Reading for Chapters 4, 5, & 6

Camera Lenses & Focusing on the Object

The lens of a camera must form a real image of the object that is being photographed – the rays of light from a point on the object must actually strike the film at a single point in order to trigger the chemical reaction that records the image. Although the lens of a camera is a complex multi-element design to minimize aberrations, its general properties are equivalent to a positive lens of comparable focal length. We can therefore understand the performance of a camera lens using the principles we have outlined in previous topics.

The ray-tracing rules we have established define a unique relationship between the focal length of the lens and the positions of the object and the image. In other words, when a lens of some focal length is placed at a given distance from the film plane then this distance and the focal length of the lens determine the distance to an object that is in sharp focus. In principle, objects at any other distance are not clearly focused. In order to bring an object at some other distance into clear focus, it is necessary either to move the lens or to change its focal length. Either of these changes is possible, and both are supported by more expensive cameras. Simpler and cheaper cameras usually have the ability to move the lens but not to change its focal length, while neither the focal length nor the lens position can be changed on the very simplest cameras.

In order to get a feeling for how much the lens must be moved to bring different objects into focus, the following table is computed using the algebraic form of the ray-tracing equations for a single positive lens. (See Appendix E, page 418 in the textbook, or just trust me.)

Distance of the object in front of the lens in units of the focal length	Position of the image behind the lens in units of the focal length
10	1.11
100	1.01
1,000	1.001

10,000	1.0001

For example, suppose we are using a lens whose focal length is 50 mm, which is a typical value for the lens on a 35 mm camera. The first entry in the table above says that an object 500 mm (just under 20 inches) in front of the lens (10x50) is imaged 55.5 mm behind the lens (1.11x50). If this object is to be imaged on the film plane, the lens must be located 55.5 mm in front of the film. The next entry says that an object that is 5000 mm (5 m) in front of the lens (about 16 feet) will be imaged 50.5 mm behind the lens. The lens must therefore be moved closer to the film plane in order to focus on this more distant object. The same thing is true for the other entries – the lens must be moved closer to the film plane as the distance to the object increases. Note that the changes become smaller and smaller as the distance to the object increases. In order to bring objects from 20 inches to infinity into clear focus, this particular lens must be able to move about 5.5 mm – from 55.5 mm in front of the film to essentially 50 mm in front of the film plane.

The lens motion required to bring these objects into focus is surprisingly small, and it is not difficult to design a camera that allows the lenses to move by these amounts relative to the film plane. On the other hand, the position of the image changes more and more rapidly as the objects gets closer to the lens, and most lenses cannot move far enough to focus on close-in objects for that reason. For example, an object that is 5 focal lengths in front of the lens has an image that is 1.25 focal lengths behind the lens, while an object that is 2 focal lengths in front of the lens has an image that is 2 focal lengths behind it. These motions cannot be accommodated by standard lens mounts, although there are special-purpose mounts that can do this.

Depth of Field

In the previous discussion (Topic 29. Camera Lenses, Part 1. Focusing on the Object), we showed how the position of the lens had to be moved relative to the film on the image plane in order to bring objects at different distances in front of the lens into sharp focus. However, simply changing the position of the lens is not adequate for real-life situations, since anything more complicated than the very simplest of objects will have components that are at different distances from the lens. In principle, all of these different components cannot be brought into sharp focus on the film plane simultaneously. The following figure illustrates the problem.

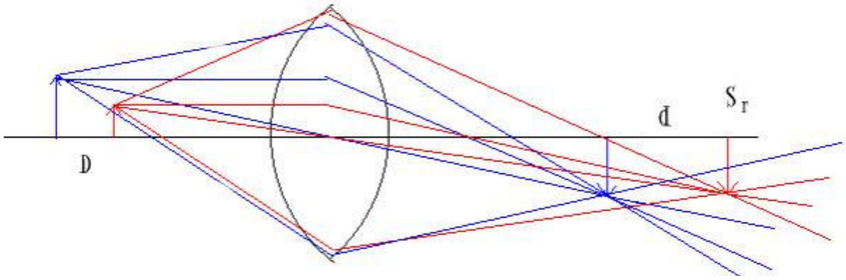

If the lens is adjusted so that the image of the red object is in sharp focus (that is, the lens is adjusted so that the film plane is at distance S_r behind the lens) the rays from the blue object strike the image plane in an extended region rather than at a single point. Depending on the goal, the lens can be moved to bring either of the two objects into sharp focus (but not both simultaneously) or it can be adjusted to a point midway between the two images so that both objects are about equally blurred.

The magnitude of this problem is determined by two independent factors. The first is the angle at which rays from the lens converge to the point of sharpest focus. This angle determines how much blurring will result from a given difference in the positions of the images from the different parts of the object. (The rays diverge on the far side of the point of sharpest focus at the same angle, so that the blurring is equally large at some point in front of the actual focus or at a point an equal distance behind it.) This angle is illustrated in the following figure.

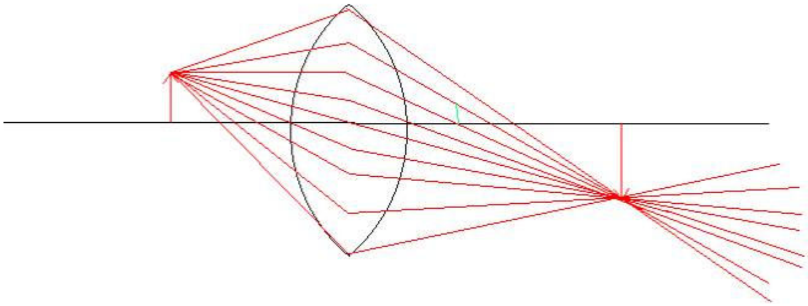

The size of this angle is determined by the ratio of the diameter of the lens to the distance between the lens and the image. Since objects are usually many focal lengths in front of the lens, the distance between the lens and the image is approximately equal to the focal length of the lens (see the previous topic). The ratio of the diameter of the lens to its focal length has appeared before as the parameter that governs the applicability of the paraxial ray approximation, and its inverse is the f-number.

The magnitude of this angle is clearly determined by the extreme rays – the rays that strike the lens at its edges, so that the problem can be reduced either by increasing the focal length of the lens (leaving everything else unchanged) or by inserting a stop or diaphragm that blocks the edges of the lens and prevents these extreme rays from reaching the image. Since the f-number is the ratio of the focal length to the diameter, reducing the effective diameter of the lens leaving the focal length unchanged increases the effective f-number of the lens. (Increasing the focal length leaving the diameter unchanged obviously has the same effect.) This increase in f-number reduces the angle at which the rays converge to (and diverge from) the point of sharpest focus and reduces the blurring for any given configuration.

The second factor that governs the magnitude of this blurring is the relationship between D, the distance between the two points in the object space, and d, the distance between the corresponding images. Decreasing d, keeping everything else the same, will decrease the size of the problem because the same angular divergence of the various rays produce a smaller divergence of the rays on the image plane.

Therefore: the blurring due to depth of field can be reduced by using a larger f-number and leaving everything else unchanged.

For a given distance D between two points on the object, the corresponding distance between the two images decreases as the objects move further away from the lens. This point is illustrated in the following table, which gives the variation in the ratio d/D as the object moves further away from the lens.

Distance of the object from the lens	d/D
5	0.06
10	0.01
100	0.0001
1000	10^{-6}

For example, if an object is 5 m (about 16 feet) in front of a lens whose focal length is 50 mm (a distance of 100 focal lengths), the ratio of d/D is 0.0001 – the images of two objects are spaced only 0.0001 times the distance between the two objects themselves. As you can see from the table, the fraction d/D decreases as the *square* of the distance from the object to the lens.

Therefore: the blurring due to depth of field will be reduced as the object moves further away from the camera, leaving everything else unchanged.

The final factor in estimating the depth of field of a real lens is the fact that no lens can focus the rays from a single point on an object to a single image point. Even the best lens produces a somewhat blurry image because of residual uncorrected aberrations, because of scattering of light in the lens, and because of dispersion (the fact that the ratio of the wavelength of the light to the size of the aperture is not exactly zero so that light always has some residual wave-like character). The practical depth of field of a lens is then set by the balance between the imperfections in the size of the image point due to all of these problems and the additional blurring caused by depth of field. That is, the practical depth of field of a lens is set as the point at which the blurring due to the depth of the real object is not greater than the blurring that is inherent in the performance of the lens itself when it is imaging a perfectly flat object. For any given lens and object, the two factors above play important roles – the depth of field can be increased by increasing the f-number of the lens or by moving the camera further away from the object. The following table shows the actual depth of field for a typical lens. That is, the table shows the range of object distances over which the image is essentially in "perfect" focus because the blurring due to the differences in the object distances is not greater than the blurring due to the other imperfections in the optical system. This lens has a focal length of 50 mm and a maximum aperture of f/1.7 (In other words, the maximum aperture is 50/1.7=29.4 mm). These parameters are typical of the lens that is used on an ordinary 35 mm camera.

Although the details of this table will vary from lens to lens, the general features are universal. In particular, when a lens is used at a small aperture opening corresponding to f/11 or f/16 and is focused on an object 10 m (about 30 feet) away, almost all objects (except for those that are quite close to the camera) will be in focus. This is how most cheap cameras are configured – they operate at a fixed lens aperture of f/11 or f/16 with a fixed focus lens adjusted so that objects about 5-10 m away are in sharpest focus. The depth of field at this small aperture is so large that almost everything else is acceptably in focus without the need for any adjustments.

The Effect of Focal Length

The focal length of the lens used in a camera affects two aspects of the image: the field of view and the magnification.

The field of view is the size of the object that can be contained on the film and the magnification is the ratio of the size of the image to the size of the object that produced it. These two quantities are related to each other – increasing the magnification while leaving the size of the image unchanged inevitably decreases the field of view (since the increase in the magnification implies that the same sized image will be produced by a smaller object), and vice-versa.

The field of view is usually expressed as an angle. There are a number of different ways of estimating this quantity, but all of them yield about the same value. Since we are going to use the approximation that the camera lens is a simple single thin lens, the calculation will not be exactly correct anyway. As with previous calculations, the agreement between this simple model and a real-lens is surprisingly good.

If the object is not too close to the lens, then the image is always formed approximately at the focal distance back from the lens. The size of the image on the image plane is usually limited by the size of the film, and the angular size of the field of view is given by those rays that enter the lens at its center (and so are not deviated) and which then strike the edges of the film as shown in the following figure:

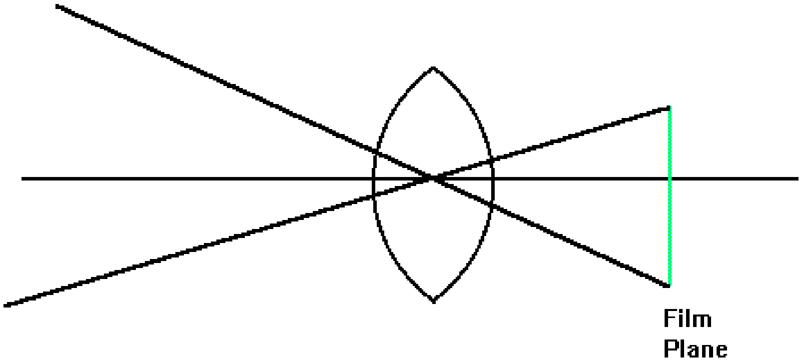

The angles on the object and image side of the lens are the same, so that the tangent of one-half of the angle of the field of view is approximately the ratio of the half-width of the film to the focal length of the lens. For example, the half-width of 35 mm film is 18 mm. The horizontal field of view of a 50 mm lens is then twice the arc-tangent of 18/50 or about 40 degrees. The film is not square, so that the field of view in the vertical direction is not the same. The half-height of the film is only 12 mm, so that the vertical field of view is twice the arc tangent of 12/50 or about 27 degrees.

(Another method for calculating the field of view would be to consider those rays that pass through the first focal point on the object side of the lens. These rays are refracted so that they emerge parallel to the axis on the image side of the lens. The angular spread of these rays provides an alternate method for estimating the field of view. This estimate is essentially the same as the estimate using the method described above unless the object is quite close to the lens. The two estimates differ in this situation because the lens must be moved away from the film plane in this situation so that the distance from the lens to the film in the first figure is somewhat greater than the focal length.) See the following figure, which compares the field of view calculated in this way with the calculation above when the object is far away.

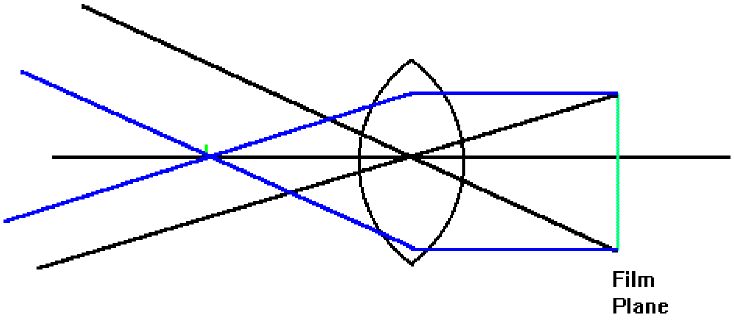

The field of view angle decreases with increasing focal length, but the relationship is through the arc tangent function and is not exactly linear (except at very small angles). The horizontal field of view of a telephoto lens whose focal length was 135 mm would be twice the arc tangent of 18/135 or about 15 degrees.

The magnification of a lens can be estimated using the same sort of analysis. The following figure shows the image produced by an object.

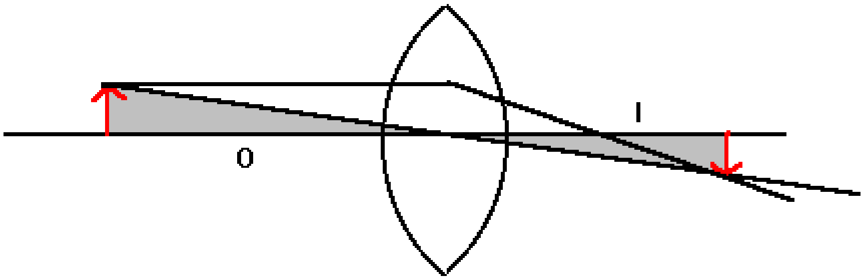

The two shaded triangles are similar to each other, so that the corresponding sides are proportional. The magnification is the ratio of the size of the image to the size of the object that produced it. Since the triangles are similar, that ratio is also equal to the ratio of the distance between the lens and the image (“I” in the figure) and the distance between the lens and the object (“O” in the figure). Since the distance “I” is approximately the focal distance, the magnification of a lens for a fixed object depends directly on the focal length and is equal to the ratio of the focal length to the distance of the object from the lens.

The magnification increases directly with the focal length when the object, the position of the camera, etc. are unchanged. Thus the same object appears larger when imaged by a camera with a telephoto lens because the focal length of that lens is larger than the focal length of a normal lens.

Although a film image is purely two-dimensional, the variation in the magnification of a lens with object distance can distort perspective – that is, how we estimate the size of something in an image by comparing it to other things in the same image. For example, if an object is sufficiently large so that the distances between different parts of the object and the lens are significantly different, then the parts that are closer to the lens will have a greater magnification than the parts that are further away. This effect can be seen when photographing a tall building from a point near its base. The bottom of the building is significantly closer to the camera than the top, is magnified by a larger amount and appears disproportionately too large as a result.

Focusing the Lens & Rangefinders

In cameras where the lens can be moved to bring the object into focus on the film plane, there is usually a mechanism that indicates when the proper position of the lens has been reached. One system that is used for providing this indication is the optical rangefinder. A number of different versions of this device have been developed, but most of them are based on the design shown in the following figure. (The optical components that are shown in the figure are usually completely independent from the lens that is used to take the picture.)

The goal is to focus the camera on the object by observing its image on the viewfinder screen.

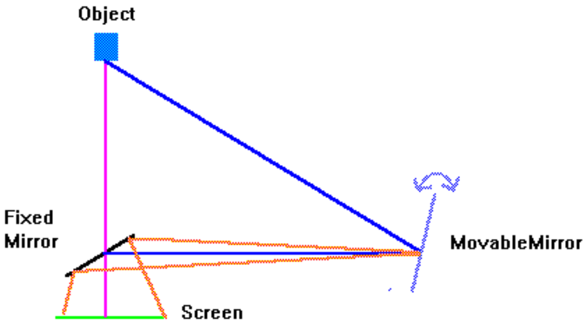

The user sees two images of the object on the screen inside of the camera. One image is produced by rays that travel along the direct path from the object to the screen as shown in purple. Rays traveling along this path strike the screen after passing through the fixed mirror, which transmits 50% of the light that is incident on it and reflects the other 50%. The second image is produced by rays traveling along the blue path. These rays are reflected twice before hitting the screen: once by a movable mirror that can rotate about an axis perpendicular to the figure and the second time by the fixed mirror. As the movable mirror is rotated, the blue image moves across the screen, and there is one position of the movable mirror for which the blue image exactly coincides with the purple image. (At other positions, the rays are reflected to other positions on the screen as shown by the orange rays, and the user sees a displaced image.) The position at which the two images overlap is determined by the geometry of the mirrors and by the distance of the object from the camera, so that the angle of the movable mirror can be calibrated to read the distance to the object directly.

The lever that rotates the movable mirror is coupled to the system that positions the lens, so that rotating the mirror until the two images overlap simultaneously adjusts the lens to focus on the object at that distance. This "coupled rangefinder" system is widely used on moderately-priced cameras.

The rangefinder is most sensitive to objects that are relatively near to the camera, and becomes less sensitive as the object moves further away because the corresponding changes in the angular position of the mirror become smaller and smaller at greater distances. This is usually not a serious problem, because the depth of field of the lens is also increasing as the object moves further away, so that the focusing becomes less critical at these distances.

Bi-Prisms

A second method for detecting when the lens has been focused on the desired object uses a pair of prisms mounted in opposite directions on the image screen. This configuration is sometimes called a biprism, although that name is also used for other configurations of two prisms that have nothing to do with focusing a camera.

The method depends on the properties of a single thin prism as shown in the following figure.

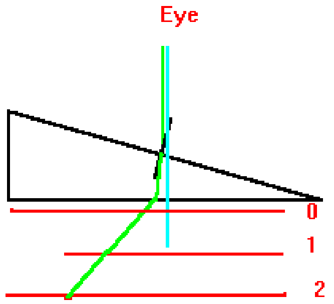

A user whose eye is looking down through the prism from above sees an object below the prism as if it were located along the straight path shown by the blue line. In fact, the ray that reaches the user's eye was refracted twice by the prism and actually followed the path shown by the green line. The angle of divergence between the blue path and the green path is a constant value – it depends only on the index of refraction of the prism and on its shape.

If the user is looking at an object located at position "0" then the user actually sees a point on the object that is to the left of the straight line path because the ray that actually enters the user's eye was refracted as shown. As the object moves away from the bottom of the prism towards positions "1" and "2" the point on the object whose rays strike the user's eye moves further and further to the left as shown in the figure. The user thinks that the object is sliding to the right, since parts of it that were further and further to the left are coming into the view point which looks to be directly under the user's eye. On the other hand, when the object moves upward so that it is directly underneath the prism, the difference between the green ray and the blue ray is very small so that the user sees an object point that really is directly on the straight line of sight. (The difference between the blue path and the green path has been exaggerated in this figure by separating the two rays for clarity.)

We now combine two identical prisms of this type with one of them slanted to the right and one of them slanted to the left. We arrange the configuration so that the top half of the object sends it rays through one of the prisms and the bottom half of the object sends its rays through the other one.

As the object moves away from the bottom plane of the prisms, the two halves of the image appear to move in opposite directions, since the prisms tilts are opposite. Conversely, as the object moves closer to the bottom plane of the prism, the two halves of the image move closer together. If the prisms are thin and no too tilted, the two halves of the image will line up exactly with each other when the object plane is exactly at the common base plane of the two prisms.

In practice, the two prisms are located on the imaging screen (this screen is where the main camera lens forms its image while the picture is being composed), and the image on this screen serves as the object for the two prisms. When the image produced by the main camera lens is located exactly on the imaging screen, the two halves of this image transmitted through the two prisms appear lined up. When the image produced by the main camera lens is not focused on the imaging screen, the two halves of the image transmitted through the two prisms are displaced in opposite directions, and the image seen by the user appears to be cut in half at its center with the upper and lower halves displaced relative to each other.

Some imaging screens contain additional very small pairs of these prisms called micro-prisms. Each pair functions in the same way – the image appears to be smooth and continuous when all of the little pieces are not displaced relative to each other and appears to be broken up into tiny disjoint parts when the image is out of focus.

Shutters

The shutter of a camera has only one function – to control the amount of the time that the film is exposed to the scene. This is necessary for two reasons: (1) if the object is moving or the scene is changing in some other way, then the shutter can be used to limit the exposure to a time interval that is short enough to “freeze” the motion, and (2) the exposure time controls the amount of light that strikes the film to be sure that the film is properly exposed. See the next topic (Topic 35. Film speeds and exposure calculations) for more details about this.

Early cameras used leaf shutters that were located in front of the lens or between two elements of a multi-element lens. As the shutter opened, a circular hole was formed that increased in diameter until the shutter was fully opened and then decreased in diameter until the shutter was fully closed again. Since the shutter was not in the image plane, the opening and closing of the shutter affected the entire image in the same way. That is, the opening of the shutter gradually allowed more and more off-axis rays to reach the film plane, but all of the film plane was exposed to the incident light at all shutter positions. The shutter function was often combined with the lens opening adjustment into a single device whose opening speed set the exposure time and whose maximum diameter set the f/number. Note that the shutter has to move faster and faster at shorter exposure times, and this requirement set a mechanical limit on the shortest exposure times that could be obtained.

Although shutters of this type have some theoretical advantages, they are difficult to use in cameras with removable lenses. Since the shutter is an integral part of the lens, both the shutter and its trigger mechanism must be built into all of the different lenses that will be used. They are also difficult to use in single lens reflex cameras, since those cameras use a through-the-lens viewfinder to compose the image. The viewfinder obviously couldn’t work if the shutter was closed. All single lens reflex cameras and most other cameras with removable lenses therefore use a focal-plane shutter.

A focal plane shutter is located just in front of the film. It consists of two pieces of opaque material called the "curtain." When the shutter is to open, the first piece of material slides horizontally across the film plane, gradually exposing the film from one side to the other (or sometimes from top to bottom). After a preset time, the second piece of material starts to move across the film plane gradually closing the opening again. At relatively slow shutter speeds (typically about 1/60 s and longer), the second piece of the curtain does not start moving until after the first piece has reached the other end of the film. The entire image area is therefore exposed for some period of time before the second curtain starts to close the opening. At higher shutter speeds, the second curtain starts moving to close the opening before the first curtain has finished opening it. The result is a moving "slot" which exposes different parts of the film at slightly different times. In most cameras, the curtains move at the same speed for all exposure times, and the width of the slot is varied to set the required exposure time. The shortest exposure time is then set by the speed of the curtains and the minimum width of the slot. Most cameras of this type can support exposure time down to about 1/1000 s.

Focal plane shutters cannot be used at all speeds with flash bulbs or electronic flash units. This is because the duration of the flash is very short – typically about 0.001 s. If the shutter is operating in its "moving slot" mode, then only a small piece of the film will receive the light of the flash unit – the piece of film that is being exposed when the flash is fired. In order for the entire image to receive the light from the flash unit, the shutter must be set to a slow enough speed so that it is running in its fully open mode. That is, the mode in which the second curtain does not start closing the opening before the first curtain has finished opening it. The flash fires in this instant after the first curtain is fully open and before the second one starts to close. The maximum speed of the shutter in which this is true varies from one design to another but is typically about 1/60 s, and this speed is usually indicated on the shutter speed selector in some way.

The optics of the eye

1. The lens

Although the lenses of the eye are not simple spherical lenses, they are thicker in the center than at the edges. Therefore, they are positive and converging, and form a real, inverted image on the retina – the optically sensitive screen that is located at the back of the eye. The image on the retina is formed by two lenses – the cornea, which is a converging lens whose focal length is fixed, and the eye-lens, which is a positive lens whose focal length can be varied by changing its shape. The focal length of the combination varies among different people, but is typically about 20 mm.

The optical strength of the cornea is much greater than the strength of the eye-lens, and most of the focusing of the lens combination in humans is due to this part of the optical system. The focal length of the cornea depends on its internal index of refraction and on the fact that the material in front of it is air. The focal length is therefore very different when the cornea is in water, since the refraction at the front surface of the cornea is much less in this situation. (This is why it is easier to see underwater when you are wearing goggles, since the goggles produce an air space in front of the cornea.) Fish and other animals that see in the water therefore have eyes whose lenses are much more curved and bulging than human eyes. In all cases, the index of refraction of the lens material is much closer to the index of refraction of water, so that a lens of the same shape has a much smaller optical strength. Lenses that must operate in water are therefore much more curved than a lens of equivalent focal length designed to work in air.

Neither surface of the eye-lens is exposed to air, and its focal length is determined by its index of refraction and by the indices of refraction of the materials that surround it. The index of refraction of the eye-lens is not homogeneous but is smallest at the surface and increases somewhat with depth.

The eye focuses on objects at different distances by changing the focal length of the eye lens. This change is realized by changing its shape – decreasing its radius of curvature (and therefore shortening its focal length) as the object comes closer to the eye. The range over which this works varies among different people, but is typically from infinity to the *near point*, which is about 25 cm (10 inches) for most people. This variation in the focal length of the eye-lens is called *accommodation*, and this ability to vary the focal length of the eye-lens decreases slowly with age. Since the eye is focused on distant objects when the muscles are relaxed, the decreases in the strength of the muscles and in the flexibility of the lens with age generally means that people lose the ability to focus on nearby objects as they grow older. They therefore often tend to become *far-sighted* as a result.

2. The iris and the pupil

The iris is a circular diaphragm, and its color is what we call the color of a person's eyes. The color of the iris has no role in the optical characteristics of the eye. The central opening of the iris is called the pupil. The diameter of the opening changes in response to changes in the ambient light level. It varies by about a factor of 4 in people: from about 2 mm in diameter in bright light to about 8 mm in diameter in the dark. The variation is much larger among nocturnal animals, such as cats. The f/number of the human eye therefore varies from about f/2.5 (20/8) in dim light to about f/10 (20/2) in bright light. The variation in ambient light level is much larger than this factor of about 16 in light-gathering ability, and the eye uses other mechanisms to adjust to changes in the intensity of the ambient light. However, the increase in f/number as the ambient light increases is important in increasing the depth of field of the optical system, and it is one of the reasons that it is easier to read and do other tasks that require precise focusing when the light is brighter. As we have discussed previously, many other optical aberrations are also smaller at larger f/numbers, since the paraxial-ray approximation is more accurate in this case.

The largest residual aberration is usually chromatic aberration, which is due to the variation in the index of refraction of the eye lenses with wavelength. When the ambient intensity is high, the use of the cones to process the image tends to minimize this problem by selectively ignoring the blue and violet portions of the image, where the chromatic aberration is the largest, and concentrating instead on the yellow and green portions, which are at the center of the visible spectrum. However, objects illuminated with light containing lots of violet and yellow (or ultra-violet and blue) can appear fuzzy because the images in the two different colors cannot be brought to a sharp focus simultaneously.

3. The retina

The retina is the optically-sensitive portion of the eye. It is the screen on which images are formed, and is connected to the brain via the optic nerve. The retina is composed of about 130 million sensors called rods and cones. The rods are more sensitive to low light levels, whereas the cones operate at higher light levels and are also responsible for color vision.

The cones are concentrated at the fovea, a point near the optical axis of the eye. Since this point is close to the optical axis of the eye, the various lens aberrations are smallest here, and the image quality is therefore best at this point. In addition, the density of cones is very high in this region, and the resolution of the eye is therefore greatest when the image falls on this area.

The rods predominate at other places on the retina, especially at the edges, which have mostly rods and few cones. The rods are most sensitive to low light levels, but cannot distinguish color. Therefore perception of color gets poorer as the ambient light level decreases. Since there are more rods on the periphery of your retina than at the center, peripheral vision tends to be more sensitive at lower levels of ambient light. The image quality is poorer, however, both because of the lower density of rods in these regions and because the aberrations are relatively larger this far from the optical axis.

The sensitivity of both the rods and the cones can be adjusted in response to changes in ambient light intensity, and this is one of the ways the eye adjusts its sensitivity. These changes are much larger than the change in effective f/number produced by changes in the size of the pupil. This effect is called *adaptation*. It is a relatively slow process – the change in the sensitivity of the retina to a sharp change in ambient light intensity typically takes 15 or 20 minutes.

In addition to being more sensitive at lower levels of ambient light, the rods are also more sensitive to blue light than to red light. If an object is illuminated with only red light, the rods will tend to see the scene as “dark” and will increase their sensitivity as a result, while the cones will tend to see the scene as lighter and will become active because of that. This trick is often used in illuminating the instrument panel of a car at night and in movie theatres. In both cases, the goal is to provide the best features of night vision using the rods and day vision using the cones simultaneously.

Correcting common vision defects

In order to understand common vision defects, it is helpful to understand how the “normal” eye works. The muscles that control the focal length of the eye-lens are fully relaxed when the eye is focused on a very distant object. The muscles contract and shorten the focal length of the eye-lens to bring nearer objects into focus. The two limits of this range are called the far-point, which is the distance of an object that is in focus when the muscles are relaxed, and the near-point, which is the distance at which an object is in focus when the muscles have shortened the focal length of the eye-lens as much as possible.

The distance between the eye-lens and the retina is about 20 mm. When an object is very far away from the eye, the image is located essentially at the focal point. Therefore the focal length of the cornea and the eye-lens should be about 20 mm when the muscles of the eye are relaxed. Since the strength of a lens is the reciprocal of its focal length in meters, the strength of the cornea and the eye-lens in this situation is about 1/0.020 = 50 diopters.

When an object is located at the near-point (the closest point at which an object can be brought into clear focus on the retina), the focal length of the cornea and the eye-lens must be changed so that the image is formed on the retina, which is still 20 mm away. If we take 25 cm as a typical value for the near-point, the algebraic form of the ray-tracing rules shows that the focal length of the cornea and the eye-lens must now be about 18.52 mm. The strength of the cornea and the eye-lens must now be about 1/0.01852 = 54 diopters. Thus the muscles of the eye must provide an accommodation range of 4 diopters. (This required range is actually independent of the size of the eye or the focal length of the eye lens when it is focused on a very distant object. It depends only on the value for the near-point.)

1. myopia or “near-sightedness”

In this condition, the strength of the cornea and the eye-lens combination is too great even when the muscles of the eye are completely relaxed. In other words, the focal length of the cornea and the eye-lens combination is always less than the distance to the retina, so that images of distant objects are formed in front of the retina even when the eye is totally relaxed. However, objects that are closer can be brought into focus. In many situations the focal length of the cornea and the eye-lens is so short that objects closer than the conventional near point of 25 cm can be brought into focus, which is why this condition is called “near-sightedness.” People with this condition can thread needles and read books, but they often cannot read distant street signs.

Since the problem is that the strength of the eye-lens and the cornea combination is too great, the solution is to provide eye glasses (or contact lenses) with negative lenses. The negative lens weakens the strength of the cornea and the eye lens just enough so that the resulting focal length when the eye muscles are relaxed matches the distance back to the retina so that distant images are now in focus. Since the eye glass lenses are negative lenses, they are thinner in the middle than at the edges. It is easy to identify this kind of eye glass lenses, since acting by themselves they do not form a real image of an object at any distance.

2. hyperopia or “far-sightedness”

In this condition, the strength of the cornea and the eye-lens combination is too weak when the eye muscles are totally relaxed, so that the image of a distant object is formed behind the retina. It is usually possible to bring distant objects into clear focus by using the eye muscles to shorten the focal length of the eye-lens, but this accommodation is usually not sufficient to bring nearby objects into clear focus. Therefore people with this condition can focus on distant objects just fine, but cannot focus on an object at the conventional near point. People with this condition may not be able to focus on a book unless it is quite far away from the eyes, and they usually cannot thread a needle or do other close-up work.

The solution is the opposite of the previous case. The eye glass lenses are positive lenses which strengthen the cornea and eye lens just enough so that the resulting focal length when the eye muscles are relaxed matches the distance back to the retina. Since these eye glass lenses are positive lenses, they are thicker in the middle than at the edges. Acting by themselves they can form a real image of a distant object at their focal point, so that it is easy to distinguish between eye glasses for myopia and hyperopia, since the latter will form a real image of a distant object while the former will not.

3. astigmatism

Astigmatism is the situation in which the focal length for the cornea and the eye-lens for an object oriented in some direction is not the same as for another object located in a perpendicular direction. In other words, the eye cannot bring the vertical and horizontal lines in a “+” symbol in sharp focus at the same time. (The axes of differing focal lengths need not be exactly horizontal and vertical.) This condition is often combined with one of the two previous conditions, but it need not be – it can exist even if the cornea and the eye lens have the proper average focal length and range of accommodation.

Since the problem is that the cornea or the eye lens are not symmetrical, the solution is to use eye glasses whose lenses are not symmetrical in a complementary way. An extreme version of this is a cylinder lens – a shape that is circular in one direction where it acts as a normal lens and flat in the other direction where it has no focusing properties. The strength and direction of the cylinder are designed to compensate for the asymmetry of the cornea or the eye lens. The curved portion of the cylinder lens may be positive or negative, depending on the detailed asymmetry in the eye.

It is easy to recognize eye glasses designed to deal with astigmatism, since acting by themselves they produce distorted images of any object. This is because they have different focal lengths (and therefore different magnifications) in different directions. Although people with mild hyperopia or myopia can often function without wearing glasses, people with astigmatism usually must wear glasses all the time. Otherwise they often get eye strain headaches caused by the eye muscles trying (and failing) to focus on the vertical and horizontal portions of an object at the same time.

The cylinder lens may be combined with an additional positive or negative lens if the person also has hyperopia or myopia.

4. decreased accommodation or “presbyopia”

This condition is typical of middle age. The eye muscles gradually weaken with age, so that the range of accommodation decreases. People with this condition cannot bring both near objects and far objects into focus. The weakening of the eye muscles often causes the focal length of the eye lens to increase as well, so that many people of middle age tend to become far sighted.

Since the problem is inadequate accommodation, no single lens can correct it, and people with this problem usually need bifocals, which are glasses with two different lens strengths: one for near and one for distant objects. The usual arrangement is that the bottom half of the lens is the near strength and the top half is the far strength.

5. cataracts

This condition is also common among older people, although it can happen at any age. The lens of the eye becomes clouded and partially opaque. The result is that objects at all distances become blurred. The cloudiness of the lens often scatters light in all directions, so that people with cataracts are more easily blinded by bright lights. The lens often becomes yellow as well, and tends to absorb blue light.

The solution is to replace the lens with a fixed-focus lens of the correct strength. Since the replacement lens cannot be deformed, the result is that people with replacement lenses have no accommodation at all and must wear bifocals (or even trifocals).

The simple magnifying glass

In this section we describe how a single positive lens can be used as a simple magnifying glass. Recall that the magnification of a lens is the ratio of the size of the image, s_i to the size of the object, s_o. As we have shown before (Topic 31. Camera Lenses, Part 3. The Effect of Focal

Length), the magnification of a simple lens is the ratio of the distance from the lens to the image, d_i divided by the distance from the lens to the object, d_o:

$$m = \frac{s_i}{s_o} = \frac{d_i}{d_o}$$

The distance, d_i , between the eye-lens and the retina is fixed by the size and shape of the eye. Therefore, the magnification can be changed only by changing the distance between the object and the eye, d_o. Since this distance is in the denominator of the fraction above, the magnification increases as an object is brought closer to the eye. The maximum magnification will be realized when the object is brought to the near point. The magnification could be increased by bringing the object still closer, but the strength of the eye-lens can not be made great enough to bring the object into focus at these short distances. Since the distance between the eye-lens and the retina is about 20 mm and the near-point is about 25 cm, the maximum magnification of the eye is about 2/25= 0.08. (Note that this magnification describes the relationship between physical size of the image on the retina and the physical size of the object. All animals learn to relate the physical size of an image on the retina back to the physical size of the object that produced it.)

If a positive lens is placed in front of the eye, the combination of this lens and the eye itself result in a lens whose strength is the sum of the strengths of the two lenses. (See Topic 27. Compound Lenses). In other words, the focal length of the combination is smaller than the focal length of the eye itself. This compound lens produces an image which is in front of the retina and is therefore out of focus. However, since the image is now in front of the retina, the image can be moved back onto the retina by bringing the object still closer to the eye. This will bring the object back into focus on the retina and simultaneously increase the magnification. Thus the simple magnifying glass works by allowing the object to be brought closer to the eye while still being kept in focus.

Since the focal length of the eye-lens can be changed by the muscles in the eye, a calculation of the magnification of the additional lens in front of the eye depends on how the eye-lens is configured. These results are easy to calculate using the algebraic form of the ray-tracing rules.

There are two extreme situations: (1) the eye lens is configured with as short a focal length as possible so that it is still focused on the object at the near-point. If a person then puts an additional positive lens whose focal length is f (in cm) in front of the eye, and if the eye-lens remains set to its shortest possible focal length, then the magnification (relative to the magnification of the eye by itself as described above) is

$$m = 1 + \frac{25}{f}$$

In order to be in sharp focus in this situation, the object must be placed at a distance of

$$d_o = \frac{25}{25 + f} f$$

Since the denominator of the fraction is always greater than the numerator, the value of the fraction is always less than 1, so that the object distance is less than the focal length of the positive lens that has been added. For example, if the focal length of the additional lens is 5 cm, the magnification is 1+5=6, and the object will be in focus when it is about 4.2 cm in front of the eye.

The second situation is when the eye lens is configured to be as relaxed as possible so that it is focused on a very distant object. If a person then puts an additional positive lens whose focal length is f (in cm) in front of the eye, and if the eye-lens remains set to its longest possible focal length, then the magnification (relative to the magnification of the eye itself as described above) is

$$m = \frac{25}{f}$$

In order to be in sharp focus in this situation, the object must be located at a distance equal to the focal length of the additional lens. That is,

$$d_o = f$$

Since this distance is somewhat greater than the distance in the first situation, it should not be surprising that the magnification is somewhat smaller. Using the same additional lens as above, the magnification would be 5 and the object would be in clear focus when it was 5 cm in front of the eye.

The Compound Microscope

The simple magnifying glass can magnify an image by a relatively small factor – usually no more than about 5. Although larger magnifications are possible, the aberrations of the lens become more important as the magnification increases. The magnifying glass also becomes more and more difficult to use, since the object must be moved closer and closer to the eye as the magnification increases.

The simple magnifier can be combined with other lenses in a number of configurations. Although all of the configurations are similar, each of these configurations is designed for a specific purpose, and the details are different.

The first such configuration is the compound microscope, which is designed to provide greater magnification than can be provided by a single lens. The microscope is designed to magnify objects that can be brought close to the device. The magnified image is inverted, but this is usually not a serious problem for most applications.

In its simplest form, the microscope consists of two positive lenses, which are mounted at either end of a tube as shown in the following diagram:

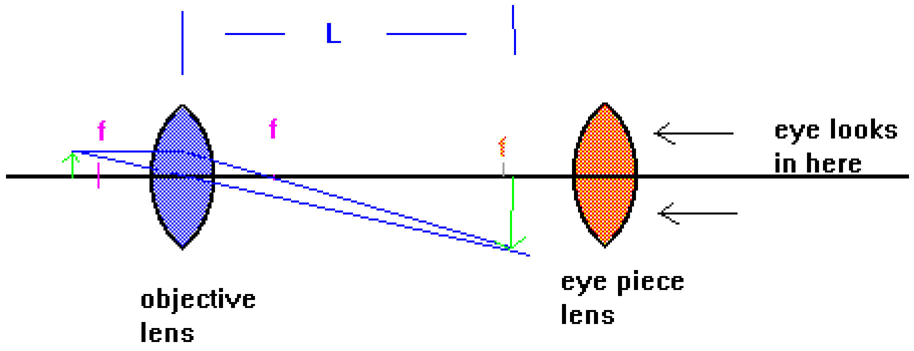

The first lens has a very short focal length, and the object is placed in front of this lens just beyond the focal distance. Since the lens is a positive lens, it forms a real, inverted image at a distance L behind the lens as shown. Since the distance from the first lens to the image is much greater than from the object to the first lens, this image is magnified by the ratio of these two distances. The size of the length L is set by the geometry of the device, and is often about 16 cm.

For example, if the focal length of the first lens is 1 cm and if L is 16 cm, then, using the algebraic form of the ray-tracing rules, the object must be 1.07 cm in front of the lens (that is, just greater than the focal length of 1 cm). The magnification of the first lens is therefore 16/1.07 or about 15. More generally, the magnification of the first lens is approximately L/f_o, where f_o is the focal length of the objective lens in cm. The magnification could be increased by increasing L, but there is a limit to how big L can become before the microscope tube is too long to be used conveniently.

The second lens uses this real image as its object and acts as a standard simple magnifier. As we showed in the previous section, the magnification of this lens depends somewhat on whether the image it produces is viewed when the eye muscles are relaxed so that the eye lens is focused at infinity or contracted so that the eye lens is focused at the near point. In either case, the magnification is approximately $25/f_e$, where f_e is the focal length of the lens in the eye piece in cm.

The eye piece lens is often equipped with an adjustment that can move it in and out along the length of the axis of the instrument. This adjustment compensates for the different strengths of the eye lenses of different users. As we showed in the previous section, the combination of the lens in the eye piece of the microscope and the lens in the eye itself must produce a real image on the retina, and the eye lens in the microscope is adjusted until this image is in sharp focus there.

The overall magnification of the microscope is the product of the magnifications of the two lenses, and overall magnifications of 100X are relatively easy to obtain.

The magnification of each lens increases as its focal length is made shorter. Since the f/number is the ratio of the focal length to the diameter, decreasing the focal length also decreases the f/number. A smaller f/number means that the lenses capture more light from the object, which is a clear advantage. However, a smaller f/number also means that the depth of field will be correspondingly reduced. Since the amount of light is usually not a problem in a laboratory environment, the advantage of greater sensitivity is often not very important, whereas any decrease in the depth of field can be serious.

Therefore, the real trade off in the design of a microscope usually is the balance between magnification and depth of field. A common solution to this problem is to decrease both the focal length and the diameter of the objective lens so that the f/number, which is the ratio of these two quantities, remains more or less constant. Most high-magnification objective lenses are very small in diameter for this reason. The idea is to keep the focal length as short as possible to maximize the magnification while keeping the diameter as small as possible to increase the f/number and thereby give the user more depth of field.

The astronomical telescope

The astronomical telescope is similar in design to the compound microscope discussed in the previous topic. However, unlike the microscope, which is designed to view objects that are small and nearby, the telescope is designed to view objects that may be large, but are very far away.

As with the previous instruments, the magnification of the astronomical telescope is with respect to the performance of the unaided eye. If the distance between the eye lens and the retina is d, and if the object to be viewed is at a much greater distance D away, then the magnification of the eye lens itself is simply d/D, which is a very small number. The magnification of the telescope will be relative to this value.The astronomical telescope consists of two lenses as shown in the following figure. The first lens, called the objective, has a focal length that is as long as possible. The focal distance of this lens is shown as “f” in the figure.

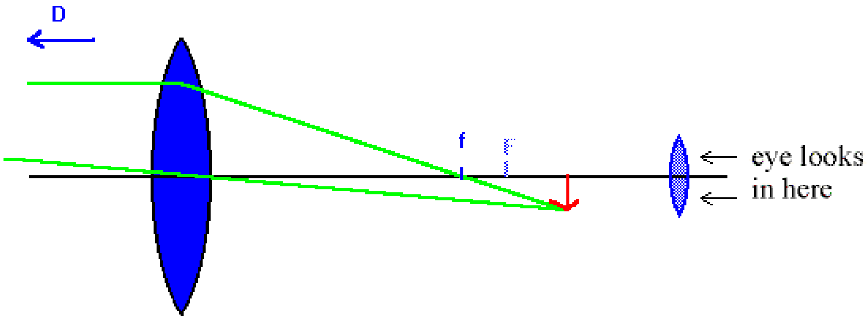

Since the object is a very great distance away, the first lens forms a real, inverted image just beyond its focal point. (The distance beyond the focal point of the objective lens is exaggerated in the figure above.) The magnification of the first lens is approximately f/D. Since D is a very large distance, this magnification is generally a very small number too. Nevertheless , although f/D may be a small number, it is already larger than the magnification of the eye lens by itself, which is only d/D, where d is the distance between the eye lens and the retina. The reason for this improvement is that the objective lens can have a focal length that is much greater than the focal length of the eye lens, which is roughly the diameter of the eye itself.

The real image formed by the objective lens acts as the object for the second lens, which is the usual simple magnifier. The second lens is positioned so that the image formed by the first objective lens is essentially at the focal distance of the second lens, which is shown as “F” in the figure. (The distances are exaggerated in the figure above.) Since the eye is positioned just behind the eye-lens of the telescope, the combination of the eye lens in the telescope and the lenses in the eye take an object that is essentially a distance F in front of the lens combination and produce an image on the retina a distance d behind the combination. The magnification of this combination is therefore d/F. The eye-lens in the telescope often has a small adjustment range to compensate for the different strengths of the eye lens and cornea in different people.

The magnification of the combination is the product of the two, or

$$m = \frac{d}{F} \times \frac{f}{D} = \frac{f}{F} \times \frac{d}{D}$$

The magnification is thus the product of two terms which can be re-arranged as shown. The first is the ratio of the focal lengths of the two lenses in the telescope and the second is the magnification of the eye lens acting by itself. The telescope therefore improves on the eye itself by the factor f/F. To maximize this value, the focal length of the first lens must be as long as possible and the focal length of the second lens must be as short as possible. The second criterion is the usual result for a simple magnifier.

From the figure, the separation between the lenses is the *sum* of the two focal lengths, so that increasing the focal length of the first lens increases the length of the instrument by the same factor. Thus a high-magnification version of this design tends to be a very long device – while decreasing the second focal distance can compensate to some extent, the second focal distance is usually so short already that its value does not make a significant contribution to the total length of the device. The length is sometimes made somewhat smaller by folding the length using prisms or mirrors. This is usually done in binoculars.

There is the usual trade-off between magnification, which increases as the focal length of the first lens is made longer and light-gathering power, which depends on the inverse of the f/number of the first lens and therefore decreases as the focal length is made longer so that the f/number is made greater. As with the microscope, many designs increase the diameter of the lens to keep the f/number at a reasonable value, but here the tendency is in the opposite direction from the microscope – that is, towards *increasing* the focal the length of the first lens (to increase the magnification) and simultaneously *increasing* its diameter (to keep the f/number and therefore the light-gathering ability at a reasonable value). The design criteria in the microscope tended to push both values in the opposite direction – towards smaller focal lengths and smaller diameters.

Note that the image produced by this telescope is inverted. This is usually not serious for astronomical observations, but a telescope designed for terrestrial use (binoculars, for example) usually have a second positive lens in the optical path that inverts the image again so that the observer sees an erect image of the distant object. This inversion can also be done with other optical elements (prisms, for example).

The Galilean telescope

The Galilean telescope was first invented by Galileo in about 1609. Unlike the astronomical telescope, which is made using two positive lenses separated by the *sum* of their two focal distances, the Galilean telescope has a positive objective and a negative eye lens, which are separated by the *difference* between their two focal distances. As in the previous discussion, we consider an object at a very great distance, D, from the eye. The magnification of the eye itself is d/D, where d is the distance between the eye-lens and the retina.

The objective lens, whose focal length is f, performs the same function as in an astronomical telescope: acting alone, it would form a real, inverted image of the distant object, as shown in the following diagram:

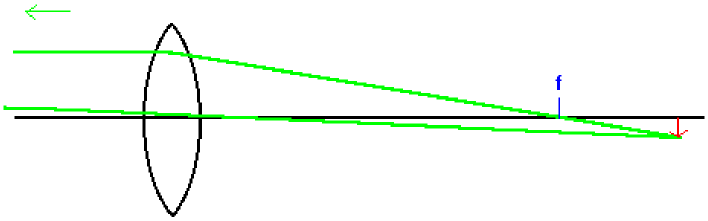

The second negative lens is positioned so that its second focal point is exactly coincident with the second focal point of the objective lens. Thus the rays that are aimed at producing the real, inverted image shown above, actually hit the second lens before they get to the image.

We can use the standard ray tracing rules to construct the image produced by the second lens, keeping in mind that the top ray in the figure above, which was headed for the second focal point of the first lens, is also headed for the same focal point of the second lens, since the two are coincident by construction. Using the ray-tracing rules for a negative lens, this ray is refracted parallel to the axis, and this is shown in red.

The dark green ray is another ray that forms the image produced by the first lens. Its direction is known from the figure above, since it originates at the object, is refracted by the first lens and ends at the image of the first lens. Note that it is also drawn to pass through the center of the second negative lens. The second lens has no effect on this ray, since it passes through its center.

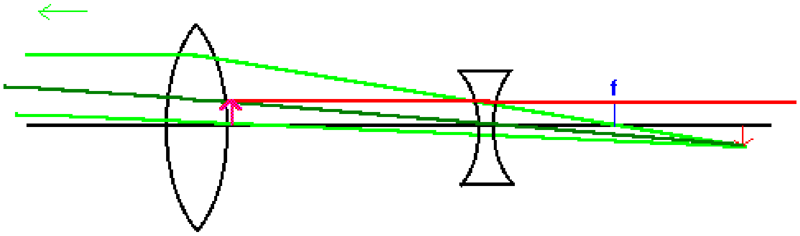

The dark green ray and the red ray appear to come from the purple image as shown. Therefore, the effect of the telescope is to produce a virtual, erect image of the distant object. Note that the image is erect rather than inverted. This virtual image is then viewed by the lens in the eye to form a real image on the retina.

Using the algebraic form of the ray-tracing rules, we can show that the magnification of this arrangement (relative to the unaided eye) is the same as for an astronomical telescope whose lenses have the same numerical focal lengths.

This arrangement is smaller than an astronomical telescope of the same magnification, since the lenses are closer together. In addition, the image is erect and so does not need a separate inversion as with the astronomical design. This configuration is often used in "opera glasses" and other simple binoculars where small size is an important requirement.

Reflecting telescopes

All of the refracting telescopes we have discussed in the previous topics share a common limitation – the f/number of the telescope is limited by the size of the objective lens, and it is difficult and expensive to build high-quality lenses with large diameters. In addition, the weight of the lens increases at least as the square of its diameter, so that supporting a large lens at its edge is a hard job. Since the ray-tracing rules for lenses and mirrors are very similar, any optical design that can be realized with lenses can usually be realized with mirrors as well.

For example, an astronomical telescope constructed using two positive lenses that are separated by the sum of their two focal lengths can also be constructed using two concave mirrors with the same separation. This design is usually called a "Gregorian" configuration as shown in the following figure.

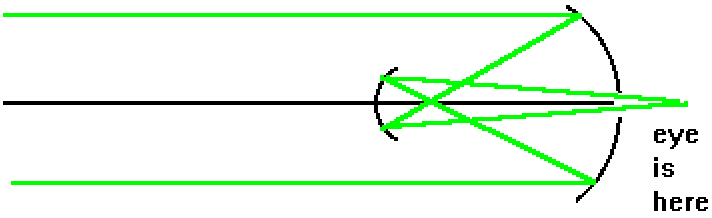

The basic idea is the same: the first mirror forms a real, inverted image near its focal point, and the rays from this image strike a second mirror where they are directed to the eye of the observer. Sometimes a simple magnifier lens is inserted into this path. Note that the second mirror partially blocks the field of view, but this is usually a small obstruction relative to the size of the main mirror. This configuration also requires drilling a small hole in the main mirror.

In addition to being easier to construct and support front-surface mirrors have no chromatic aberration, which is an advantage in many astronomical applications. Most large telescopes use mirrors rather than lenses for these two reasons.

In addition to the “Gregorian” design shown in the figure above, there are a number of other configurations, which differ primarily in the shape and location of the second mirror. Instead of using two concave mirrors as above, a “Cassegranian” design replaces the second small mirror with one that is convex. This concave-convex design is similar in concept to the Galilean telescope, which uses a positive-negative lens combination. There are also designs which use a combination of a large mirror to form the primary image and secondary lenses to focus this image into the eye or onto a sheet of film.