Chapter 2

Optical triangulation: limitations and prior work

Active optical triangulation is one of the most common methods for acquiring range data. Although this technology has been in use for over two decades, its speed and accuracy has increased dramatically in recent years with the development of geometrically stable imaging sensors such as CCD's and lateral effect photodiodes.

Researchers and manufacturers have used optical triangulation scanning in a variety of applications. In medicine, optical triangulation has provided range data for plastic surgery simulation, offering safer, cheaper, and faster shape acquisition than conventional volumetric scanning technologies [Pieper et al. 1992]. In industry, engineers have used triangulation scanners for applications that include postal package processing [Garcia 1989] and printed circuit board inspection [Juha & Souder 1987]. Triangulation scanners also provide data to drive computer graphics applications, such as digital film-making [Duncan 1993].

Figure 2.1 shows a typical system configuration in two dimensions. The location of the center of the reflected light pulse imaged on the sensor corresponds to a line of sight that intersects the illuminant in exactly one point, yielding a depth value. The shape of the object is acquired by translating or rotating the object through the beam or by scanning the beam across the object.

In this chapter, we begin with an overview of optical triangulation configurations (Section 2.1). Next, we discuss the limitations of traditional methods in both qualitative and

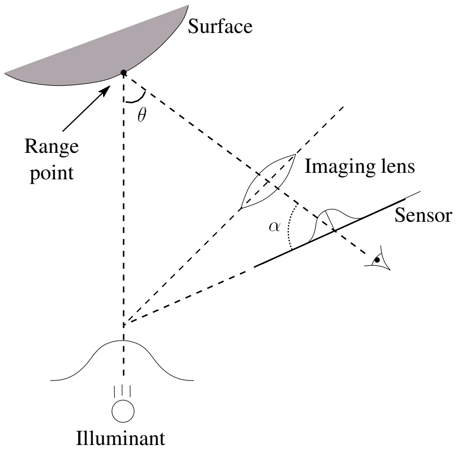

Figure 2.1: Optical triangulation geometry. The illuminant reflects off of the surface and forms an image on the sensor. The center of the illuminant maps to a unique position on the sensor based on the depth of the range point. In order for all points along the center of the laser sheet to be in focus, the angles θ and α are related to one another by Equation 2.1.

a quantitative contexts (Section 2.2). Finally, we describe previous efforts to evaluate and correct for the errors inherent in traditional triangulation methods (Section 2.3).

2.1 Triangulation Configurations

The range acquisition literature contains many descriptions of optical triangulation range scanners, of which we list a handful [Rioux et al. 1987] [Hausler & Heckel 1988] [Mundy & Porter 1987]. Several survey articles have also appeared [Jarvis 1983] [Strand 1983], including Besl's excellent survey [Besl 1989] which describes numerous optical range imaging methods and estimates relative performances. The variety of optical triangulation configurations differ primarily in the structure of the illuminant (typically point, stripe, multi-point, or multi-stripe), the dimensionality of the sensor (linear array or CCD grid), and the scanning method (move the object or move the scanner hardware).

2.1.1 Structure of the illuminant

The structure of the illuminant can take a variety of forms. A beam of light forms a spot on a surface and provides a single range value. By passing the beam through a cylindrical lens, a light stripe can be projected onto an object to collect a range profile (Figure 1.2b). Researchers have also tried projecting multiple spots or multiple stripes onto an object for more parallelized shape acquisition, though multiple steps are usually required to disambiguate the light reflections. When imaging the reflected light onto a sensor with lenses, the single point and stripe illuminants offer the advantage that, at any instant, all intersections of the light with the object must lie in a plane. Since lenses image points in a plane to points in another plane, the sensor can be oriented to keep the beam or sheet of light in focus, thus reducing depth of field problems. When the focal plane is tilted, the image plane must also be tilted so as to satisfy the Scheimpflug condition [Slevogt 1974]:

$$\tan \alpha = M \tan \theta \tag{2.1}$$

where θ and α are the tilt angles of the focal and image planes, respectively, as shown in Figure 2.1, and M is the magnification on the optical axis. The resulting triangulation geometry has the property that the focal, image, and lens planes all intersect in a single line. Note that multi-point and multi-stripe systems generally cannot take advantage of this optimal configuration, because the illumination usually does not lie in a plane.

2.1.2 Type of illuminant

The type of illuminant can be either coherent or incoherent. Coherent illuminants such as lasers offer several distinct advantages over their incoherent counterparts. First, lasers can be held in tight focus over a long range. Second, since laser light is tuned to a single wavelength, the sensor can be coated with a bandpass wavelength filter, decreasing sensitivity to ambient light. In addition, the optics do not have to be corrected for chromatic aberrations, and elements such as prisms can be used with ease. Lasers used in triangulation also do not typically have problems with heat dissipation, whereas incoherent illuminants of sufficient power frequently do. The disadvantages of using lasers are laser speckle (random coherent interference due to surface roughness) and the need for special safety precautions among lasers operating at visible and ultraviolet wavelengths. We will discuss the impact of laser speckle on range accuracy in Sections 2.2.4 and 3.4.

2.1.3 Sensor

Sensors for optical triangulation systems also come in a variety of forms. For narrow beam illumination, point (zero-dimensional) sensors such as photodiodes or line (one-dimensional) sensors such as lateral effect photodiodes and linear array CCD's are sufficient, though point sensors must be scanned to provide another dimension. For light stripe, multi-point, and multi-stripe systems, a two-dimensional sensor is necessary and typically comes in the form of a CCD array, though point and line sensors can also be scanned to provide the required dimensions.

2.1.4 Scanning method

The method of scanning an object is also a matter of choice. During a single range imaging sweep, the illumination-sensor system can be stationary while a platform translates and rotates the object through the field of view, or the object may be stationary while the scanner moves. Alternatively, both the object and the scanning equipment may be stationary while rotating mirrors sweep the illuminant and sensor viewing directions across the object. For this latter scanning approach, the sensor viewpoint must sweep in synchrony with the illuminant to ensure that the sensor tracks the moving illuminant and keeps it in focus (see, for example, [Rioux et al. 1987]). In addition, an optical triangulation scanner can collect range data in continuous sweeps such that consecutive range points are adjacent in space (barring surface discontinuities). Alternatively, the scanning may proceed in spatially disjoint steps where the object moves an arbitrary amount relative to the scanner between acquisitions. In Chapters 3 - 4, we will be concerned with scanners which perform continuous sweeps or which move in small steps relative to the object, because these types of motion permit correction of systematic artifacts.

2.2 Limitations of traditional methods

For optical triangulation systems, the accuracy of the range data depends on proper interpretation of imaged light reflections. The most common approach is to reduce the problem to one of finding the "center" of a one dimensional pulse, where the "center" refers to the position on the sensor which hopefully maps to the center of the illuminant. Typically, researchers have opted for a statistic such as mean, median or peak of the imaged light as representative of the center. Each of these statistics gives the correct answer when the surface is perfectly planar. In this case, the sensor records a compressed or expanded image of the illuminant's shape, depending on the orientation of the surface with respect to the illuminant and sensor. The location of the center of the imaged pulse is not altered under these circumstances.

2.2.1 Geometric intuition

The general accuracy of these statistics breaks down, however, whenever the surface perturbs the shape of the illuminant. Perturbations of the shape of the imaged illuminant occur whenever:

- The surface reflectance varies.
- The surface geometry deviates from planarity.
- The light paths to the sensor are partially occluded.
- The light is coherent, and the surface is sufficiently rough to introduce laser speckle.

In Figures 2.2-2.4, we give examples of how the first three circumstances result in range errors even for an ideal triangulation system with infinite sensor resolution and perfect calibration. For purposes of illustration, we omit the imaging optics of Figure 2.1 and treat the sensor as one-dimensional and orthographic. Further, we assume an illuminant of Gaussian cross-section, and we use the mean for determining the center of an imaged pulse. Figure 2.2 shows how a step reflectance discontinuity results in range points that do not lie on the surface. A comparison of Figures 2.2a and 2.2b shows that the range error worsens with larger reflectance steps. Figure 2.3 provides two examples of shape variations resulting in range errors. Note that in Figure 2.3b, the center of the illuminant is not even striking a surface. In this case, a measure of the center of the pulse results in a range value, when in fact the correct answer is to return no range value whatever. Finally, Figure 2.4 shows the effect of

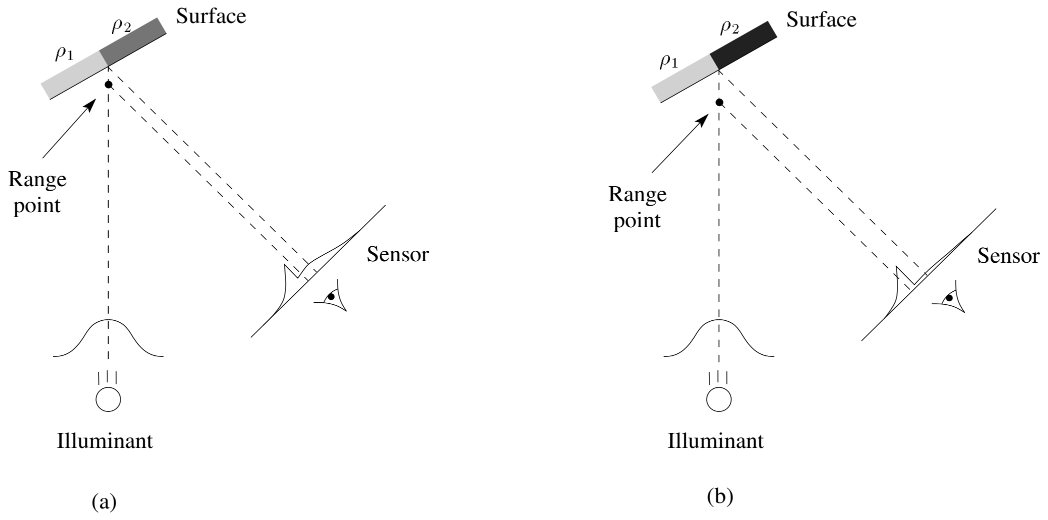

Figure 2.2: Range errors due to reflectance discontinuities. (a) The change in reflectance from light gray to dark gray distorts the image of the illuminant, perturbing the mean and resulting in an erroneous range point. (b) Same as (a), but the reflectance step is larger, causing a larger range error.

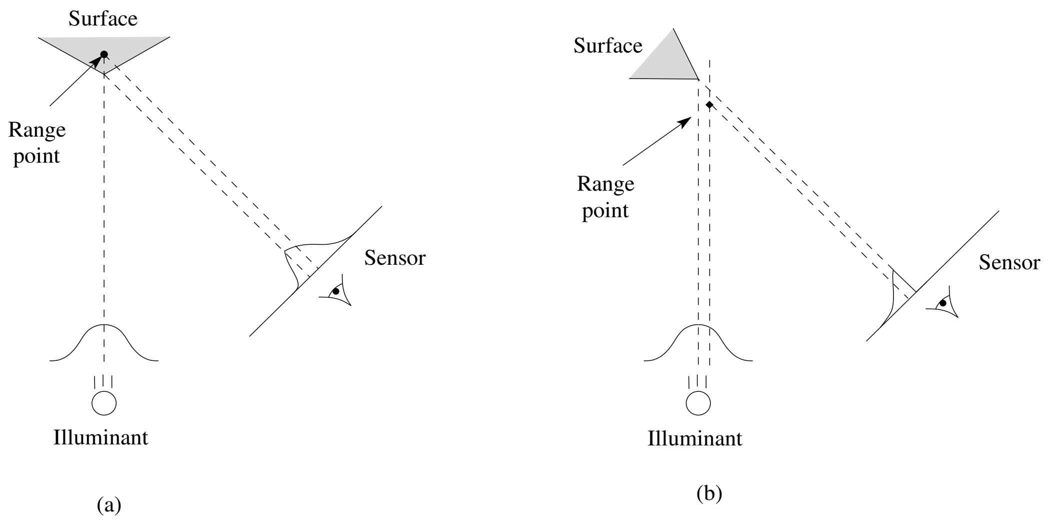

Figure 2.3: Range errors due to shape variations. (a) When reflecting off of the surfaces at the corner shown, the left half of the Gaussian is much more compressed than the right half. The result is a shift in the mean and an erroneous range point. (b) When the illuminant falls off the edge of an object, the sensor images *some* light reflection. In this case, a range point is found where there is no surface at all along the center of the illuminant.

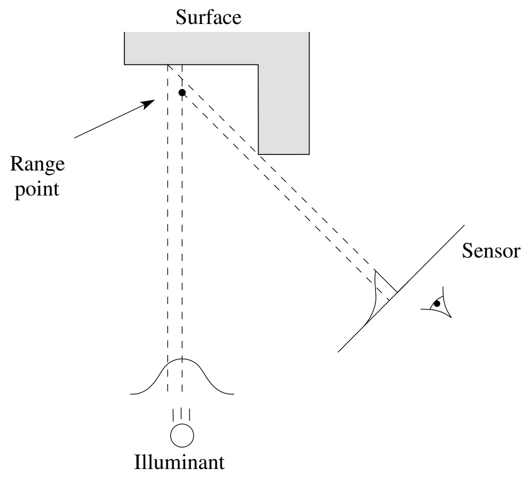

Figure 2.4: Range error due to sensor occlusion. A portion of the light reflecting from the object is blocked before reaching the sensor. The mean yields an erroneous range point, even when the center of the illuminant is not visible.

occluding the line of sight between the illuminated surface and the sensor. This range error is very similar to the error encountered in Figure 2.3b.

The fourth source of range error is laser speckle, which arises when coherent laser illumination bounces off of a surface that is rough compared to a wavelength [Goodman 1984]. The surface roughness introduces random variations in optical path lengths, causing a random interference pattern throughout space and at the sensor. Figure 2.5 shows a photograph of laser speckle arising from a rough surface[1]. For triangulation systems, the result is an imaged pulse with a noise component that affects the mean pulse detection, causing range errors even from a planar target (see Figure 2.6).

[1] Although optically smooth surfaces (i.e., mirrors) do not introduce laser speckle, they are also extremely hard to measure. Mirrored surfaces only reflect light to the sensor when the surface is properly oriented. Further, mirrored surfaces generally result in interreflections that significantly complicate range analysis. As a result, diffusely reflecting surfaces are desirable for range scanning. Diffuseness arises from surface roughness, which in turn leads to laser speckle.

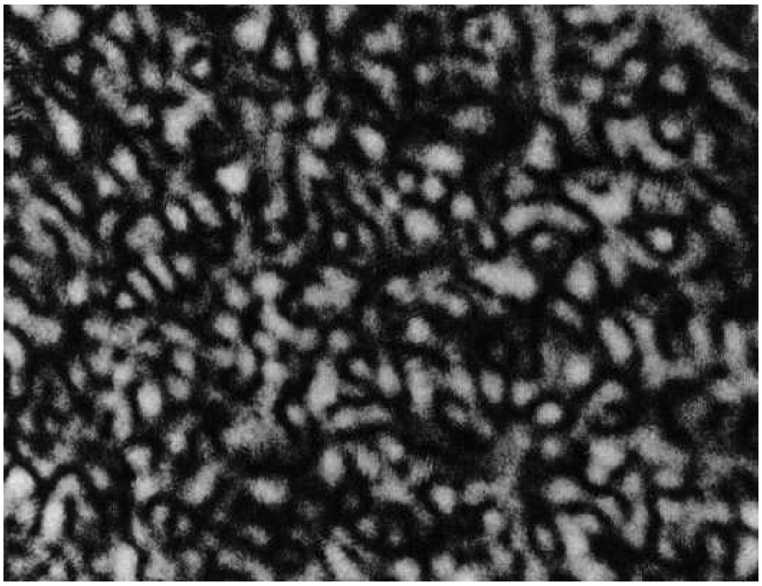

Figure 2.5: Image of a rough object formed with coherent light. Source: [Goodman 1984].

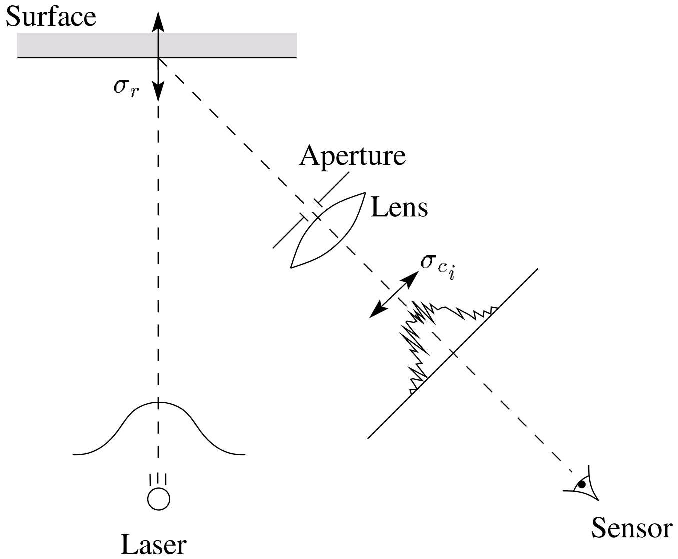

Figure 2.6: Influence of laser speckle on traditional triangulation. The image of the Gaussian is noisy, causing a random shift in the position of the mean, c_i. The uncertainty in the mean's position, σ_{c_i} maps to an uncertainty in the observed range, σ_r. Note that the direction of range uncertainty follows the center of the laser beam.

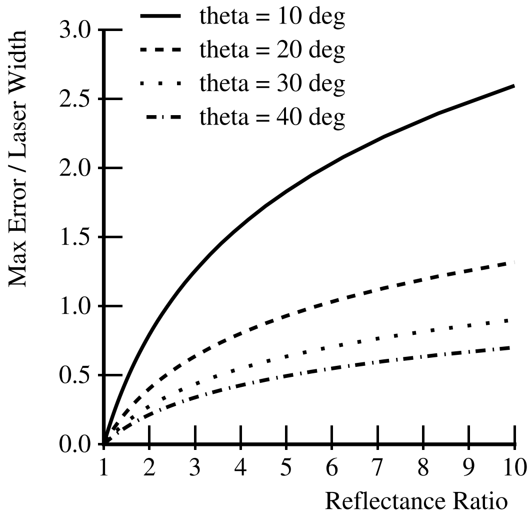

Figure 2.7: Plot of errors due to reflectance discontinuities. As the size of the reflectance step (represented as the ratio of reflectances on either side of the step), the deviation from planarity increases accordingly. Smaller triangulation angles (theta) exhibit greater errors.

2.2.2 Quantifying the error

To quantify the errors inherent in using traditional mean analysis, we numerically computed the errors introduced by reflectance and shape variations for an ideal triangulation system with a single Gaussian illuminant. We define the normally incident irradiance profile of the Gaussian illuminant as:

$$E_L(x) = I_L \exp\left[\frac{-2x^2}{w^2}\right] \tag{2.2}$$

where I_L is the power of the beam at its center, and w is a measure of beam width, taken to be the distance between the beam center and the e^{-2} point. This measure of beam width is common in the optics literature. We present the range errors in a scale invariant form by dividing all distances by the beam width. Figure 2.7 illustrates the maximum deviation from planarity introduced by scanning a reflectance discontinuity of varying step magnitudes for four different triangulation angles. As the size of the step increases, the error increases correspondingly. In addition, smaller triangulation angles, which are desirable for reducing the likelihood of missing data due to sensor occlusions, actually result in larger range errors.

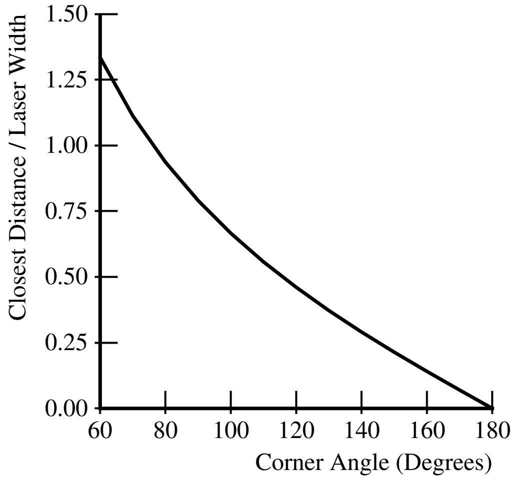

Figure 2.8: Plot of errors due to corners. The y-axis indicates the closest distance between the range data and the actual corner, while the x-axis measures the angle of the corner. Tighter corners result in greater range errors.

This result is not surprising, as sensor mean positions are converted to depths through a division by $\sin\theta$, where θ is the triangulation angle, so that errors in mean detection translate to larger range errors for smaller triangulation angles.

Figure 2.8 shows the effects of a corner on range error, where the error is taken to be the shortest distance between the computed range data and the exact corner point. The corner is oriented so that the illumination direction bisects the corner's angle as shown in Figure 2.3a. As we might expect, a sharper corner results in greater compression of the left side of the imaged Gaussian relative to the right side, pushing the mean further to the right on the sensor and pushing the triangulated point further behind the corner. In this case, the triangulation angle has little effect as the greater mean error to depth error amplification is offset almost exactly by the smaller observed left/right pulse compression imbalance.

We can estimate what the absolute values of these errors are for a typical triangulation system. For example, the triangulation angle of our scanner hardware is 30^o, and the laser width is about 1 mm across a depth of 0.3 m. For this configuration, we would expect errors of 0.8 mm for a 10:1 reflectance step and 0.5 mm for a corner angle of 110^o.

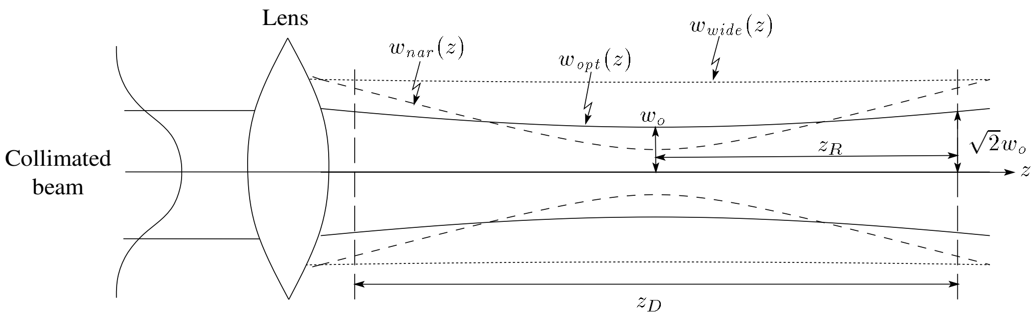

Figure 2.9: Gaussian beam optics. A collimated beam is focused by the lens to a width, w_o, and spans a desired depth of field z_D. Bringing the beam into tight focus results in a rapidly expanding beam ($w_{nar}(z)$). A wide beam expands slowly ($w_{wide}(z)$), but may be too wide over the desired depth of field. The optimal beam ($w_{opt}(z)$) expands only by a factor of $\sqrt{2}$ over the depth of field. For this optimal configuration, the distance between the beam waist and the $\sqrt{2}$ point is called the Rayleigh range, z_R.

2.2.3 Focusing the beam

One possible strategy for reducing the errors described in the previous section is to narrow the illuminating beam. Unfortunately, there are practical limitations to implementing this strategy. First, in order to take advantage of the smaller beam width, the sensor resolution must increase. This leads to expensive sensors. Second, using lenses to focus a Gaussian beam arbitrarily over the field of view is not possible due to diffraction limits. The best candidate for a tightly focused illuminant over a large depth of field is a laser beam. The irradiance profile of a laser beam is typically Gaussian as described by Equation 2.2, and the width of the beam expands from its narrowest point as [Siegman 1986]:

$$w(z) = w_o\sqrt{1 + \left(\frac{\lambda z}{\pi w_o^2}\right)^2}$$

where w_o is the beam waist or width of the beam at its narrowest point, z is the distance from this point, and λ is the wavelength of the laser. Figure 2.9 shows how the beam expands for various beam waists. Note that as the waist becomes narrower, the beam expands faster. For a given depth of field, z_D, we want the beam to be as tight as possible to limit the effects of triangulation errors. If the waist is too narrow, then the beam expands too rapidly and becomes too large at the edges of the field of view. If the waist is too large, it expands slowly, but is still too large at the edges of the field of view. The optimum beam width is attained when the waist obeys the relation:

$$w_o = \sqrt{\frac{\lambda z_D}{2\pi}}$$

In this case, the beam reaches a maximum width of $w = \sqrt{2}w_o$ and the value $z_R = z_D/2$ corresponds to what is known as the Rayleigh range of the beam. The variation in width is thus:

$$\sqrt{\frac{\lambda z_D}{2\pi}} \leq w \leq \sqrt{\frac{\lambda z_D}{\pi}}$$

The best width to range ratio at the edges of the field of view is:

$$\frac{w}{z_D} = \sqrt{\frac{\lambda}{\pi z_D}} \tag{2.3}$$

The discussion of range errors due to reflectance and shape variations in the previous section shows that the errors of an optical triangulation system are on the order of the width of the illuminant. Thus, the beam width to range ratio is a measure of the relative accuracy of the system. Equation 2.3 tells us that once we have selected a range of interest, diffraction limits impose a bound on the relative accuracy in the presence of reflectance and shape variations.

Note that Bickel et al. [1985] have explored the use of axicons (e.g., glass cones and other surfaces of revolution) to attain tighter focus of a Gaussian beam. The refracted beam, however, has a zeroth order Bessel function cross-section; i.e., it has numerous side-lobes of non-negligible irradiance. The influence of these side-lobes is not well-documented and would seem to complicate triangulation. In addition, once the sensor resolution is fixed, then arbitrarily narrowing the beam actually becomes undesirable. If the image of the beam spans only one or two pixels, then the computed mean will not attain sub-pixel accuracy. See Section 3.5 for a discussion of the beam width in a signal processing context.

2.2.4 Influence of laser speckle

Even under the most ideal circumstances of scanning a planar object with no reflectance variations, the accuracy of optical triangulation methods using coherent illumination are still limited by an unavoidable interference phenomenon: laser speckle. In this section, we describe how laser speckle arises and how it results in range errors.

In order to analyze the effects of laser speckle, we must consider the wave nature of light and, in particular, the effects of diffraction. One of the most powerful tools for studying these effects is scalar diffraction theory. Scalar diffraction theory treats light as a scalar field rather than as a coupled electric and magnetic field. Experiment has shown that this approximation is valid as long as (1) the diffracting aperture is large compared to a wavelength and (2) the diffracted fields are not being observed too close to the aperture. Both of these criteria are satisfied in the current context.

Throughout this section we adopt the notation of Goodman [1968]. For a monochromatic wave, the scalar field at a point in space may be written:

$$u(\mathbf{x}, t) = U(\mathbf{x}) \cos[2\pi\nu t + \phi(\mathbf{x})]$$

where $U(\mathbf{x})$ and $\phi(\mathbf{x})$ are the amplitude and phase of the wave at position $\mathbf{x}$, and ν is the optical frequency. It proves convenient to re-write this expression as:

$$u(\mathbf{x}, t) = \mathrm{Re}[\mathbf{U}(\mathbf{x}) \exp(-i2\pi\nu t)]$$

where $\mathbf{U}(\mathbf{x})$ is a complex function (called a phasor):

$$\mathbf{U}(\mathbf{x}) = U(\mathbf{x}) \exp[-i\phi(\mathbf{x})]$$

and $Re[]$ is the real part of its argument. When solving problems in scalar diffraction theory, it suffices to represent the field with $\mathbf{U}(\mathbf{x})$, because the time dependence is the same everywhere and is known a priori. Researchers sometimes refer to $\mathbf{U}(\mathbf{x})$ as the field "amplitude", though it should be remembered that it is a complex quantity. Note that when taking physical measurements, we evaluate the intensity of a field, given by:

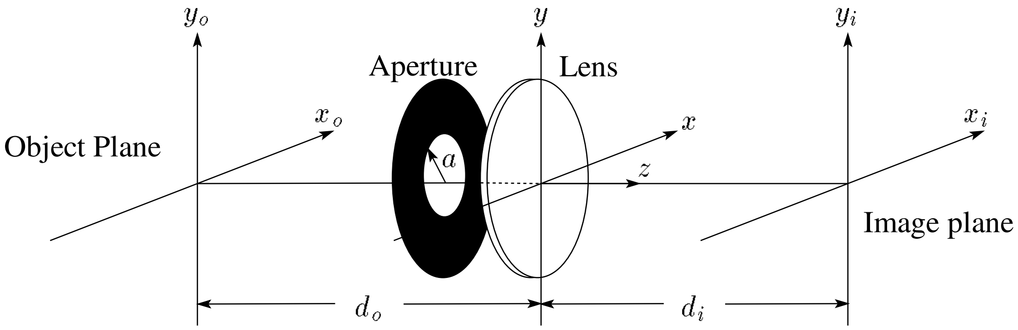

Figure 2.10: Imaging geometry and coordinate systems.

$$I(\mathbf{x}) = |\mathbf{U}(\mathbf{x})|^2 = \mathbf{U}(\mathbf{x})\mathbf{U}^*(\mathbf{x})$$

Now consider the imaging geometry shown in Figure 2.10. A striking result from scalar diffraction theory is the fact that the image formed by a lens is the same as the image described by geometrical optics, filtered by an amplitude spread function (also known as a point spread function):

$$\mathbf{U}_i(x_i, y_i) = \mathbf{U}_g(x_i, y_i) * h(x_i, y_i) \tag{2.4}$$

where $\mathbf{U}_i$ is the resulting image field, and $\mathbf{U}_g$ is the geometrical optics image, which is in turn given by the relation:

$$\mathbf{U}_g(x_i, y_i) = \frac{1}{M}\mathbf{U}_o\left(-\frac{x_i}{M}, -\frac{y_i}{M}\right) \tag{2.5}$$

where $\mathbf{U}_o$ is the field in the object plane, and M is the magnification factor. The amplitude spread function is simply the Fourier transform of the lens pupil:

$$h(x_i, y_i) = \iint_{-\infty}^{\infty} P(\lambda d_i s_x, \lambda d_i s_y) \exp[-i2\pi(x_i s_x + y_i s_y)] ds_x ds_y \tag{2.6}$$

where $P(x, y)$ describes the shape of the pupil (usually a circular disc), d_i is the distance to the image, and λ is the optical wavelength. For a circular lens aperture, the pupil function is unity with a radius of a and zero outside. The amplitude spread function for such an aperture

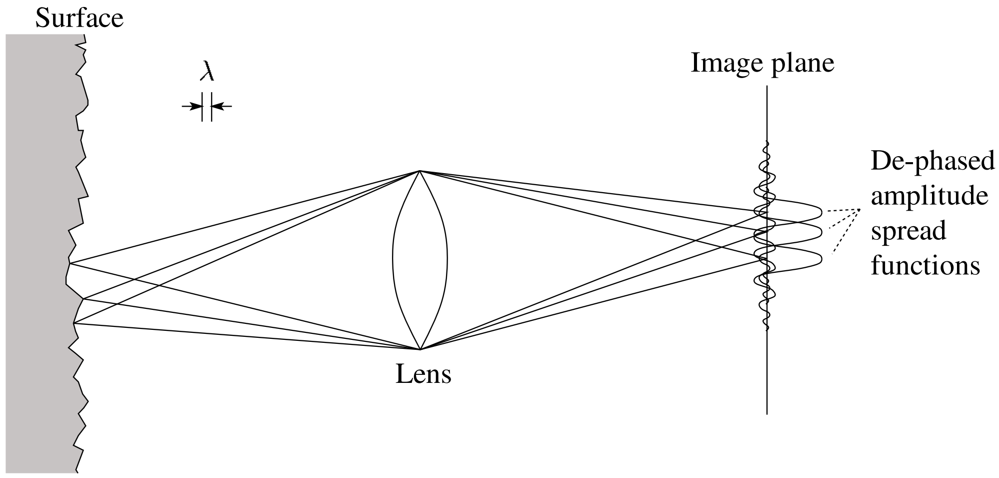

Figure 2.11: Speckle due to diffraction through a lens. The image of the surface is filtered by the amplitude spread function. Because the surface is rough compared to a wavelength, the image contains noisy phase terms that influence nearby points on the surface through the amplitude spread function. The result is laser speckle. Source: [Goodman 1984].

resembles a circularly symmetric "sinc" function[2], where the first zero crossing is at:

$$r \approx 0.61 \frac{\lambda d_i}{a}$$

Note that the imaging equations (2.4 - 2.6) relate object and image *fields*, not intensities. The intensities must be computed by finding the square magnitude of the fields.

Speckle arises in an image when the surface is rough compared to the optical wavelength. The result is random variations in the relative phases of imaged points. If the imaging system were to exhibit no diffraction effects (i.e., if the amplitude spread function were an impulse function), then these phases would have no impact, because they would vanish when computing the intensities. However, real imaging systems do exhibit diffraction effects, and the amplitude spread function serves to distribute the random variations across the image plane. The result is laser speckle. Figure 2.11 illustrates the this phenomenon.

In the remainder of this section, we attempt to quantify the effect of laser speckle on

[2] $\mathrm{sinc}(x) \equiv \frac{\sin(\pi x)}{\pi x}$

range accuracy. To do so, we consider the case of a two-dimensional world, which permits us to remove references to the y coordinate in the equations above. This simplification will aid in the discussion of the issues of laser speckle, but will not significantly alter the basic results. Note that in this two-dimensional world, the amplitude spread function for an imaging aperture is simply the "sinc" function. In addition, we will treat the triangulation imaging geometry as though it were an axially aligned imaging geometry. This approximation introduces some error, because some of the imaged light originates from regions outside of the plane in focus (the center of the illuminant), but for narrow illuminants, this error is negligible [Francon 1979].

For a Gaussian illuminant incident on the surface of an object, the field at the surface of the object takes the form:

$$\mathbf{U}_o(x_o) = G_o(x_o)\exp[i\phi_o(x_o)] \tag{2.7}$$

where $G_o(x_o)$ is the Gaussian illuminant and $\phi_o(x_o)$ represents the random phase variations introduced by the roughness of the surface. We can now write the field at the image plane as:

$$\mathbf{U}_i(x_i) = \{G_g(x_i)\exp[i\phi_g(x_i)]\} * h(x_i)$$

where $G_g(x_i)$ and $\phi_g(x_i)$ are the images of the Gaussian and and phase terms, respectively, as given by geometric optics. From this equation, we can compute the intensity measured at the image plane:

$$I_i(x_i) = \mathbf{U}_i(x_i)\mathbf{U}_i^*(x_i) \tag{2.8}$$

Equations 2.7 - 2.8 describe the physical image of a Gaussian illuminant after experiencing random phase variations and diffraction "filtering." Errors in determining the mean of this image result in range errors in triangulation systems using traditional analysis. The mean is defined by:

$$c_i = \frac{\int_{-\infty}^{\infty} x_i I_i(x_i) dx_i}{\int_{-\infty}^{\infty} I_i(x_i) dx_i} \tag{2.9}$$

For a centered Gaussian, this mean should be exactly zero; however, the random phase variations introduce statistical variations in the mean that we can quantify as:

$$\sigma_{c_i} = [\langle c_i^2 \rangle - \langle c_i \rangle^2]^{\frac{1}{2}} = \langle c_i^2 \rangle^{\frac{1}{2}} \tag{2.10}$$

where the expected value of the mean is assumed to be zero. Figure 2.6 illustrates the influence of laser speckle in a triangulation system. The error in the mean maps to a range error through the relation:

$$\sigma_r = \frac{M\sigma_{c_i}}{\sin\theta}$$

In other words, once we have computed the variations in detecting the mean at the sensor, we must map it into a range error through multiplication by the imaging system's magnification factor, M, and division by the triangulation factor, $\sin\theta$.

Computing the standard deviation described by Equations 2.9 and 2.10 analytically has proven to be extremely difficult. Instead, we have performed numerical simulations. These simulations require some characterization of the phase function, $\phi_o(x_o)$, which describes the roughness of the surface. Researchers typically model this function to be proportional to the height profile, $\beta(x_o)$ of the surface:

$$\phi_o(x_o) = \frac{2\pi}{\lambda}\beta(x_o)$$

This $\beta(x_o)$ is typically taken to be a random process with Gaussian statistics at each point on the surface and a Gaussian autocorrelation spectrum. The pointwise statistics are thus parameterized by the surface height variance, R_q, also called the roughness, and the correlation length, α [Fujii et al. 1976]. As a result, the parameter space we must explore is four dimensional: a function of the width of the laser, the width of the point spread function, the surface roughness and the correlation length. To date, we have thoroughly studied the deviation of the mean for surfaces rough compared to a wavelength with zero correlation length, and we have arrived at a qualitative understanding of the behavior when the surface roughness parameters are varied. For uncorrelated rough surfaces, the variation of the mean follows the relation:

$$\sigma_{c_i} \approx \frac{1}{2}\sqrt{\frac{w_i \lambda d_i}{2a}}$$

where w_i is the width of the image of the Gaussian illuminant, d_i is the distance from the lens to the aperture, and a is the aperture radius. Mapping this into a range error we obtain:

$$\sigma_r \approx \frac{M}{2\sin\theta}\sqrt{\frac{w_i \lambda d_i}{2a}}$$

Noting that $w = Mw_i$ is the width of the illuminant in object space, and $d_o = Md_i$ is the distance to the object, the range error can be re-written as:

$$\sigma_r \approx \frac{1}{2\sin\theta}\sqrt{\frac{w \lambda d_o}{2a}} \tag{2.11}$$

This relation yields several insights. First of all, the range error is fundamentally independent of the resolution of the sensor; it is an inescapable consequence of the speckle phenomenon. Secondly, we expect that narrowing the aperture will increase the speckle noise and thereby increase the range error. Finally, the range error varies with the square root of the width of the laser as seen from the sensor. Narrowing the laser ought to decrease range noise.

Numerical experiments with various surface roughness and correlation length parameters have yielded some quantitative results that follow qualitative expectations. In summary, as the correlation length grows, we observe a decrease in the computed range error. Similarly, as the surface roughness decreases, the range error also decreases. This result is not surprising; the limiting case of a perfectly smooth planar surface should result in no laser speckle and thus no range error. Increasing the correlation length and decreasing the roughness bring the results closer to the perfectly smooth case.

To give a sense for the size of the range errors, we can substitute some typical values into Equation 2.11. The triangulation system used for experiments in Chapter 4 has a 30 degree triangulation angle, an aperture radius of about 5 mm, a stand-off of 1 meter and uses a HeNe laser source (λ = 632.8 nm). We would therefore expect the range error due to reflections from a "perfectly" rough surface to be approximately 0.25 mm. This error corresponds to roughly half the resolution of the sensor being used. Real surfaces, however, are not perfectly rough, so we would expect smaller range errors in practice.

2.3 Prior work on triangulation error

Several researchers have observed the limitations in triangulation accuracy due to reflectance and shape variations [Buzinski et al. 1992] [Soucy et al. 1990] [Mundy & Porter 1987]. In addition, researchers have studied the effect of laser speckle on range determination and have indicated that speckle is a fundamental limit to the accuracy of laser range triangulation, as described in the previous section, though its effects can be reduced with well-known speckle reduction techniques [Baribeau & Rioux 1991] [Dorsch et al. 1994]. These speckle reduction methods, however, typically require spatial averaging, resulting in a trade-off between noise and spatial resolution. Using a less coherent light source appears to be a better alternative, as it leads to less speckle noise without loss of resolution. Regardless, changing the viewpoint generally results in a different speckle pattern, so multiple range images acquired from different viewpoints can be combined to reduce noise without loss of resolution. We demonstrate this last idea in Section 8.3.

Mundy & Porter [1987] attempt to correct for variations in surface reflectance by noting that two imaged pulses, differing in position or wavelength are sufficient to overcome the reflectance errors on a planar surface, though some restrictive assumptions are necessary for the case of differing wavelengths, and they do not address errors due to shape variations. Kanade et al. [1991] describe a rangefinder that finds peaks in time for a stationary sensor with pixels that view fixed points on an object as the illuminant sweeps across the scene. This method is very similar to the spacetime analysis presented in this thesis for solving some of the problems of optical triangulation; however, the authors in [Kanade et al. 1991] do not indicate awareness that their design addresses these problems. Further, we show that the principle generalizes to other scanning geometries.