

3D measurements from imaging laser radars: how good are they?

Martial Hebert and Eric Krotkov

In this paper we analyse a class of imaging range finders – amplitude-modulated continuous-wave laser radars – in the context of computer vision and robotics. The analysis develops measurement models from the fundamental principles of laser radar operation, and identifies the nature and cause of key problems that affect measurements from this class of sensors. We classify the problems as fundamental (e.g. related to the signal-to-noise ratio), as architectural (e.g. limited by encoding distance by angles in $[0, 2\pi]$), and as artifacts of particular hardware implementations (e.g. insufficient temperature compensation). Experimental results from two different devices – scanning laser rangefinders designed for autonomous navigation – illustrate and support the analysis.

Keywords: range finders, mobile robots, CW lasers, calibration, range data

Range sensing is a crucial component of any autonomous system. It is the only way to provide the system with three-dimensional representations of its environment. The classical computer vision approach to range sensing is to use passive techniques such as stereo-vision, or shape from X. However, those techniques are not yet sufficiently reliable or fast to be used in many applications, most notably real-time robotic systems. Active sensors, which generate the illumination instead of using only the ambient illumination, have received increasing attention as a viable alternative to passive sensors. Their importance was recognized relatively early. For example, Nitzan *et al.*¹ describe a system for interpreting indoor scenes that uses range and intensity from a laser ranging system. Many such sensors were developed² and many have been used in real computer vision and robotics applica-

tions such as obstacle avoidance^{3,4} and autonomous navigation⁵ of mobile robots. Surveys of rangefinding sensors and their applications in robotics can be found elsewhere⁶⁻⁸. A review of their use in autonomous navigation of mobile robots can be found in Hebert *et al.*⁹.

Although significant theoretical results have been derived in the area of laser radar characterization¹⁰, experimental work is needed to evaluate performance in robotics applications. In this paper, we characterize and evaluate a class of sensors, the imaging laser radars*. Those sensors have been proposed as a good compromise between accuracy, resolution, and speed requirements, especially in the context of mobile robotics. Our intent is to present measurement and noise models for those sensors, to identify problems that are specific to this class, and to provide experimental data to support our conclusions. Our emphasis is on identifying limitations and capabilities that have an impact on the use of standard image analysis algorithms for those sensors.

We concentrate on laser radars because of our experience with two such sensors, Erim and Perceptron, in our work in mobile robotics. The theoretical and experimental results presented in this paper use in part results from earlier analyses for the Erim sensor¹¹⁻¹⁴, and for the Perceptron sensor^{15,16}. However, we suggest that more analyses of this type are needed in order to better grasp the state of sensor technology from the point of view of computer vision and robotics.

The paper is organized in three parts. The principle of the sensors, are described along with the theoretical models of noise and measurement geometry, and the two sensors that we use in our experiments. Some problems that are specific to this class of sensors are then described. Those problems can significantly impact the quality of the data and limit the use of those

The Robotics Institute, Carnegie Mellon University, Pittsburgh, PA 15213, USA

*An optical-wavelength radar is also called Lidar, which is an acronym for Light Detection And Ranging.

sensors. We distinguish between problems that are inherent to the physics of the laser radars, and problems specific to the hardware currently available for robotics applications. Experimental results obtained from actual sensors are given. The experiments illustrate the measurement models and problems introduced in the previous two sections.

SENSORS

In this section we address imaging laser radars, covering both their principles and practical characteristics. First, we describe the general principle of operation, define a sensor reference frame, and present measurement models for the range and intensity data. Then, we describe two particular sensors that we used for experimentation.

Principle of operation

The basic principle of a laser radar is to measure the time between transmitting a laser beam and receiving its reflection from a target surface. Three different techniques can be employed to measure the time of flight, which is proportional to the range: pulse detection, which measures the time of flight of discrete pulses; coherent detection, which measures the time of flight indirectly by measuring the beat frequency of a frequency-modulated continuous-wave (fm-cw) emitted beam and its reflection; direct detection, which measures the time of flight indirectly by measuring the shift in phase between an amplitude-modulated continuous-wave (am-cw) emitted beam and its reflection.

Experimental devices have been developed using both pulse and coherent detection technologies (for a survey, see Besl²). They are not yet widely in use in computer vision and robotics applications. In this paper, we concentrate on am-cw laser radars.

For am-cw laser radars, the range to a target is proportional to the difference of phase; if $\Delta\phi$ is the difference of phase, then the range is $r = \frac{\lambda}{4\pi} \Delta\phi$, where λ is the wavelength of the modulation. Since the phase is defined modulo 2π , the range is defined modulo r_a , where $r_a = \lambda/2$ is the distance (or *ambiguity interval*) corresponding to the maximum phase difference of 2π . Therefore, an inherent limitation of this principle of operation is that it cannot measure range uniquely, i.e. it measures range only within an ambiguity interval. In practice, it is not possible to distinguish between ranges r and $r + r_a$ without employing external constraints, or two beams with different modulation frequencies.

We digress briefly to consider constraints imposed by autonomous navigation. First, the ambiguity interval should equal or exceed the maximum range of interest. Typically, this maximum range depends on vehicle velocity and the type of terrain. For example, for a maximum range of 20 m, $r_a = 40$ m and $\lambda = 40$ m. Second, safety considerations limit the power and wavelength of the laser diode used to generate the laser beam. The power of the diode is typically 100–300 mW, and the wavelength of the carrier signal is typically 700–1000 nm (near-infrared).

For many applications, *imaging* laser radars are

essential. Imaging sensors generate a dense set of points structured as an image. Typically, image generation is achieved by two mechanically controlled mirrors that raster-scan the beam across a scene, measuring the range at regularly sampled points. Note that regular space in image space generally causes irregular space in the scene. For real-time applications such as autonomous navigation, the image acquisition time is limited. The combination of electro-mechanical parts and time constraints adds new complexity and new sources of errors that do not exist in non-imaging sensors, such as surface probes and surveying devices.

In addition to range, am-cw scanners measure the 'strength' of the reflected beam, thus generating a second image that some call the *reflectance* image. To avoid confusion with surface reflectance, we will refer to it as the *intensity* image. This image is similar to a TV camera intensity image, but is registered with the range image, and does not depend on the ambient illumination. Although relatively little work has been done using Lidar intensity images, we have found them to be useful for automatically determining the position and orientation of the sensor with respect to a walking robot¹², for tracking roads from a robot truck⁹ and for object recognition¹⁷.

Sensor reference frame

It is useful to convert the range pixels to points in space expressed with respect to a reference frame. In this section we define a standard reference frame attached to the scanner, and the relation between a pixel (row, column, range) and the coordinates of the corresponding point.

Figure 1 illustrates the reference frame. As shown, the y-axis coincides with the direction of travel of the laser beam projected through the central point of the scanner (i.e. the principal ray). The angle θ (azimuth) corresponds to a rotation about the z-axis. The angle ϕ (elevation) corresponds to a rotation about the x-axis.

Given the sensor measurement (u, v, d) (i.e. row, column, range), the transformation to spherical-polar




Figure 1. Reference frame for scanning laser range finder. R : denotes range. To be consistent with the majority of the text, we should denote range by r

coordinates is:

$$\varphi = u\Delta_\varphi + \varphi_0, \quad \theta = v\Delta_\theta + \theta_0, \quad r = \Delta_R d \quad (1)$$

where:

Δ_φ is the angular increment, in degrees/row, of the nodding mirror,
 Δ_θ is the angular increment, in degrees/column, of the panning mirror,
 φ_0 is the initial orientation, in degrees, of the nodding mirror,
 θ_0 is the initial orientation, in degrees, of the panning mirror, and
 Δ_R is the scanner range resolution in metres/grey-level.

Given the spherical polar coordinates φ , θ , r , the transformation to Cartesian coordinates is given by:

$$x = r \sin \theta, \quad y = r \cos \theta \cos \varphi, \quad z = r \cos \theta \sin \varphi \quad (2)$$

Measurement models

An am-cw range sensor would approach perfection if it emitted a zero-width laser beam and observed the returned signal through an infinitely small receiver. In reality, the beam subtends a non-zero angle, and the receiver detects signals subtended by a solid angle we will call the instantaneous field of view (IFOV). Assuming a circular field stop, the beam projects to an ellipse on the target surface, the *footprint* of the beam. Every point within the intersection of the footprint and the IFOV contributes a range value and an intensity value to the final range and intensity measurements.

We may model a (range, intensity) pair as a complex number z , or equivalently as a vector (see Figure 2a). The phase of z represents the sensed range (or phase shift), and the magnitude of z is the sensed intensity. According to this model, the range measured at a pixel is the integral of $k_0 z$ over the IFOV of the receiver, where k_0 is scalar and depends on parameters of the instrument (transmitted power, and the electro-optics of the receiver) and parameters of the environment (the angle α between the surface normal and the direction of measurement, the reflectance ρ of the target surface, and the range r as shown in Figure 2b).

Ultimately, the fidelity of the range measurement depends on the power of the signal reaching the photodetector, which in turn depends on k_0 . More precisely, the time-average radiant flux \bar{F}_P (in watts) reaching the photodetector can be shown¹⁰ to be:

$$\bar{F}_P = k_1 \frac{\rho \cos \alpha}{r^2} \quad (3)$$



Figure 2. Model of (range, intensity) pair as complex number

where k_1 is a function of the transmitted radiant flux, the capture area of the receiver, and filter bandwidths. Assuming that the output power of the photodetector is proportional to \bar{F}_P , the output signal is proportional to $\rho \cos \alpha / r^2$.

When the received power is small, the output signal is small, and noise is significant. There are many sources of noise, including photon noise, laser noise, ambient noise, dark-current noise, secondary emission noise, and subsequent amplifier noise.

Nitzan *et al.*¹ identify photon noise as the dominant source, and assume that photoemission is a Poisson process to obtain the signal-to-noise ratio:

$$\text{SNR} = k_2 \left(\frac{\rho \cos \alpha}{r^2} \right)^{1/2} \quad (4)$$

Assuming a modulation index of 100%, they go on to derive the approximate effect of photon noise on the measured range value. Their analysis indicates that the standard deviation of the range can be estimated by:

$$\sigma_r \approx \frac{R_a}{\sqrt{2\pi/\text{SNR}}} \quad (5)$$

Combining (4) and (5) yields:

$$\sigma_r \approx \left(\frac{r_a}{\sqrt{2\pi k_2}} \right) \left(\frac{r^2}{\rho \cos \alpha} \right)^{1/2} \quad (6)$$

where the first term depends only on the physical, optical, and electronic characteristics of the sensor, and the second term depends only on the observed scene.

This equation states that σ_r is approximately linear in r , or equivalently, that the variance σ_r^2 is approximately quadratic in r . Our experimental results follow this model. In particular, Figure 12 shows that σ_r^2 is quadratic in r , consistent with (6).

Examples

We used two sensors for experiments: Erim and Perceptron, whose geometric parameters are listed in Table 1. Other sensors based on the same principle exist^{2,18}. The two sensors that we use are typical of the operation and performance of this type of sensor.

The Erim scanning laser rangefinder is designed for applications in outdoor autonomous navigation. Several versions of the scanner exist. We refer to the version used for research on autonomous land

Table 1. Nominal values of sensor parameters

Parameter	Units	Erim	Perceptron	Description
φ_{fov}	deg	30	60	Vertical FOV
θ_{fov}	deg	80	60	Horizontal FOV
N_{rows}	pixel	64	256	Rows
N_{cols}	pixel	256	256	Columns
Δ_φ	deg	0.47	0.24	Vertical step ¹
Δ_θ	deg	0.32	0.24	Horizontal step ²
r_a	—	64 ft	40 m	Ambiguity interval
N	bit	8	12	Number of bits/pixel
Δ_R	—	3.0 in	0.98 cm	Range unit ³

$${}^1\Delta_\varphi = \frac{\varphi_{fov}}{N_{rows} - 1}, \quad {}^2\Delta_\theta = \frac{\theta_{fov}}{N_{cols} - 1}, \quad {}^3\Delta_R = \frac{r_a}{2^{N_{range}} - 1}$$




Figure 3. Erim intensity (top) and range images. The scene contains a tree (visible to the left), and a person (visible in the upper centre) on a path

navigation^{3, 9}, which is the successor of a sensor used for legged locomotion¹⁹.

The Erim scanner used a 100 mW laser diode operating at 830 nm. The sensor volume is roughly 90×90 cm, and the mass is about 45 kg. The scanning mechanism consists of a vertical rotating mirror for horizontal scanning, and a horizontal nodding mirror for vertical scanning. The acquisition rate is 0.5 s for a $64 \times 256 \times 8$ bit image. Table 1 summarizes the characteristics of the sensor. Figure 3 shows a range image (bottom) and a reflectance image (top) from the Erim scanner. The images show an outdoor scene of a road surrounded by a few trees. The sharp discontinuity at the top of the range image is due to the ambiguity interval of about 20 m.

The Perceptron scanning laser range finder is more recent than the Erim, and has higher resolution and bandwidth. It is currently used for terrain mapping in support of legged locomotion²⁰.

The sensor volume is roughly $50 \times 45 \times 35$ cm, and the mass is about 30 kg. The sensor uses a 180 mW laser diode operating at 810 nm. The image acquisition rate is 0.5 s for a $256 \times 256 \times 12$ bit image. The scanning mechanism is similar to the Erim design, except that the vertical field of view and the vertical image size are programmable. Table 1 summarizes the operating characteristics of the sensor as specified by the manufacturer¹⁶.

Figure 4 shows the eight most significant bits of a

typical range image (right) and intensity image (left) from the Perceptron. The distance between the scanner and the floor is approximately 4 m.

PROBLEMS

In this section we describe four effects that may lead to corrupted or degraded data. Two effects, mixed pixels and range/intensity crosstalk, are due to fundamental limitations of am-cw laser radars, although they are sometimes compounded with problems in the design of the actual sensors that we used. The two other effects, distortion due to scanning and range drift, are more specific to the particular sensors that we use. However, it is important to be aware of those effects since they do affect the quality of the data. Furthermore, simple cost-effective remedies do not seem to exist at the moment even though those problems are theoretically avoidable.

Mixed pixels

Significant problems occur at pixels that receive reflected energy from two surfaces separated by a large distance. From an image analysis point of view, we would like the range at such points to be measured on either of the surfaces, or at least to fall in between in some predictable way. The fact that the range is measured by integrating over the entire projected spot leads to a phenomenon known as *mixed pixels* in which the measured range can be anywhere along the line of sight. In practical terms, this means that occluding edges of scene objects are unreliable, and that phantom objects may appear due to mixed measurements that are far from the real surfaces. This is a problem inherent to direct detection am-cw laser radars and it cannot be completely eliminated.

Figure 5 shows the geometry of the measurement at an occluding edge: two objects at distances D_1 and D_2 from the sensor ($D_1 < D_2$) are separated by a distance D , and a point is measured at the edge between the two surfaces. Due to the angular width of the beam, the range is formed by integration over a spot that contains reflections from both surfaces. The relevant parameter is the ratio p between the spot surface due to the near




Figure 4. Perceptron intensity (left) and range images. One of the authors is visible in the intensity image standing in front of box-like targets used for calibrating the sensor. The images have been enhanced for printing




Figure 5. Surface 1 occludes surface 2 (wide view)

surface and the total area. The combination of the two measurements is better explained using a model in the complex plane like the one presented below. Each portion i of the spot generates a measurement that can be represented by a complex number z_i . The phase of z_i , φ_i is proportional to the range D_i , and its magnitude depends on the reflectivity ρ_i of the material. The resulting measurement is given by the sum $z = z_1 + z_2$. The finally measured range is proportional to the phase of z .

The phenomenon becomes clear with this formulation: the phase of z can be anywhere between φ_1 and φ_2 , depending on the ratio of the lengths of z_1 and z_2 which depends on p , ρ_1 , and ρ_2 . The effect of mixing on sensed range becomes even more unpredictable when D is greater than half the ambiguity interval of the sensor. In that case, the angle φ between z_1 and z_2 is greater than π , and the resulting measurement is such that $\varphi_2 < \varphi < 2\pi$ or $0 < \varphi < \varphi_1$, depending on p and the reflectivities. In the first case, the range may be anywhere *behind* the furthest object. In the second case, the range may be anywhere *in front of* the nearest object. In practice, the latter situation occurs only in specific combinations of range and scene geometry. In practice, this phenomenon is observed as soon as there are discrete objects in the scene. Figure 6 shows the effect of mixed pixels in a real image. The top panel shows an Erim image with a small window (shown as a white rectangle) containing an object (a tree) and its left and right edges. The bottom panel shows an overhead view of the 3D points in this region as calculated by Equation (2), and using the range from the image pixels. The points at the centre of the distribution correspond to the smooth surface of the object. Away from the centre are two lines of mixed pixels that appear at the object's edges. In this example, all mixed pixels are located between the two targets, tree and background surface. This result is typical of the mixing effect in laser radars.




Figure 6. Mixed pixels in real image. Top: selected region in Erim image; bottom: overhead view of the corresponding points

The mixed pixel effect complicates the processing and interpretation of the range images. Its main consequence is that strong range edges are unreliable. One approach to the problem is to apply a median filter to the image. The filter removes most of the mixed pixels since they appear as spurious readings only along the direction of an edge. Another approach is to transform all the points to a 3D coordinate frame. There, mixed pixels appear as isolated points, which makes them easier to remove than in image space. Since these two approaches remove mixed pixels whose range is far from the true value, they will not remove those in the vicinity of the actual edge. Thus, it is not possible for them to remove reliably all the mixed pixels from the data.

Scanning pattern

With currently available commercial technology, imaging can be achieved only by scanning the beam using rotating and nodding mirrors. The scanning mechanism introduces additional errors into the sensor. They are probably the hardest to quantify and to correct. We identify three sources of error: synchronization, distortion, and localization. They result in a correct range measurement being stored at the wrong pixel in the image. Limitations due to the scanning mechanism are not inherent to the am-cw technology, but are due to the lack of alternatives to electro-mechanical scanning devices.

Synchronization

The main problem, *synchronization*, is due to the fact that three systems (horizontal mirrors, vertical mirrors, and range measuring system) must be synchronized exactly. In particular, the motion of the two mirrors must be exactly synchronized with the sampling of the range measurements. A small error in synchronization results in an error in either φ or θ , depending on which mirror is affected. Even though the range measurements themselves are not affected, the angular errors will propagate to the coordinates computed from equation (2). For example, poor synchronization may occur at the top of the image because the nodding mirror takes a finite amount of time to go from zero speed at its starting position to its normal scanning speed. During this interval of time, a few scanlines that are not correctly sampled are collected. Figure 7 shows a Perceptron image in which a rectangle in the scene projects to a skewed shape instead of a rectangle. This is due to the poor synchronization of the mirrors at the beginning of the scan. In this particular case, the problem can be fixed easily by starting the data acquisition a few lines below the starting position of the



Figure 7. Synchronization error in Perceptron image. The black wire-frame rectangle marks the correct position

nodding mirror. In general, it is not possible to perfectly synchronize the mirrors. In general, there is a discrepancy between the nominal values of the scanning angles and the actual values. This error is difficult to quantify. We describe an experimental setup for estimating the angular error distribution later.

The distribution of the scanning errors due to poor synchronization remains relatively constant as long as the sensor is stationary. The effect of them on the image may be predicted, and appropriate corrections can be made by simple processing of the range images. For example, ignoring the first few rows would eliminate most of the unreliable measurements. However, additional image *distortion* may occur in some applications. In autonomous navigation applications, for example, the sensor may be subject to motion, shocks and vibrations. This may produce large errors in scanning angles because of the mirrors getting out of phase. Those errors are hard to quantify. This effect is unavoidable except through the use of an external stabilization device.

The distribution of the scanning errors due to poor synchronization remains relatively constant as long as the sensor is stationary. However, additional image *distortion* may occur in some applications in which the scanner is subject to external motion.

Localization

Even with perfect synchronization and no distortion due to external motion, sensor geometry may be incorrect if one uses a simplified model of the scanning mechanism. The ideal geometric model given above assumes implicitly that all the directions of measurements originate from a single point. This is equivalent to saying that the mirrors are infinitely small and are located exactly at the same point. Some dimension and do not coincide in space. Therefore, the measurement directions do not intersect at a single point. In reality, the measurements intersect along a line instead of intersecting at a point. This is a problem of *localization* of the origin of the sensor which translates into a systematic error in the conversion from (r, θ, ϕ) to (x, y, z) . It could be solved by modelling the scanning mechanism or by calibration. For medium-range scanners such as Perceptron and Erim, the localization error is small compared to the errors on range and scanning angles.

Range/intensity crosstalk

Ideally, we would like the range measurements to be completely independent of the reflective properties of the observed object. Unfortunately, they do influence the range measurements and can even render range useless in some cases. This *crosstalk* effect between range and intensity is due to a number of causes.

The first cause for the crosstalk effect is a fundamental property of direct am-cw range measurement. The standard deviation given by equation (6) depends on the reflectance of the observed material. Roughly speaking, the lower the intensity, the higher the range noise. This affects only the variance of the measurement, not its mean value.

Another source of crosstalk is in the implementation of the detection electronics. Typically, the receiver electronics operate optimally only in a narrow range of intensities compared to the large dynamic range of

intensities that can be observed. As a result, surfaces that reflect intensities outside of the optimal operating range will produce noisy or even erroneous range measurements. This effect can be reduced by dynamically adjusting the operating range according to the intensity. The Perceptron scanner implements such a solution. However, the low intensities (dropout) and the high intensities (saturation) still produce spurious range readings. There are ways to increase the dynamic range of the receiver but they are not implemented in most scanners available to date.

The crosstalk effect becomes more noticeable at edges or on textured surfaces. In those cases, the beam illuminates a region that contains points of different surface reflectance. Since those points may generate slightly different range measurements, the situation is similar to the one discussed above, except that the mixing is due to the variation of reflectance within the spot instead of the variation of range. A consequence of this effect is that the behaviour of the range measurement at the edge between two surfaces with different reflectance may be unpredictable.

The crosstalk problem cannot be completely eliminated. However, some additions to the basic sensor design can diminish its effects. For example, the Perceptron scanner adjusts dynamically the operating range of the receiver, and uses a lookup table built using an off-line calibration procedure to correct the range as a function of intensity.

Figures 8–10 illustrate the crosstalk effect. In order to quantify the crosstalk effect, we designed an experiment in which a target with low reflectance is observed against a background of higher reflectance (see Figure 8). Considering one scanline in the range image, the dark target is located between columns 121




Figure 8. Experimental setup to study range/intensity crosstalk




Figure 9. Mean range and intensity for black and white targets. —: mean range (m); - - -: mean intensity (grey)




Figure 10. Variance of range and intensity for black and white targets. —: variance of range (cm²); - - -: variance of intensity (grey²)

and 135. We computed the mean and variance of the range and intensity distributions at each pixel in the scanline by taking 100 images of the scene. Figure 9 shows the mean values as a function of the column number. The mean intensity drops sharply at the edges of the black target and remains at a low level between them, as expected. The mean range remains roughly constant except for a sharp discontinuity at each edge. The reason is that the intensity from both materials is mixed at the edges, therefore the range is not properly corrected. Figure 10 shows the variance of the range and intensity distributions. This clearly shows a sharp increase in σ^2 between the high intensity background and the black target, as expected from the theoretical expression of range noise.

Range drift

We observed a significant drift of range measurements over time. To illustrate this effect, we placed a target 6 m from the origin of the Perceptron scanner and acquired one image per minute over 24 hours, during which the scene was static.

Figure 11 plots the sensed range at one target pixel. The figure shows a dramatic variation over time. Between hours 0 and 3, the ranges climb approximately one metre, as if the sensor were translating away from the target. After this four hour 'warm-up' period, the sensed ranges reach a plateau where they remain, with apparently random variations, for the rest of the day.

We hypothesized that some of the variations might be due to temperature changes. To test this hypothesis, we placed an electric heater directly behind the scanner, and repeated the above trial, acquiring images




Figure 11. Range measurements vary over time. The target is a sheet of cardboard 6 m in front of the scanner

at 2 Hz over two hours (14,000 images). Without the heater, the temperature was 21°C, and we observed approximately constant range measurements. We turned on the heater and after 30 minutes the temperature climbed to 45°C. During this time, the sensed ranges fell about 40 cm, a substantial decline. When we turned off the heater, the sensed ranges gradually increased until they regained their original level. This demonstrates conclusively that the temperature changes cause the distribution of range values to translate by significant amounts.

It is clear that heat cannot directly affect the phase shift which carries the range information. Therefore, the drift observed in those experiments is due to poor temperature compensation in the electronics used in the scanner. Improvements in the electronics have reduced this effect considerably. An important lesson is that such effects may dominate the errors due to the physics of the measurements.

ACCURACY AND PRECISION

We have introduced a theoretical framework that leads to a characterization of expected sensor accuracy for am-cw laser radars. However, it is important to verify that the sensors do indeed follow the theoretical model. In particular, real sensors include effects such as those described above that are not part of the theoretical framework. Actual sensor accuracy is important in determining what algorithms and what applications are appropriate for a given sensor. In this section we describe a series of experiments designed to measure range accuracy for the Erim and Perceptron sensors under different conditions and to compare it with the predicted theoretical values. Following Besl¹, we distinguish between *accuracy*, the difference between measured range and actual range, and *precision*, the variation of measured range to a given target. To separate the errors due to scanning and the errors due to actual range measurements, we distinguish between range precision and angular precision.

Accuracy

To determine the accuracy of the range measurements is to identify the distance between them and ground truth ranges. For a target point lying in the direction (φ, θ) , let $r_{\varphi, \theta}$ be the range measurement reported by the scanner, and let $d_{\varphi, \theta}$ be the true distance from the geometric origin, measured with a tape measure. Under ideal conditions, we expect to observe a linear relationship:

$$r_{\varphi, \theta} = ad_{\varphi, \theta} + r_0 \quad (7)$$

where r_0 is the *offset distance* from the origin to the (conceptual) surface corresponding to a range measurement of zero, and a is the slope.

To determine the parameters a and r_0 , we acquire range measurements of targets at six known distances between 6 and 16 m, and fit a line to the data. We illustrate the results for the Perceptron scanner in Table 2, which shows the extracted parameters from five trials under different conditions. Because of range drift (c.f. above), we do not assign high confidence to the

Table 2. Accuracy results for different targets and lighting conditions. The table shows accuracy results for various combinations of targets (one untreated cardboard slab, one cardboard slab painted black, and a planar piece of wood) and lighting conditions (sunny, cloudy, with and without room lights)

	α (deg)	r_0 (m)	Error (m)
Black (sunny, lights)	44.59	-1.32	0.28
Black (cloudy, lights)	44.34	-0.73	0.20
Black (cloudy, no lights)	44.11	-0.81	0.29
Cardboard (cloudy, lights)	44.92	-1.11	0.12
Wood (cloudy, lights)	45.57	-1.26	0.09

particular slope, intercept, and rms error entries in the table.

Nevertheless, the variation with surface material and lighting conditions is obvious. This suggests that the accuracy of the scanner depends significantly on variables in equation (7), including surface material, ambient illumination, and temperature. It also suggest that the effect of those variables is amplified by the particular hardware used in those sensors, as described above. We conclude by remarking that preliminary analysis of this and other data suggests that the accuracy does not depend significantly on target distance.

Range precision

To determine the precision of the range measurements is to identify by how much repeated range measurements vary. Here, we quantify the precision as the standard deviation of a distribution of measurements.

We have conducted a number of experiments in which we take 100 images at each target position, and compute the standard deviation of the depth measurements. In the experiments, we have examined how precision changes as a function of ambient illumination conditions, surface material of the target, distance from the scanner to the target, and beam incidence angle at the target. In this section, we report on the effect of ambient illumination for the Perceptron scanner, and refer readers interested in the other properties to Appendix A of Kweon⁵.

To study the effect of ambient light on sensed range, we place a target (in this experiment, a cardboard slab painted black) at a known distance, take 100 images,




Figure 12. Range variance under different lighting conditions. □: sunny, room lights; ○: cloudy, room lights; ◇: cloudy, no lights

and compute the variance in range at particular pixels. We repeat this procedure for six target distances between 6 and 16m under different indoor lighting conditions.

Figure 12 plots the results, which show that the precision decreases with intensity of illumination. The results strongly support the conclusion that the brighter is the ambient light, the larger are the temporal variations in the range measurements. As in the case of accuracy, above, the effect of the variables in equation (6) is clearly visible in those experiments, and it is amplified by the particular hardware implementation used. The results also illustrate the dependence of precision of the square of the target distance (c.f. equation (6)).

Angular precision

To measure the angular precision of a scanner, or its repeatability in pixel position, we fix the scanner configuration and scene, and take a series of images. Naturally, we expect a static point in the scene to project to one and only one pixel position in each of the images. We position the scanner in front of a vertical wall on which we have drawn white circles of radius 12cm, surrounded by black squares. We acquire a sequence of 10^4 images. From each intensity image we extract the white circle by thresholding, and then compute its centroid. We compute the standard deviations of these centroids for several row and column positions.

We find that for both scanners, the standard deviation of the centroid is an order of magnitude smaller than the nominal horizontal and vertical angle increments, varies little over time, and varies little over different pixels in the image. The Perceptron has significantly better angular precision than the Erim.

These findings suggest that the angular precision of the scanners is not a limiting factor. However, the experimental setting – stationary sensor and scene, gathering information over a region – represents a best case, for which the angular precision should be zero. Thus, the findings do not reveal the limitations introduced by relative sensor motion, and do not justify neglecting angular variations as a source of random disturbances in the measurement process.

DISCUSSION

In this paper we have examined in detail the 3D measurements supplied by amplitude-modulated laser radars. We presented measurement models for this class of sensors, identified problems unique to the technology, and problems specific to the implementation of sensors currently available to the robotics community, presented experimental performance results, and related our practical experience with the sensors.

How good are the three-dimensional measurements? In terms of speed and reliability for medium-range operations, we are not aware of any sensors with superior performance. So the short answer is that they are very good.

The special problems – mixed pixels, crosstalk, deviations from the scanning pattern, and range drift –

make it necessary to pre-process the images. For some problems (e.g. mixed pixels) no solutions exist. Thus, range data interpretation algorithms, like so many others in machine perception, must tolerate spurious data.

The quantitative performance, in terms of accuracy and precision, is highly variable. It depends on geometric factors such as incidence angle and target distance. We understand these reasonably well. But it also depends significantly on non-geometric factors such as temperature, ambient illumination, and surface material type. Even though detailed measurement models which include some of those variables have been developed, there is still a significant discrepancy between the theoretical sensor characteristics and the observed performance.

ACKNOWLEDGEMENTS

The authors wish to thank M. Blackwell, R. Hoffman, and I. Kweon who performed many of the experiments described in this paper and kept the scanners in working condition. This work was supported in part by NASA under Grant NAGW-1175 and by DARPA through ARPA No. 4976 monitored by the Air Force Avionics Laboratory under contract F33615-87-C-1499. The views and conclusions contained in this document are those of the authors and should not be interpreted as representing the policies of NASA, DARPA, or the US government.

REFERENCES

- 1 Nitzan, D, Brain, A and Duda, R 'The measurement and use of registered reflectance and range data in scene analysis', *Proc. IEEE*, Vol 65 No 2 (February 1977) pp 206-220
- 2 Besl, P 'Active, optical range imaging sensors', *Machine Vision & Applications*, Vol 1 (1988) pp 127-152
- 3 Dunlay, R and Morgenthaler, D 'Obstacle detection and avoidance from range data', *Proc. SPIE Mobile Robots Conf.*, Cambridge, MA (1986)
- 4 Veatch, P and Davis, L 'Efficient algorithms for obstacle detection using range data', *Comput. Vision, Graph. & Image Process.*, Vol 50 (1990) pp 50-74
- 5 Kweon, I *Modeling Rugged Terrain with Multiple Sensors*, PhD thesis, School of Computer Science, Carnegie Mellon University (January 1991)
- 6 Everett, H 'Survey of collision avoidance and ranging sensors for mobile robots', *Robotics & Autonomous Syst.*, Vol 5 (1989) pp 5-67
- 7 Jain, R and Jain, A *Analysis and Interpretation of Range Images*, Springer-Verlag, New York (1990)
- 8 Nitzan, D 'Assessment of robotic sensors', *Proc. Int. Conf. Robot Vision and Sensory Controls*, London, UK (April 1981) pp 1-11
- 9 Hebert, M, Kanade, T and Kweon, I *3-D Vision Techniques for Autonomous Vehicles*, Technical Report CMU-RI-TR-88-12, Robotics Institute, Carnegie Mellon University (August 1988)
- 10 Pratt, W *Laser Communication Systems*, Wiley, New York (1969)
- 11 Environmental Research Institute of Michigan, *Proc. Range and Reflectance Processing Workshop*, Ann Arbor, MI (1987, limited distribution)
- 12 Krotkov, E *Laser Rangefinder Calibration for a Walking Robot*, Technical Report CMU-RI-TR-90-30, Robotics Institute, Carnegie Mellon University (December 1990)
- 13 Watts, R, Pont, F and Zuk, D *Characterization of the ERIM/ALV Sensor - Range and Reflectance*, Technical Report, ERIM, Ann Arbor, MI (1987)
- 14 Zuk, D, Pont, F, Franklin, R and Larowe, V *A System for Autonomous Land Navigation*, Technical Report IR-85-540, ERIM, Ann Arbor, MI (1985)
- 15 Kweon, I, Hoffman, R and Krotkov, E *Experimental Characterization of the Perceptron Laser Rangefinder*, Technical Report CMU-RI-TR-91-1, Robotics Institute, Carnegie Mellon University (January 1991)
- 16 Perceptron, Inc., Farmington Hills, Michigan, *LIDAR Scanning System* (U.S. Patent Appl. No. 4226-00015)
- 17 Hebert, M and Kanade, T '3-D vision for outdoor navigation by an autonomous vehicle', *Proc. Image Understanding Workshop*, Cambridge, MA (1988)
- 18 Odetics, Inc., Anaheim, California *3D Laser Imaging System* (1989)
- 19 Zuk, D and Delleva, M *Three-Dimensional Vision System for the ASV*, Final Report No. 170400-3-f, DARPA, Defense Supply Service, Washington (January 1983)
- 20 Hebert, M, Krotkov, E and Kanade, T 'A Perception system for a planetary explorer', *Proc. IEEE Conf. on Decision and Control*, Tampa, FL (December 1989) pp 1151-1156