Machine Vision Lens Fundamentals

All the information a machine vision system collects comes through the lens. The correct choice of lens can reduce image processing requirements, and improve system performance and robustness. Software cannot correct the effects of a poorly chosen lens.

This primer provides the technical and practical information you need to choose a lens for a machine vision system. First we review design principles, providing simple formulas that form the basis of further calculations. From models, we proceed to a discussion of real-world lenses and practical parameters. A discussion of special lenses completes this primer.

FIRST ORDER DESIGN THEORY

To establish an understanding of theoretical principles, we'll review a few basic lens definitions and parameters. We then examine the thin lens model. The thin lens model describes a lens with no limitations, one that can be used at any magnification and work at any conjugate. But the thin lens model does not provide the complete picture, since real lenses do have limitations. Following this theoretical discussion, we examine real lenses and their parameters, as well as special lenses.

Camera Format

The camera format provides the dimensions of the image sensor. Lenses, by design, provide images over a limited area. Be sure the lens covers an area as large or larger than the camera format.

Field of View (FOV)

The FOV is the object area that is imaged by the lens onto the image sensor. It must cover all features to be measured, with additional tolerance for alignment errors. It is also good practice to allow some margin (e.g., 10 percent) for uncertainties in lens magnification. Features within the field of view must appear large enough to be measured. This minimum feature size depends on the application. As an estimate, each feature must have 3 pixels across its width, and 3 pixels between features. If there are more than 100 features across a standard camera field, consider using multiple cameras.

Magnification

The required magnification is :

$$\text{mag} = \frac{W_{camera}}{W_{FOV}}$$

where W_{camera} is the width of the camera sensor and W_{FOV} is the width of the FOV. Note that the required magnification depends on the camera sensor size.

Working Distance

The working distance is the distance from the front of the lens to the object. In machine vision applications, this space is often needed for equipment or access. In general, a lens that provides a long working distance will be larger and more expensive than one that provides a shorter working distance.

Thin Lens Model

To understand machine vision lenses, we start with the thin lens model. It is not an exact description of any real lens, but shows lens principles. It also provides terms with which to discuss lens performance. A ray, called the chief ray, follows a straight line from a point on the object, through the center of the lens, to the corresponding point on the image (figure 1). The lens causes all other rays that come from this same object point and that reach the lens to meet at the same image point as the chief ray. Those rays which pass through the edge of the lens are called marginal rays.

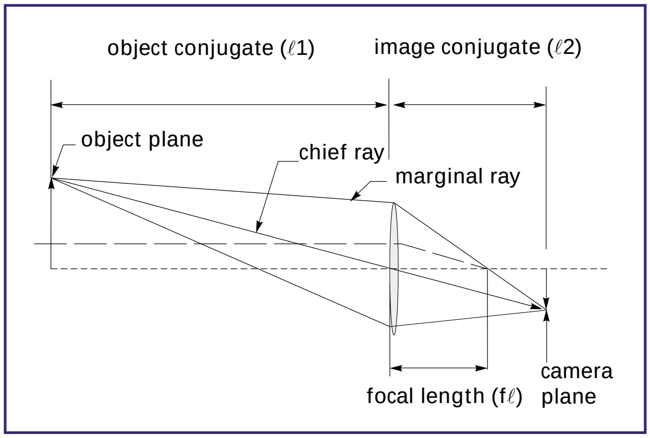

Figure 1. **Thin lens model**

The distance from the object plane to the lens is called the object conjugate. Likewise, the distance from the lens to the camera plane is called the image conjugate. These conjugates are related by the lens maker's formula:

$$\frac{1}{f\ell} = \frac{1}{\ell 1} + \frac{1}{\ell 2}$$

Focal Length

If we let the object conjugate get very large, we see

$$\frac{1}{f\ell} \approx \frac{1}{\ell 2} \Rightarrow \ell 2 \cong f\ell$$

In other words, the focal length is the distance between the lens and the camera plane when the object is at infinity. For photographic lenses, the objects are usually far away, so all images are formed in nearly the same plane, one focal length behind the lens.

From figure 1 and geometry, we can see that

$$m = \frac{\ell 2}{\ell 1}$$

The magnification is the ratio of the image to the object conjugates. If the focal length of a lens increases for a specified magnification, both object and image conjugates increase by the same ratio.

f/Number

The f/number describes the cone angle of the rays that form an image (figure 2). The f/number of a lens determines three important parameters:

- Brightness of the image
- Depth of field
- Resolution of the lens

For photographic lenses, where the object is far away, the f/number is the ratio of the focal length of the lens to the diameter of the aperture. The larger the aperture, the larger the cone angle

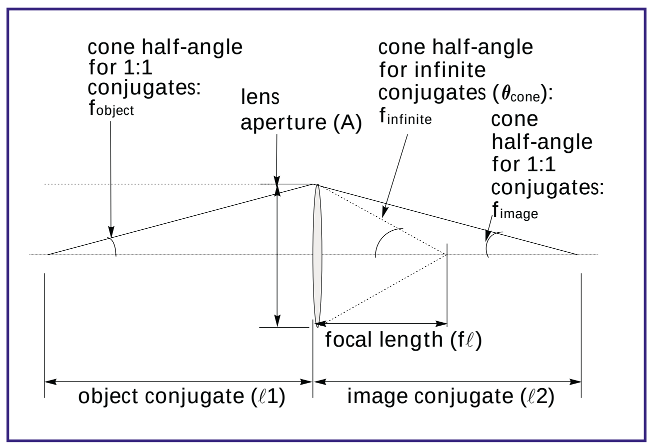

Figure 2. **f/number**

THIN LENS EXAMPLE

We need a magnification of 0.5×, with a working distance of 50 mm. We want to find the correct lens focal length and total system length (TSL). From the equations (after some algebra), we get:

$$f\ell = \frac{m}{m+1} \times \ell 1$$

so

$$f\ell = \frac{0.5}{1.5} \times 50 \text{ mm} = 16.7 \text{ mm}$$

$$\ell 2 = \ell 1 \times m = 50 \text{ mm} \times 0.5 = 25 \text{ mm}$$

$$TSL = \ell 1 + \ell 2 = 50 \text{ mm} + 25 \text{ mm} = 75 \text{ mm}$$

Therefore, we need a lens with focal length of ~17 mm. The total system length is ~75 mm.

and the smaller the f/number. A lens with a small f/number (large aperture) is said to be "fast" because it gathers more light, and photographic exposure times are shorter. A well-corrected fast lens forms a high-resolution image, but with a small depth of field. A lens with a large f/number is said to be "slow." It requires more light, but has a larger depth of field. If the lens is very slow, its resolution may be limited by diffraction effects. In this case, the image is blurred even at best focus.

The f/number printed on a photographic lens is the infinite conjugate f/number. It is defined as:

$$f_\infty = \frac{f\ell}{A}$$

where $f\ell$ is the focal length of the lens and A is the diameter of the lens aperture. When the lens is forming an image of a distant object, the cone half-angle of the rays forming the image is:

$$\theta_{cone} = \operatorname{atan}\left(\frac{1}{2 \times f_\infty}\right)$$

This infinite conjugate f/number is only applicable when the lens is imaging an object far away. For machine vision applications, the object is usually close and the cone angle is calculated from the working f/number (see box).

Numerical Aperture (NA)

For lenses designed to work at magnifications greater than 1 (for example, microscope objectives), the cone angle on the object side is used as the performance measure. By convention, this angle is given as a numerical aperture (NA). The NA (figure 3) is given by:

$$NA = \sin(\theta_{cone})$$

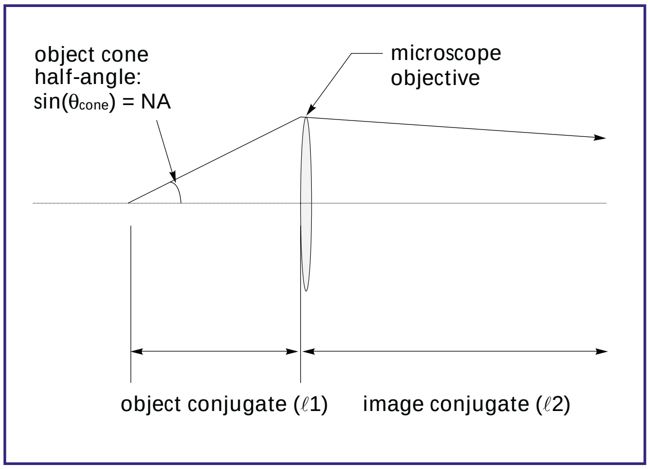

Figure 3. **Numerical aperture (NA)**

f/NUMBER (WORKING)

In machine vision, the working f/number describes lens performance:

$$f_{image} = \frac{\ell 2}{A}$$

$$f_{object} = \frac{\ell 1}{A}$$

where $\ell 2$ and $\ell 1$ are the image and object conjugates, respectively. f_{image} is called the working f/number in image space, or the simply image side f/number. Similarly, f_{object} is the object side f/number.

For close objects, f_{image} is larger than $f_{infinty}$, so the lens is "slower" than the number given on the barrel. For example, a lens shown as f/4 on its barrel will act like an f/8 lens when used at a magnification of 1.

The object-side f/number determines depth of field. It is given by:

$$f_{object} = \frac{1}{m} \times f_{image}$$

NA is related to f/number by these exact relationships:

$$\mathrm{NA} = \sin\left(\operatorname{atan}\left(\frac{1}{2 \times \mathrm{f/number}}\right)\right)$$

$$\mathrm{f/number} = \frac{1}{2 \times \tan\left(\operatorname{asin}(\mathrm{NA})\right)}$$

For NA < 0.25 (f/number >2), these simplify to:

$$\mathrm{NA} \cong \frac{1}{2 \times \mathrm{f/number}}$$

$$\mathrm{f/number} \cong \frac{1}{2 \times \mathrm{NA}}$$

REAL-WORLD LENSES

Thick Lens Model

The thin lens model treats a lens as a plane with zero thickness. To model a real-world lens, we divide this thin lens plane into two planes (figure 4). These planes contain the entrance and the exit pupils of the lens. Everything in front of the entrance pupil is said to be in object space. Everything behind the exit pupil is said to be in image space. How light gets from the entrance pupil to the exit pupil is not considered in this model.

In object space, we think of the real-world lens as a thin lens located at the entrance pupil. The entrance pupil is generally located within the physical lens, but not always. Wherever it is located, light rays in object space proceed in straight lines until they reach the entrance pupil. The effects of any elements in front of this position are taken into account when the entrance pupil position is calculated. In the same way, we think of the real-world lens as a thin lens located at the exit pupil in image space.

For many lenses, the entrance and exit pupils are located near each other and within the physical lens. The exit pupil may be in front of or behind the entrance pupil. For certain special lens types, the pupils are deliberately placed far from their “natural” positions. For example, a telephoto lens has its exit pupil far in front of its entrance pupil (figure 5). In this way, a long focal length lens fits into a short package. A telecentric lens has its entrance pupil at infinity, well behind its exit pupil (figure 6).

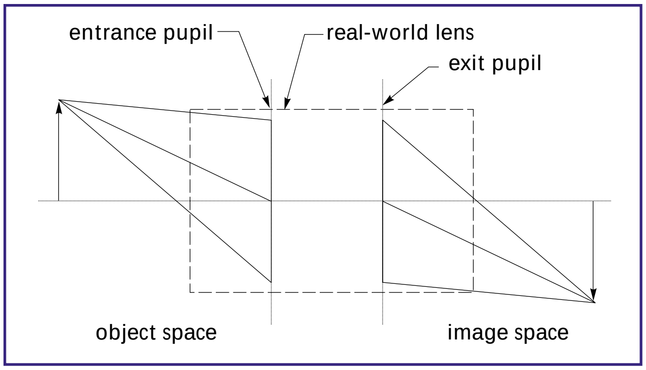

Figure 4. **Thick lens model**

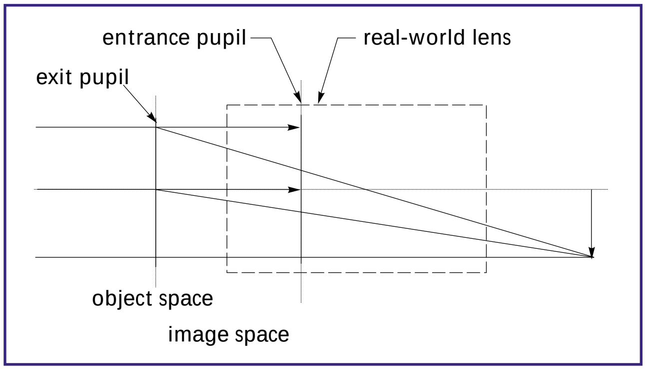

Figure 5. **Telephoto lens**

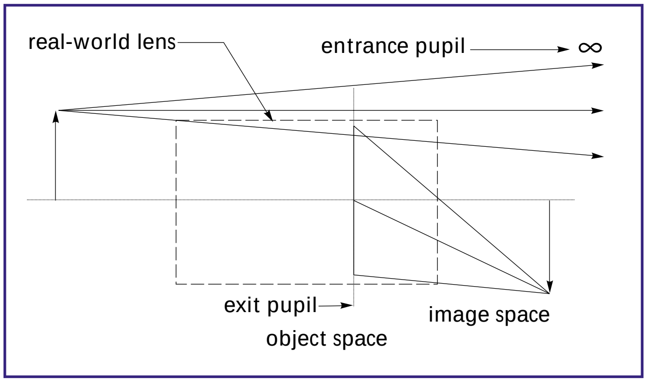

Figure 6. **Telecentric lens**

Aberrations

If real lenses followed first order theory, lens design would be easy. Unfortunately, it is difficult to make a real lens approximate this behavior. Diffraction sets a lower limit on image spot size. The differences between ideal “diffraction limited” behavior and real lens behavior are called aberrations.

The job of the lens designer is to choose glasses, curvatures, and thicknesses for the lens’ elements that keep its overall aberrations within acceptable limits. Such a lens is said to be well corrected. It is impossible to design a lens that is well corrected for all conjugates, FOVs, and wavelengths. The lens designer works to correct the lens over the small range of operating conditions at which the lens must function. The smaller the range, the simpler the design can be.

A lens that is corrected for one set of conditions may show significant aberrations when used under a different set of conditions. For example, a surveillance lens is corrected for distant objects, with magnifications of <1/10. By using extension tubes, the image conjugate of the lens can be extended so that the lens forms an image at a magnification of 1. But this image may show significant aberrations. The lens was not corrected to work at these conjugates.

Standard Lenses

Commercial lenses, produced in high volume, are by far the best value in terms of performance per price. Finding a stock lens is the most cost-effective solution to a machine vision problem. Table 1 lists various types and their range of operating conditions. Commercial lenses incorporate design and manufacturing techniques that are not available in custom designs. For example, a lens for a 35 mm single lens reflex (SLR) camera that costs one hundred dollars at the local camera store would cost ten thousand dollars to design and many thousands of dollars to manufacture in small quantities. It is always best to consider commercial lens options before starting a custom lens design.

REAL LENS PARAMETERS

Resolution

Resolution is the ability of an optical system to distinguish two features that are close together. For example, if a lens images a row of pins on an electrical connector, it must have sufficient resolution to see each pin as separate from its neighbors. A lens imaging a lot code on a pharmaceutical bottle must have sufficient resolution to distinguish one character from another. Resolution is also required to make sharp images of an edge. A lens with high resolution will show an edge transition in fewer pixels than a lens with low resolution.

There are many different definitions of lens resolution. They differ by what type of test object is measured (points, bars, sine patterns, or other objects), and by the criteria for determining when two objects are “resolved.” A practical measurement for machine vision uses three-bar targets of various spatial frequencies. A chrome-on-glass USAF-1951 target is a good test object. If the contrast between bar and space is >20 percent, the bars are considered to be resolved.

Resolution does not determine the dimensional accuracy to which objects can be measured. The position of a large object can be determined to within a fraction of a resolution spot under

Table 1. **Common Commercial Lens Types**

Lens type	Magnification	Image format	Object FOV	Focal length	Working f/number (object side)
CCTV	<0.1	1-inch CCD format: 13-mm-diameter and smaller	large	5 mm–50 mm	>20 (adjustable)
35-mm-SLR	<0.2	42-mm-diameter	large	20 mm–300 mm	>10 (adjustable)
SLR macro	~0.5–2	42-mm-diameter	15 mm–60 mm	50 mm–100 mm	>4 (adjustable)
Enlarger	2–20	500 mm	50 mm	40 mm–150 mm	>4 (adjustable)
Copier	~1:1	300 mm	300 mm	50 mm–200 mm	6 (fixed)
Microscope	5–100	requires additional lens	<2 mm	5 mm–40 mm	0.1 NA–0.95 NA (fixed)

suitable conditions. Many vision systems determine positions to one-quarter pixel. On the other hand, if the lens has distortion, or if its magnification is not known accurately, then the measured position of a feature may be in error by many resolution spot widths.

Diffraction

Diffraction limits the resolution possible with any lens. In most machine vision calculations, we consider light as traveling in straight lines (rays) from object points to image points. In reality, diffraction spreads each image point to a spot whose size depends on the f/number of the lens and the wavelength of the light. This spot pattern is called an Airy disk. Its diameter is given by:

$$D_{Airy} = 2.44 \times \lambda \times \text{f/number}$$

where D is the diameter of the inner bright spot, λ is the wavelength of light, and the f/number is the image side f/number. Since the wavelength of visible light is ~ 0.5 μm, this means the diameter of the diffraction-limited spot (in μm) is approximately equal to the working f/number.

For example, a typical CCD camera has pixels that are 10 μm square. To form a diffraction-limited spot of this diameter, the working f/number on the image side should be ~ 10. An f/22 lens forms an image spot larger than a pixel. Its image therefore appears less sharp than that of the f/10 image. An f/2 lens image will not appear sharper than an f/10 image, since the camera pixel size limits the resolution. In this case, the system is said to be detector limited.

Contrast

Contrast is the amount of difference between light and dark features in an image. Contrast (also called modulation) is defined by:

$$\text{contrast} = \frac{\text{light} - \text{dark}}{\text{light} + \text{dark}}$$

Here, "light" is the gray level of the brightest pixel of a feature, and "dark" is the gray level of the darkest pixel. A contrast of 1 means modulation from full light to full dark; a contrast of 0 means the image is gray with no features. Finer (higher spatial frequency) features are imaged with less contrast than larger features. A high-resolution lens not only resolves finer features, but generally images medium-scale features at higher contrast. A high-contrast image appears "sharper" than a lower contrast image, even at the same resolution.

Factors other than lens resolution can affect contrast. Stray light from the environment, and glare from uncoated or poorly polished optics reduce contrast. The angles of the lens and of the illumination have a great effect on contrast. The contrast of some objects is dependent on the color of the illumination.

Depth of Field

Depth of field (DOF) is the range of lens-to-object distances over which the image will be in sharp focus. It is difficult to calculate because the definition of "sharp" focus is subjective. Our calculations are based on the resolution of the camera. Some applications do not require this high resolution. It is best to test DOF with a breadboard model for any new application.

In general, the geometrical DOF (figure 7) is given by :

$$\text{DOF} = 2 \times \text{f/number}_{\text{object}} \times \text{blur}$$

where blur is the diameter of the allowable blur in object space. A conservative definition of DOF requires that the diffraction spot size created by the lens is one pixel width in diameter, and the geometric blur due to defocus is also one pixel width in diameter. With these assumptions:

$$\text{DOF}(\mu\text{m}) = 2 \times \left(\frac{W_{\text{pixel}}(\mu\text{m})}{m}\right)^2$$

Here, we set the image side f/number of the lens equal to the pixel in μm. W_{pixel} is the pixel width in μm; m is the lens magnification. Thus, for a camera with 10 μm pixels and operating at 0.5× magnification, and an image side f/number of f/10, the DOF is 800 μm, or 0.8 mm.

Holding all other factors equal, the DOF increases with the object side f/number of the lens. Using a smaller aperture opening reduces the resolution of the lens slightly, but greatly increases the DOF. A more generous definition of DOF sets the contrast of the image at the highest spatial frequency that the camera can reproduce (called the Nyquist frequency) to 50 percent at best focus, and allows this contrast to fall to 20 percent at the limits of the

DOF. In this case, we set the image side f/number equal to 1.6 × the pixel width in μm. With these assumptions:

$$\text{DOF } (\mu\text{m}) = 4.2 \times 2 \times \left(\frac{W_{pixel}(\mu\text{m})}{m}\right)^2$$

With this smaller lens aperture, the image will be slightly less sharp, but the depth of field will be four times greater.

Telecentricity

Telecentricity determines the amount that magnification changes with object distance. Standard lenses produce images with higher magnification when the object is closer to the lens. We experience this with our eyes. A hand held up near your face looks larger than when it is moved farther away. For the same field size, a longer focal length shows less magnification change than a short focal length lens.

A telecentric lens acts as if it had an infinite focal length. Magnification is independent of object distance. An object moved from far away to near the lens goes into and out of sharp focus, but its image size is constant. This property is very important for gaging three-dimensional objects, or objects whose distance from the lens is not known precisely.

A telecentric lens views the whole field from the same perspective angle. Thus, deep round holes look round over the entire field, rather than appearing elliptical near the edge of the field. Objects at the bottom of deep holes are visible throughout the field.

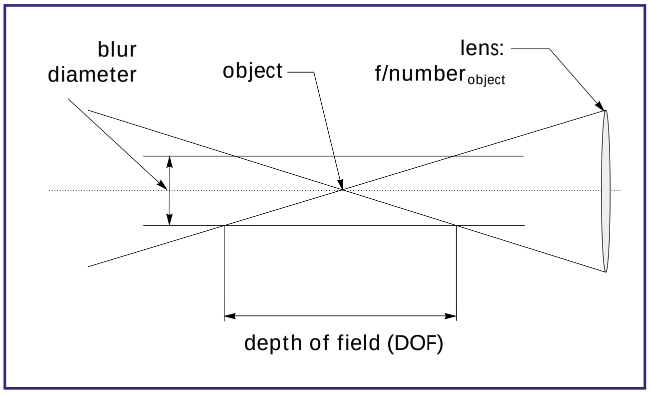

Figure 7. **Depth of field**

The degree of telecentricity is measured by the chief ray angle in the corner of the field (figure 8). In machine vision, a standard commercial lens may have chief ray angles of 10 degrees or more. Telecentric lenses have chief ray angles less than 0.5 degree. Telecentric lenses with chief ray angles of <0.1 degree are available by request.

Telecentricity does not affect the depth of field. Telecentricity is a measure of the angle of the chief ray in object space.

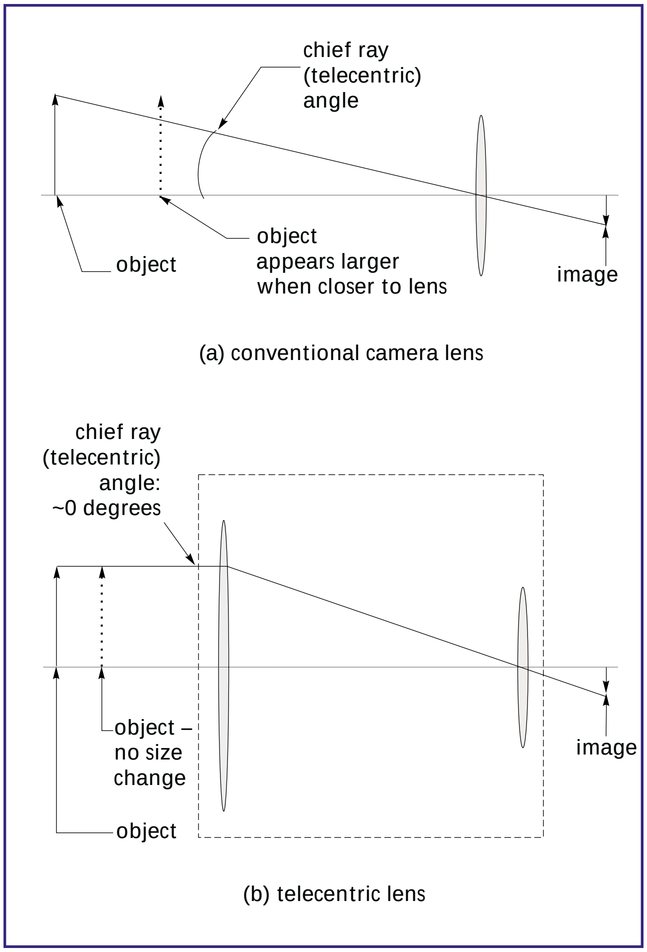

Figure 8. **Telecentricity — (a) conventional camera (b) telecentric lens**

Depth of field is determined by the angles of the marginal rays. Chief ray and marginal ray angles are independent of each other.

The objective element of a telecentric lens must be larger than the field of view. The lens must "look straight down" on all portions of the field. Telecentric lenses for very large fields are thus large and expensive. Most telecentric lenses cover fields less than 6 inches in diameter.

Gaging Depth of Field

The gaging depth of field is the range of distances over which the object can be gaged to a given accuracy (figure 9). A change in object distance changes the image magnification and therefore the measured lateral position of the object. The gaging depth of field describes how precisely the object distance must be controlled to maintain a given measurement accuracy. Telecentric lenses provide larger gaging depths of field than do conventional lenses.

Distortion

In optics, distortion is a particular lens aberration that causes objects to be imaged farther or closer to the optical axis than for a perfect image. It is a property of the lens design and not the result of manufacturing errors. Most machine vision lenses have a small amount of "pincushion" distortion (figure 10). Relative distortion increases as the square of the field, so it is important to specify the field over which field distortion is measured.

Distortion is generally specified in relative terms. A lens which exhibits 2 percent distortion over a given field will image a point in the corner of its field 2 percent too far from the optical axis. If this distance should be 400 pixels, it will be measured as 408 pixels.

Lens distortion errors are often small enough to ignore. Because distortion is fixed, these errors can also be removed by software calibration. Lenses designed to have low distortion are available.

Spectral Range

Most machine vision lenses are color corrected though the whole visible range. Filters that narrow the spectral range to a single color sometimes improve lens resolution. CCD cameras are inherently sensitive to near-infrared (IR) light. In most cases, there should be an IR filter included in the system to reduce this sensitivity. Many cameras have IR filters built in.

SPECIAL LENSES

Zoom

Zoom lenses have focal lengths that are adjustable over some range. They are useful for prototypes, where the focal length requirement is not yet determined. They can be set at focal lengths between those available with fixed lenses. Zoom lenses are larger, less robust, more expensive, and have smaller apertures than similar fixed-focal-length lenses. Also, they frequently have more distortion.

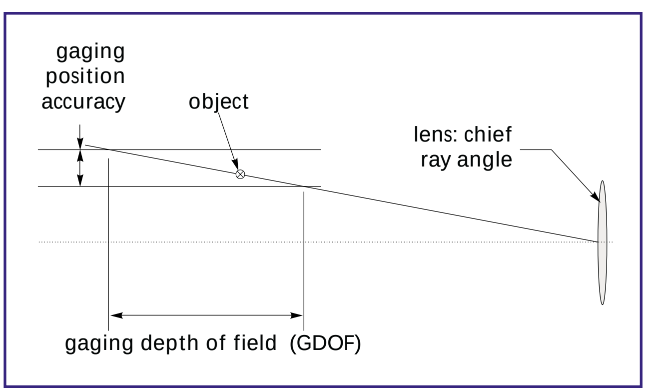

Figure 9. **Gaging depth of field**

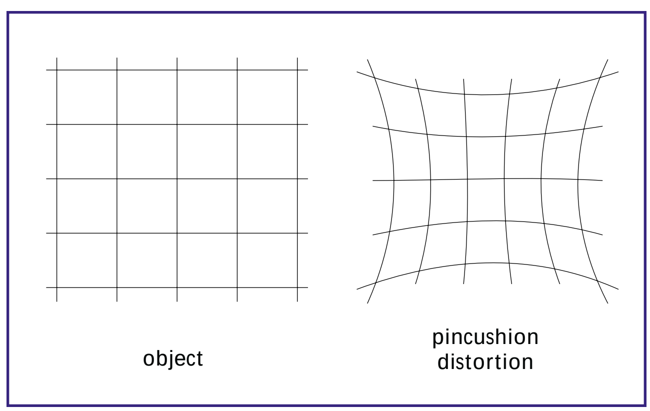

Figure10. **Pincushion distortion**

Macro

A camera lens optimized to work at magnifications near 1 is called a macro lens. A macro lens provides better image quality than a standard camera lens used with extension tubes.

Telecentric

Telecentric lenses provide constant magnification for any object distance (see Telecentricity). They are useful for precision gaging, or when a constant perspective angle across the field is desirable. Their object distance is generally less than that of a standard lenses. The magnification of a telecentric lens is fixed by its design. Because the first element must be as large as the field width, telecentric lenses tend to be larger and more expensive than standard lenses.

Close-up Lenses

Close-up attachment lenses reduce the object distance of a standard lens. The nominal magnification of a lens with a close-focusing attachment is:

$$m = \frac{f\ell_{lens}}{f\ell_{attachment}}$$

where $f\ell_{lens}$ is the focal length of the base lens, and $f\ell_{attachemnt}$ is the focal length of the attachment lens.

The object distance is approximately equal to $f\ell$ attachment. Image quality is somewhat degraded at low f/numbers.

Teleconverters (Extenders)

Teleconverters are short relay optics that fit between the lens and the camera and increase the effective lens focal length. They are usually available at 1.5× and 2× power. The penalties for their use are increased image side f/number (by the power factor), and increased distortion.

Reverse Mounting

For magnifications greater than 1, a camera lens can be used in reverse, with the object held at the usual camera plane and the camera in the usual object plane. In this case, the object distance will be short, while the lens-to-camera distance is long. Adaptors are available to hold camera lenses in this orientation.