application.
- The image should be easy to focus with a small inexpensive camera.

10.3 Catadioptric systems

For omnidirectional vision, the usual arrangement is to have one camera pointed upwards towards a mirror which reflects a 360 degree view of the objects around the robot. Popular choices for the shape of the mirror are conic sections, such as spheres, paraboloids, hyperboloids, or elliptic mirrors, and this for good reasons.

Fig. 10.3 shows the geometry of the reflection from several mirrors of revolution, focused through the pinhole F of a camera. Only a vertical slice of the mirror is shown. The first mirror (a) is planar: the image from the floor is reflected to the camera pinhole F. The image is equivalent to the one which a virtual camera with a virtual pinhole at F' would obtain. This "virtual viewpoint" is important for focusing calculations. The second and third mirror (b and c) represent conic mirrors, the first one is convex, the second one is concave. A conic mirror provides a better view of distant objects without modifying the imaging angle of the camera. Notice that in the case of the concave mirror, the left and right side of the image overlap or can even switch places, according to the angle of the cone. Finally, the fourth alternative (d) represents two overlapped curved mirrors, an elliptical concave mirror and a hyperbolical convex mirror. Both of them have the property of possessing a single "virtual viewpoint" at one of the focus. That is, all rays going through a focus are reflected on the opposite focus, where we can put the camera pinhole. Notice also, that the conic mirror does not have a single virtual viewpoint, but two for every vertical mirror slice. All virtual viewpoints for a conic mirror lay on a circle around the vertical, for all vertical slices at all angles around the vertical.

Curiously enough, the only mirrors of revolution with a single virtual viewpoint, are 3D conic sections of revolution, that is, hyperboloids and ellipsoids. Paraboloids also have a single virtual viewpoint at the focus, but they send all reflected rays to a focus located at infinity. Also plane mirrors offer a single virtual viewpoint.

Fig. 10.4 shows an example of the reflection of a checkerboard pattern on a spherical and on a parabolic mirror. The distortion introduced by the spherical mirror emphasizes the zone around the center. The parabolic mirror is more even-handed. Distant objects do not become smaller too fast. This is important if the projection will be used to navigate and we want to be able to distinguish distant objects clearly.

One important property of mirrors of revolution is that straight lines emanating from the center of the camera are straight lines in the image. That is,

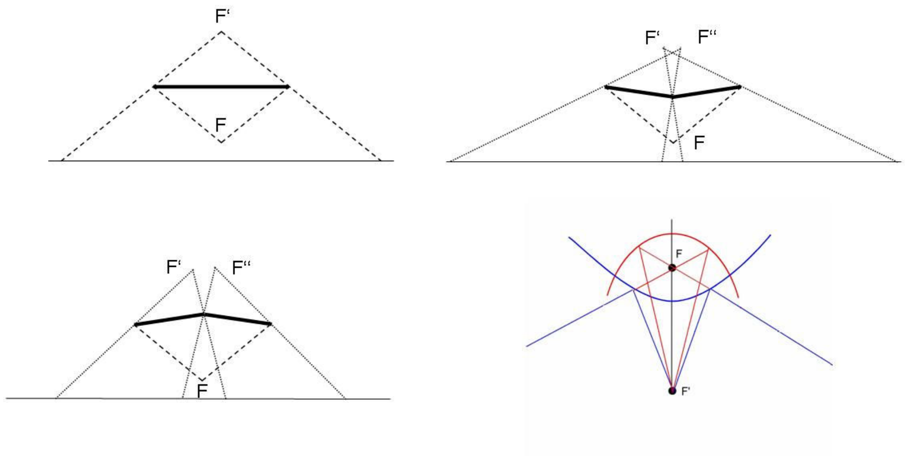

Fig. 10.3. Progression of catadioptric systems

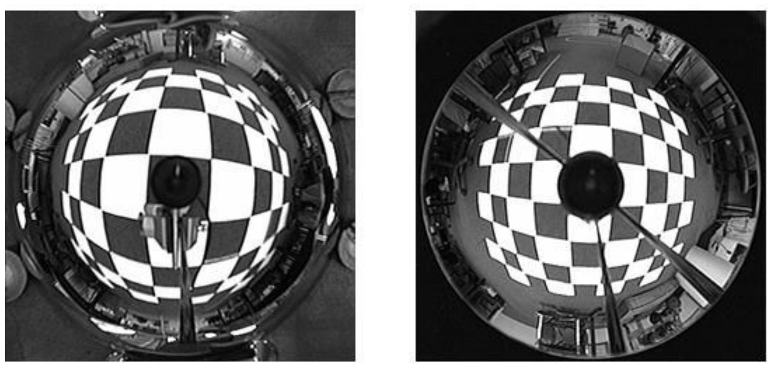

Fig. 10.4. A checkerboard pattern reflected on a sphere and on a paraboloid

the angles of lines meeting at the center of the image are invariant under mirror of revolution reflections. Although the image can become warped, angles of objects with respect to the camera can be determined directly.

The more popular catadioptric systems use a hyperbolic mirror to provide an omnidirectional image. A hyperboloid has an advantage over other kinds of shapes. Fig. 10.5 shows a hyperbola with its two foci. If we think of the upper leg of the hyperbola as a convex mirror, then any ray coming from outside and directed straight to the focus, will be reflected to the other focus. The camera can be positioned exactly there. Therefore, all incoming reflected rays will be perfectly concentrated in the camera objective and the camera will have the reflected image perfectly focused.

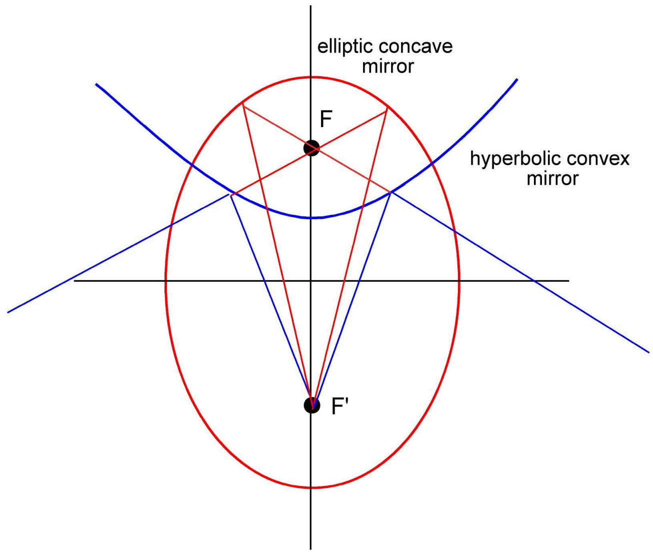

Fig. 10.5. A hyperbola and ellipse. All rays directed to the focus F are reflected towards the focus F', on the concave elliptic mirror, and on the convex hyperbolic mirror. The concave mirror requires a smaller camera opening angle for the same image.

An alternative are elliptic mirrors. The ellipse is the dual conic section of the hyperbola. As the figure shows, an elliptic concave mirror, reflects incoming rays passing through one of the foci to the other focus, where we can position the camera opening. Of course, the mirror has to be cut, in order to let light reach the upper side of the concave mirror, but the result is exactly the same as with a hyperbolic mirror. Fortunately, elliptic mirrors are easy to obtain. They are used in battery lamps and in LCD projectors. They can be bought at a fraction of the cost of a hyperbolic mirror, they can be made of glass or plastic, and are much lighter that metallic convex hyperbolic mirrors. Fig. 10.6 shows our first omnidirectional robot, built in 2000. The mirror was taken from a battery lamp costing about one dollar.

The small mirror was adapted to an omnidirectional robot which took part in the global vision challenge at RoboCup 2001 in Seattle, winning it easily.

The *distance function* is the most important feature of an omnidirectional mirror, that is the correspondence between radial distances on the imaging chip and radial distances on the floor. Fig. 10.7 shows a curve that gives the correspondence between the pixel distance from the center of a video image and the corresponding distance on the field. This distance function

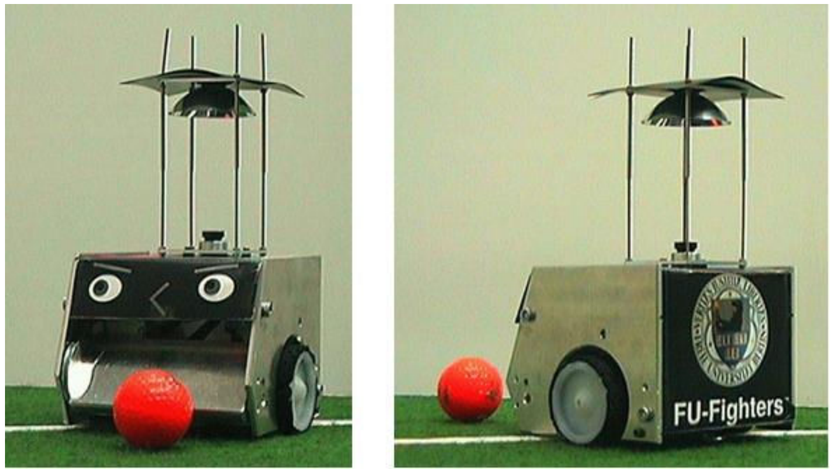

Fig. 10.6. Our first omnidirectional local vision robots

corresponds to a mirror that we designed to provide a linear mapping for the first hundred pixels from the center of the image, and the distance function of a conic mirror after 100 pixels distance. The conic mapping grows almost exponentially, until it reaches a maximum distance of 10 meters for the 200th pixel. We used a mirror with this distance function at RoboCup 20043, 2004, and 2005.

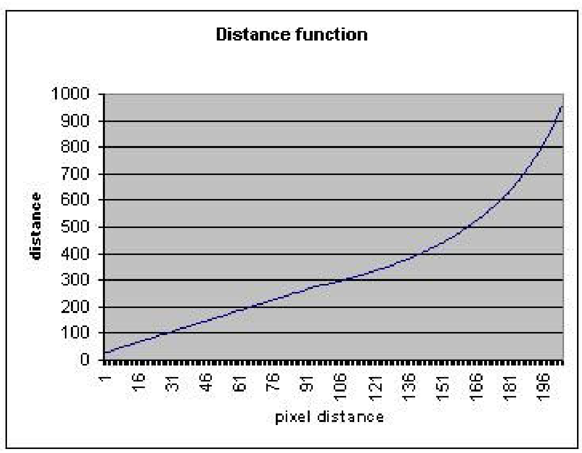

Fig. 10.7. Distance function for an omnidirectional mirror. The function has a linear phase in its first part.

The mirrors we produce are of the convex type, mainly because the machines we count on, can only mill this kind of mirrors. Figure below shows the mirrors produced at our workshops for the mid-size robots. The mirrors are cut out of an aluminum piece and are coated with a nickel alloy for better reflection

Mirrors that provide a linear distance function have been designed by several groups (cite). The image they provide looks as if it had been taken from a camera attached to the ceiling. They have the disadvantage of wasting real state on the image chip, since the image provided has to be fit inside the circular reflection, as Fig. shows.

The vision system we finally designed in 2003 had to meet the following constraints:

- The mirror could no be hung higher that 80 cm, and it had to provide an adequate view of a 12 by 8 meters field of play.
- We decided to have a linear distance function for the first three meters around the robot. This would allow us to have a very good view of the field features, especially the white lines. After the linear phase, the mirror should have the distance function of a conic mirror.
- We decided to hang the camera as far away from the mirror as possible, in order to obtain the best possible focus.

10.4 Basics of mirror design

We first have to consider the geometry of mirror reflections, in order to design our own special mirror.

An ideal pinhole camera has an infinite focus range: since all incoming rays impinge at a unique position in the CCD chip, there are no defocused images. A real video camera, however, is not a pinhole camera since low integration times can only be achieved when the optics of the camera collects more light than is possible with a pinhole. The camera lens has a certain radius and is designed to focus images on a specific plane.

If the camera is too near to the mirror, it is more difficult to get good focus. In general, the farther away the camera from the mirror, the easier it is to get an acceptable focus. But how far away the camera can go, depends on the opening angle of the optics. Cameras with a zoom are therefore easier to focus to mirrors, which can also be built smaller.

Let us consider a convex conic mirror. Its wall maintains a constant angle α with the horizontal. If a ray goes through the focal point F at an angle β with the vertical, we would like to obtain a formula for the relationship between the angle β and the distance D of the imaged point, measuring on the floor from the center of the robot.

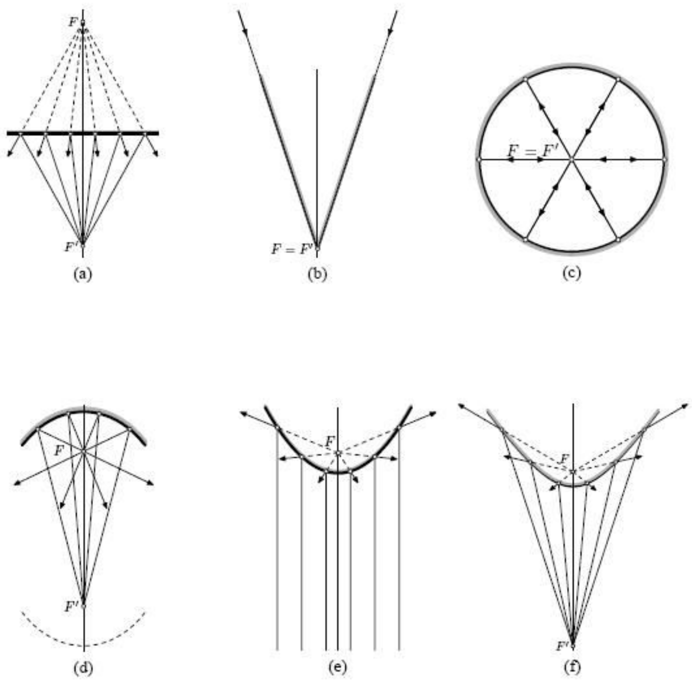

Fig. 10.8. All mirrors with a single virtual viewpoint. In the lower row, an elliptic, a parabolic, and a hyperbolic mirror. The virtual viewpoint is at a focus of the conic section. In the upper row there is a planar and circular mirror. In the circular mirror, virtual and real viewpoints coincide. The conic mirror is really just a differential surface: the camera pinhole is located at the tip of the cone, and light is reflected on the differential segments on both sides exactly at the tip of the cone, entering then the camera pinhole. The rest of the mirror could have an arbitrary shape under this arrangement.

Consider the diagram in Fig. 10.9. The camera's pinhole is located at point F. The coordinate system for the mirror shape has its origin at M. The distance between the floor and the camera's pinhole is H, and the distance between the pinhole and the tip of the mirror (located at M) is h. Assume that a ray of light L scattered by the floor enters the camera at an angle β with the horizontal, after having been reflected by the mirror. The mirror itself has an angle α with the horizontal at the point where the ray has been reflected. We can compute the distance function for this or any mirror, if we

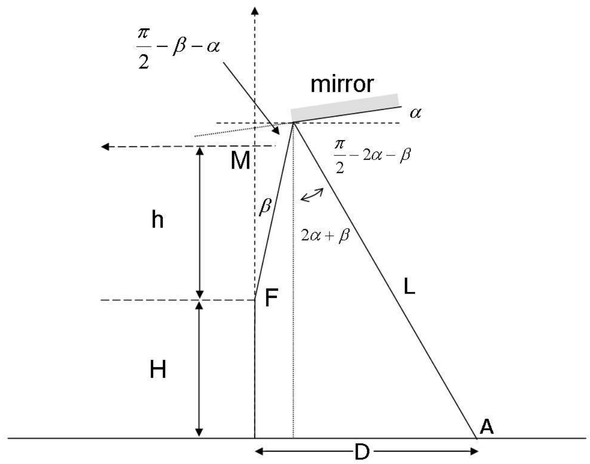

Fig. 10.9. Reflection of a point A on the floor along the ray L, reflected on a mirror towards a pinhole F.

can relate the angle β to the angle of reflection with respect to the vertical. This angle is $2\alpha + \beta$, as shown by simple geometry. The angle of reflectance on the mirror is $\pi/2 - \beta - \alpha$ because to the angle complementary to β we have to subtract α. This angle is the same for the ray L with respect to the mirror. As the diagram shows, the angle of the ray L with the horizontal is therefore $\pi/2 - 2\alpha - \beta$, and since this angle is complementary to the one we are looking for, the requested angle is $2\alpha + \beta$. Therefore, the distance D on the floor is given by

$$D = x + (H + h + f(x))\tan(2\alpha + \beta) \tag{10.1}$$

For constants H, h, and α, the distance function for a conical mirror grows as the tangent function of the angle β. This function grows slowly first (slower than the exponential function), but much faster at the end, when the argument approaches $\pi/2$. The derivative of the $\tan(x)$ function is $1/\cos^2(x)$, which diverges rapidly to infinity at $x = \pi/2$.

10.5 Computing your own mirror

In this section we show how to compute the shape of a mirror of revolution for a given distance function. The distance function $d(p)$ should be monotone. It provides the distance of a point A on the floor to the vertical axis of the camera, when A is imaged at pixel p (measured from the center of the image). In what follows we consider just a vertical slice of the mirror, since it is a mirror of revolution and has the same form around the vertical axis.

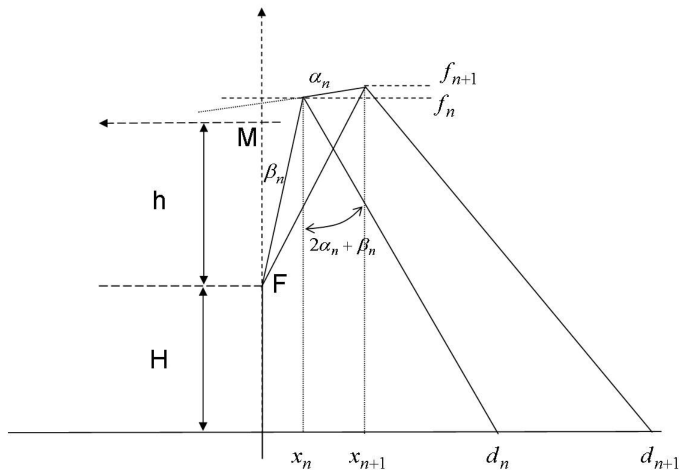

Fig. 10.10. Differential segment of a mirror with angle α_n with respect to the horizontal. Reflection of two rays on the extremes of the differential segment.

Fig. 10.10 shows the main idea of the iterative method. Assume that we compute for all pixels at a distance $p = 1, 2, \ldots, n$, from the center of the image, the angle at which a ray from pixel i leaves the camera pinhole F with respect to the vertical (this angle depends on the dimensions of the imaging chip and the distance of the chip from the camera's pinhole). Let us call those angles $\beta_1, \beta_2, \ldots, \beta_n$. Given the coordinates $p_1, p_2, \ldots, p_n$ in cm of the pixel $1, 2, \ldots, n$ on the chip (measured from the center), and the distance r of the camera pinhole from the imaging chip, then $\beta_i = \arctan(pi/r)$, for $i = 1, 2, \ldots, n$.

In what follows, we assume that we have already computed the values of x_n, and f_n, and derive from them, the angles β_i and the distance function, the next point (x_{n+1}, f_{n+1}) of the mirror's shape. The computation is started with initial values for x_1 and f_1.

Fig. 10.10 shows the ray with angle β_n reflected on a section of the mirror, and falling at a distance $x_n + d_n$ from the vertical axis of the camera. The next ray, at angle β_{n+1} falls at a distance d_{n+1}. The section of the mirror being considered is a segment starting at the point (x_n, f_n) and ending on the point (x_{n+1}, f_{n+1}). The camera pinhole is located at a distance H from the floor, and the origin of coordinates M of the mirror function f is located at a height h from the pinhole F.

First, we compute the angle α_n for the n-th mirror segment. We know also from the analysis in the previous section (Eq. 10.1) that the ray from the mirror to the point d_n makes an angle $2\alpha_n + \beta_n$ with the vertical. Therefore $d_n - x_n = (H + h + f_n)\tan(2\alpha_n + \beta_n)$, and from this we obtain

$$\alpha_n = \frac{1}{2}\left(\arctan\left(\frac{d_n - x_n}{H + h + f_n}\right) - \beta_n\right) \tag{10.2}$$

It is easy to see from the diagram in Fig. 10.10 that

$$\frac{f_{n+1} - f_n}{x_{n+1} - x_n} = \tan(\alpha_n)$$

It is also clear that

$$\frac{x_{n+1}}{h + f_{n+1}} = \tan(\beta_{n+1})$$

These are two equations with two unknowns (x_{n+1} and f_{n+1}). They provide us with an iterative method for obtaining the next point (x_{n+1}, f_{n+1}) given the previous point (x_n, f_n) and the previously computed angle α_n.

Solving the system of equations we obtain

$$f_{n+1} = \frac{f_n + \tan(\alpha_n)(h - x_n)}{1 - \tan(\alpha_n)\tan(\beta_{n+1})} \tag{10.3}$$

and

$$x_{n+1} = (h + f_{n+1})\tan(\beta_{n+1}) \tag{10.4}$$

Therefore, we can start from an initial value for (x_1, f_1) and proceed with the iterative computation procedure defined by Eqs. 10.2, 10.3, and 10.4, for obtaining the shape of the mirror for a given monotone distance function d.

As a proof of concept, we computed the distance function for a conic mirror, and then recomputed the mirror shape which produces that distance function. We obtained a conic mirror again. Fig. 10.11 shows the result and the position of the virtual viewpoint for a camera 50 cm above the ground, and a mirror 20 cm above the camera's pinhole.

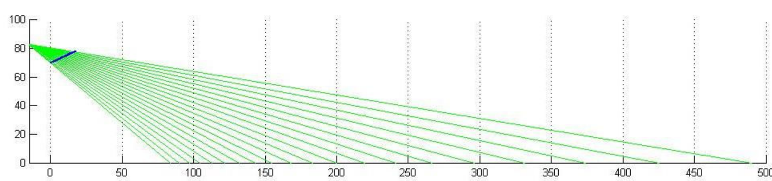

Fig. 10.11. Projection of rays from the floor on one side of a conic mirror, and continuation of the ray towards a "virtual" viewpoint.

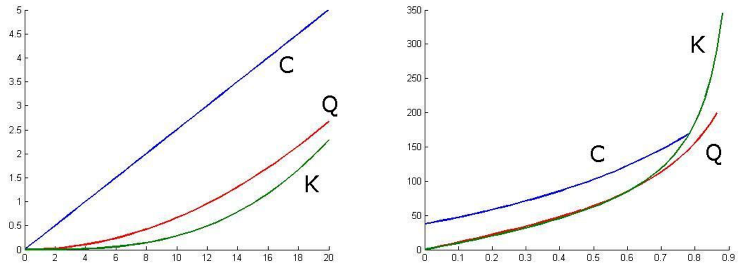

Fig. 10.12. Distance function for several mirror shapes. On the left, a slice of a conic, a quadratic, and a cubic mirror. On the right, the distance functions. The distance function for the cubic mirror grows faster at the end. The distance function for the conic mirror does not start at the origin.

10.6 Why conics are special

In this section we show why the conics of revolution are the only mirrors with a single virtual viewpoint. Remember: a camera with its pinhole at the virtual viewpoint would capture the same image as a camera with its pinhole at the point where reflected rays converge.

For the proof, refer to Fig. 10.13. The pinhole of the camera is located at F. A virtual viewpoint is located at F'. Assume that all rays which go through F, after reflection on the mirror, go through F' when prolonged. Let us call α the angle of the reflected ray with the horizontal, θ the angle of the incoming ray with the horizontal too, and β the inclination of the mirror with respect to the vertical. The angles formed by the incoming and reflected rays on the mirror, with respect to the normal to the mirror, must be equal. That is: $\alpha + \beta = \theta - \beta$. Therefore

$$2\beta = \theta - \alpha$$

We do not have the angle β directly, but the derivative of the mirror shape function $f(x)$ is equal to $\tan(\beta)$. Therefore, we are interested in an expression for $\tan(\beta)$. We know that

$$\tan(2\beta) = \tan(\theta - \alpha) = \frac{\tan(\theta) - \tan(\alpha)}{1 + \tan(\alpha)\tan(\theta)}$$

and

$$\tan(2\beta) = \frac{2\tan(\beta)}{1 - \tan^2(\beta)}$$

Therefore

$$\frac{2\tan(\beta)}{1 - \tan^2(\beta)} = \frac{\tan(\theta) - \tan(\alpha)}{1 + \tan(\alpha)\tan(\beta)}$$

From this, we obtain the differential equation

$$\frac{2f'(x)}{1-f'^2(x)} = \frac{\tan(\theta)-\tan(\alpha)}{1+\tan(\alpha)\tan(\theta)}$$

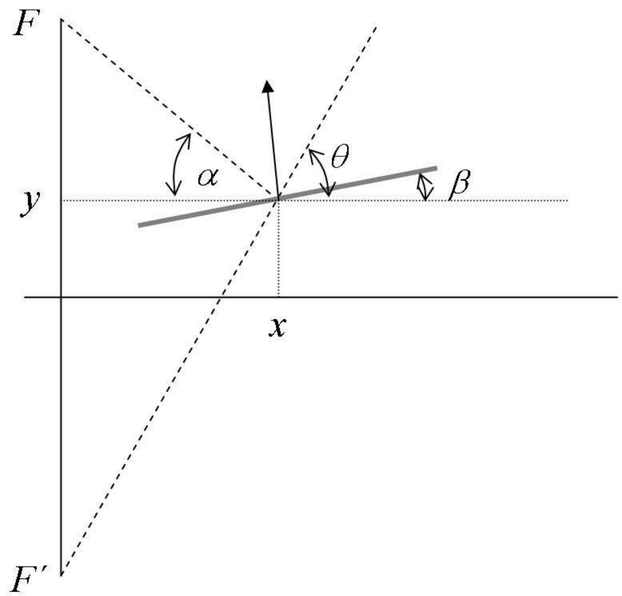

Fig. 10.13. Reflection of a ray on a mirror. The reflection goes through the pinhole F. The virtual viewpoint F' is located at the continuation of the original ray, and at $x = 0$, for a mirror of revolution.

Now, we only need expressions for $\tan(\theta)$ and $\tan(\alpha)$. From the diagram $\tan(\alpha) = (c - y)/x$ and $\tan(\theta) = (c + y)/x$. Therefore

$$\frac{2f'(x)}{1-f'(x)^2} = \frac{(c+y)/x-(c-y)/x}{1+(c+y)(c-y)/x^2}$$

or equivalently

$$\frac{f'(x)}{1-f'(x)^2} = \frac{xy}{x^2+(c+y)(c-y)}$$

(10.5)

This leads to a quadratic equation for $f'(x)$. Solving the quadratic

$$xyf'(x)^2 + (x^2 - y^2 + c^2)f'(x) - xy = 0 \qquad (10.6)$$

we obtain a first order differential equation.

The only solutions to the differential equation 10.6 are ellipses and hyperbolas (also the degenerate case of planar mirrors). The parabola is a limiting case when the pinhole of the camera is located at infinity (for a convex paraboloid mirror). See [?] for the details of how to solve the differential equation.

The same result was proved in 1990 by [?] in a more general setting, that is, of mirrors in n-dimensional space. Therefore, the conic sections, that is, the locus of all solutions to quadratic equations, are the only mirrors which satisfy the "single virtual viewpoint constraint.

However, as we saw in the previous section, it is possible to build a mirror which produces any desired monotonic distance function. Those mirrors do not necessarily have a single virtual viewpoint, but do we need it? The answer has to do with the theoretical analysis of the distortions and focusing aberrations of mirrors. Having a single viewpoint simplifies the mathematics and the analysis of aberrations. Multiple viewpoints can only be handled with numerical analysis.

10.7 An omnidirectional mirror for an autonomous car

The techniques described in this chapter can be extended to deal with mirrors other than mirrors of revolution. We have developed mirrors for a computer vision system for an autonomous car with 180 degrees of lateral vision. Fig. 10.7 shows the image obtained in a camera looking back on a mirror which is reflecting the streets in front of the car. In the sequence of images the car is approaching an intersection (the streets meet at 90 degrees). When the car drives across the intersection, the view to the sides covers 180 degrees and is quite warped, while the vertical view covers no more than 90 degrees. Along the vertical slices, the mirror images look like obtained with a planar mirror, specially towards the center of the image. On the sides, the mirror images resemble those of omnidirectional mirrors.

10.8 Catadioptric eyes in nature

Nature is the best engineer, so it should not come as a surprise that catadioptric eyes have been designed by evolution in different types of animals. Such eyes consist of a reflector which concentrates and focuses light on individual light sensors. It is also remarkable that pinhole eyes exist after all. A lens or a mirror offer the possibility of concentrating more light on the light receptors, increasing the contrast of the image. A pinhole eye is simpler and works with much less optic machinery. In this chapter we have used the pinhole camera

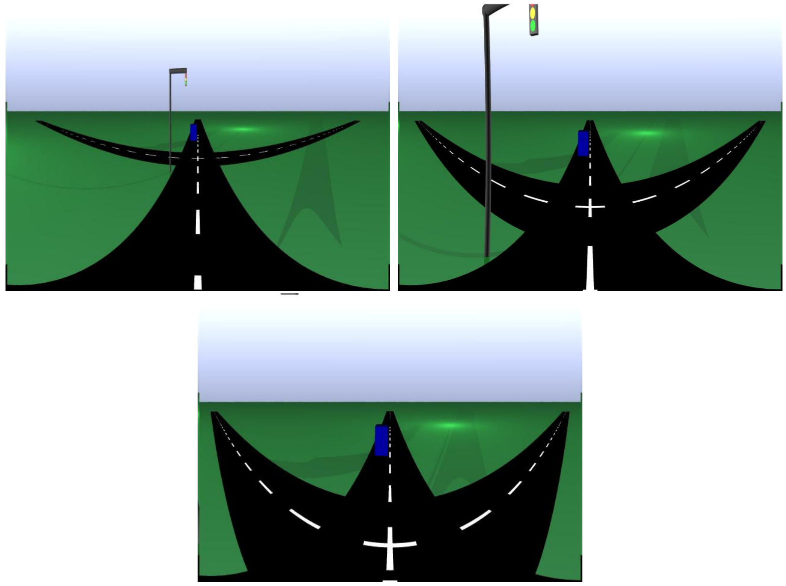

as an ideal model mainly because of its simplicity. Fig. 10.8 shows a diagram of the Nautilus eye, a cephalopod mollusc. The pinhole eye does not collapse since it is filled with sea water. The Nautilus has probably a blurred vision, which is nevertheless useful in its marine environment.

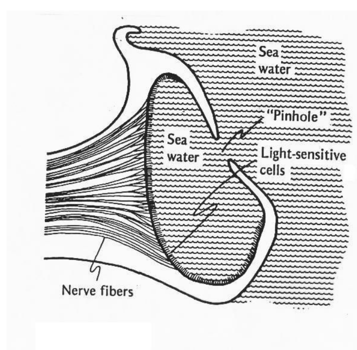

Fig. 10.14. The pinhole eye of the Nautilus, a prehistoric creature.

Our first example of a catadioptric eye is the scallop eye. Scallops do not image the environment with anything near the resolution of the vertebrate eye, but they can sense light and shadows, and since the scallop eyes are arranged along the border of the clamshell, 60 to 100 of them, they can be "triggered" sequentially by a shadow moving sequentially in front of the shell. Scallops can also probably image objects at a distance, because they close their shell when divers swim by. Fig. 10.15 shows the position of the scallop eyes along the rim of their shell.

Fig. 10.15. The eyes of the scallop, visible as dark spots on the rim of the shell.

In the diagram of Fig. 10.16 we can see a photograph of the scallop eye. The retina is clearly visible right behind the lens. It seems that the problem with the scallop lens is that it is made of a material with low refractive index, so that the retina would have to be located very far behind the lens in order for the image to be focused. The alternative is to focus the light back with the help of a spherical mirror, so that the focused image forms just behind the lens and on the light receptors of the retina. In this arrangement the light goes through the receptors unfocused and comes back focused, so that scallop eyes must form images with low contrast,as if looking through a fog. The lens of the scallop eye has a peculiar form which apparently has the function of correcting the spherical aberration of the mirror. The mirror itself has a silvery color and is made of similar material as that contained in the shiny inner surface of shells (and in pearls).

The example of the scallop is interesting because in this case evolution traded-off perfecting the lens against perfecting a mirror reflector. Lens and mirrors are equivalent, they can focus light in the same way.

That mother nature also found out the advantage of using parabolic light collectors much earlier than Newton, is evidenced by the reflecting eyes of the Gigantocypris, small deep-sea crustaceans which in that environment try to collect as much light as possible. Their eyes look like car light reflectors, but instead of a light bulb, the reflectors have a bulb covered with light sensors! Fig. 10.17 shows a photograph of the two eyes, and a diagram of its parabolic shape. A biologist wrote about these eyes: "The paired eyes have huge metallic-looking reflectors behind them, making them appear like the

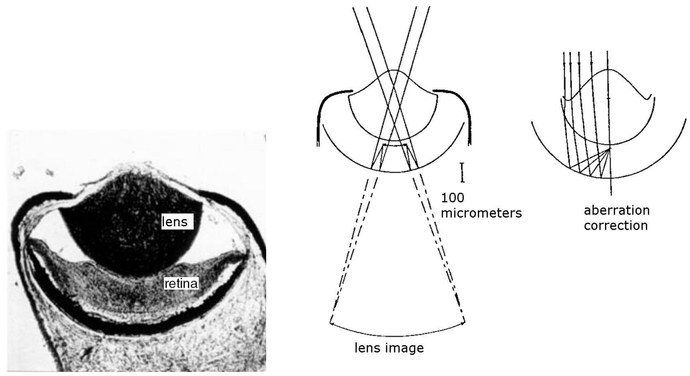

Fig. 10.16. The scallop eye in cross section (left). The back reflecting surface focuses light directly on the light receptors (right). The lens helps to correct the spherical aberration of the mirror.

headlamps of a large car; they look out through glass-like windows in the otherwise orange carapace and no doubt these concave mirrors behind serve instead of a lens in front (Hardy 1956)."

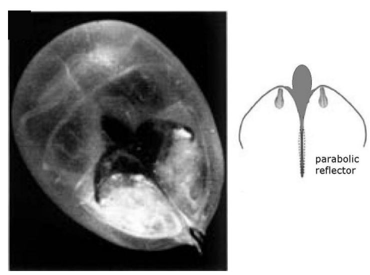

Fig. 10.17. Photograph of the Gigantocypris eyes, and diagram of the parabolic reflectors. The bulbs at the focus of the paraboloid are covered with photoreceptors.

However, the prize for the most remarkable catadioptric eye design must really go to shrimps and lobsters. Their eyes resemble the composite eyes of flies or bees. However, there is an important difference: light falling on each segment of the composite eye is not focused with a lens, but is just reflected down on the retina. The reflection occurs on the sides of the elementary eyes, which have a square cross section and are prims mirrors which reflect light on their four sides. Fig. 10.18 shows a photograph of the composite eye, of the square cross section of the individual prisms. The angle of the prisms walls are oriented so as to reflect and focus light on the retina. However, when light arrives at a slight skewed angle with respect to the plane of the diagram, its direction must be corrected, and this is done thorugh a double reflection on the prisms.

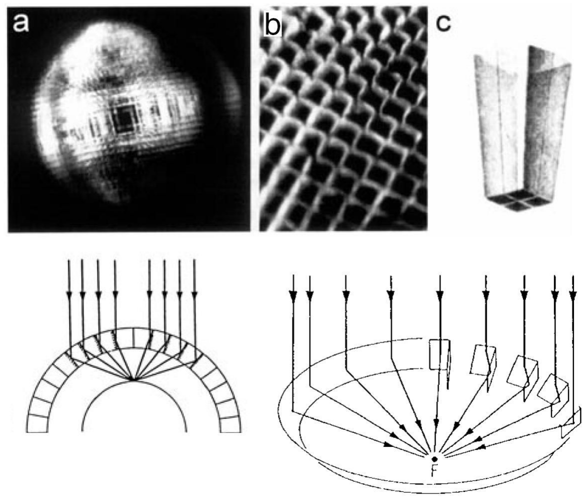

Fig. 10.18. The eyes of shrimps and lobsters are composite eyes. Individual eyes are prisms with reflecting walls. The walls focus the light on the retina. The corner-mirror construction can collect more light on the same retinal spot.

Fig. 10.19 shows how a corner mirror works. Any ray of light with arrives to one side of the mirror, skewed, is reflected on the other side and is sent back in the same angle as it arrived. The angle changes direction twice, rotating twice by 90 degrees. The effect is just like from a planar mirror, but we do not have to worry about the angle of light ray with the normal to the mirror surface. It is like an "omnidirectional" planar mirror. Such a corner mirror was left by American astronauts on the moon. Aiming a laser from the earth, its reflection could be detected on earth easily, since the laser does not have to be in perfect alignment with the mirror. An accurate measurement of the distance of the earth to the moon was obtained with this laser reflection.

Finally, even vertebrates and fish have developed catadioptric solutions to the problem of gathering the maximum amount of light in the retina, specially at night. The surface of the retina in cats and sharks is covered by reflecting cells which send back the light which went by the photoreceptors without being absorbed. These photons get a second chance to be detected.

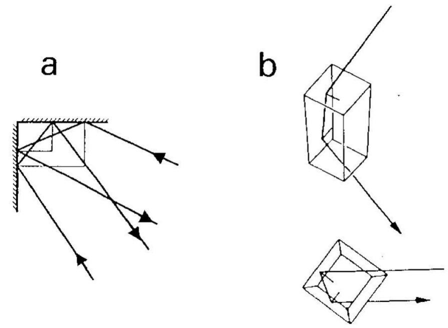

Fig. 10.19. Operation of a corner mirror. Any ray falling on its surface is reflected back in the same direction.

The reflected photons are narrowly focused so that if they are not detected again, they will leave the eye through the lens, so that its backscattering in the eye does not diminish the image contrast. That's why the eyes of cats look bright at night. The eyes of sharks, spidersm, and many other night active animals reflect light on the "tapetum", the quasi-mirror on the back of the retina. Fig. 10.20 shows a diagram of the angle at which the shark's tapetum reflects light back. The arrangement is catadioptric.

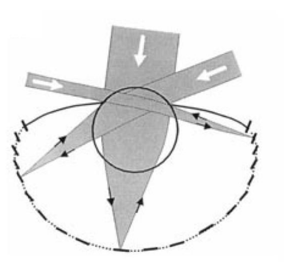

Fig. 10.20. Arrangement of the reflecting surfaces in the tapetum of the shark's eye. reflected light is focused and can be detected or can leave the eye through the lens.

10.9 Final comment

Optical lenses work because they can deflect light and make it converge where it can form an image. Mirrors can do exactly the same trick, they change the path of a ray of light, but reflecting it back instead of letting it go through. Other than that, they share the image formation equation with lenses which can focus light:

$$\frac{1}{d} + \frac{1}{s} = \frac{1}{f}$$

where d is the distance of the object to the lens/mirror, s the distance of the virtual image, and f the distance of the focus. Parabolic mirrors image objects at infinity ($d = \infty$) at the focus ($s = f$). The parabolic mirror is therefore the equivalent of the traditional collecting lens, as realized by Isaac Newton when he invented the telescope which bears its name. It is most remarkable that during the evolutionary history of the earth, nature has given a try to almost all camera designs which engineers have developed, from pinhole to catadioptric cameras.