

PROCESSING SINGLE RANGE PROFILES FROM A WRIST MOUNTED LASER RANGE FINDER

Colin Archibald

Laboratory for Intelligent Systems
National Research Council of Canada
Ottawa, Ontario K1A 0R8

Abstract

Current robot vision systems consist of pseudo realtime feedback, where a scene is analysed while the robot is in a resting state, or otherwise occupied. For many applications a tight feedback loop with the robot servos being driven by sensor data is desirable. To this end, a lightweight, robot wrist mountable laser range finder has been developed at NRC. This sensor provides single scan lines of range data at a rate of approximately 13 calibrated profiles per second. Processing these scans to provide scene information is a unique problem and considerably different from intensity image processing, and should also be recognised as distinct from processing entire 3-D range images. Although there is only a single raster scan of data, explicit cues to the content of the 3-D scene are available. This paper describes a library of utilities written to extract information from the range profiles to be used in real time feedback control of the robot. Results of various operators, and combinations of operators, are presented. Potential applications and future directions are discussed.

KEYWORDS: robot vision, range data, laser range finder, profile interpretation, sparse range image.

1. Introduction

New sensors, and reasons for applying the sensors, are continually demanding innovative methods of extracting information about the real world from their signals. At NRC's Laboratory for Intelligent Systems we have taken the rather pragmatic approach of developing and applying sensors that provide a maximum of immediate, explicit 3-D information about the scene. The purpose of a laser range finder is to provide the location of a surface at many points, combined into a range image or a 3-D profile. A laser range profile scanner provides a single raster scan of surface information. The laser range finders developed here have been well documented elsewhere [1][2].

The surface information collected by the scanner is rarely perfect. In order to process imperfect data we need to understand the process by which it became imperfect. In intensity image processing this has been called the image formation process. Fortunately, there are only a few reasons why surface information may not be available from range data. For example, all light sensitive devices have a saturation point, and although the range finder returns depth, or range values, the electronics compute that measurement based on the position and amount of light reflected from the surface. Immediately, here are 2 reasons why no measurement would be available - too much, or too little light reflected from the scene. Secondly,

the shadow effect phenomenon is inherent to triangulation based range finding technology [3]. The shadow effect occurs when either the projected laser or the sensor is occluded from a point which the other is not occluded from. For example, if the laser is illuminating a surface point, but that point is occluded from the sensor, no range data is available at this point. These aspects of the sensors capabilities leave gaps in the range data. The gaps are not merely errors in the data, it is known that no data is available at these points for one of the reasons mentioned above. It is then necessary to determine what should be placed in these gaps. Operators using large neighbourhoods, and global operators are not usefully applied to data with gaps. A typical method of filling gaps in range data has been to interpolate between the valid data points. This is acceptable for range finders with a short depth of field, but very detrimental to large depth of field range finders where it is possible to have shadow effects covering one third of the scan or more. Although the perspective distortion appears to be corrected in the calibration, there are some imperfections from a true orthographic projection, making the interpretation of the profiles slightly more challenging.

Once the gaps in the data are 'repaired', there are a variety of common operators that can be used to extract features from the range profiles. Several operators have been experimented with and current results are described in section 3. In section 2 the sensor is described briefly, along with the other apparatus used thus far. Conclusions, potential applications of this technology, and future research toward these applications are discussed in section 4.

2. Apparatus

The laser range finder that we are currently using is an instance of those described in detail in [5]. Some information specific to our device is included here. The sensor was designed specifically for large depth of field use, from 10cm to 1 meter, with best accuracy achieved at the shorter range. To date we have calibrated the device to return reliable data in the range of 11cm. to 70cm. measured from a reference point on the frame of the sensor. Fig. 1 shows the mechanical design of the range finder. A HeNe laser source arrives via an optical fiber. There are 3 fixed mirrors and one double faceted scanning mirror on the path of the laser from optical fiber to CCD. The double faceted mirror scans the scene driven by a galvanometer, over a 40 degree field of view. This means that the field of view is approximately 10 cm in the direction of the scan (we will call this the X direction) at a standoff of 10 cm (the Z direction). Where Z is 100 cm, the field of view is 80 cm in X.




Fig. 1 Mechanical arrangement of the profile scanner.

The double faceted mirror which is used both on the path of the projected laser and the reflected light is a unique feature of this device called synchronized scanning [1]. This scanning geometry has the advantage of being remarkably immune to ambient light, since the viewing direction is always aligned to the direction of the scan. The signal to noise ratio is very high, which also allows a large depth of view with a constant laser power.

Physically, the profile scanner must be very compact in order to be useful on the wrist of a small payload industrial robot. The scanner we are using, shown in Fig. 2 on the wrist of the PUMA 560, has the following physical characteristics: width 90 mm, length 140 mm, depth 28 mm (this is 80 mm where the galvanometer protrudes), and the weight is 500 grams including the protective case.




Fig. 2 Profile scanner mounted on the wrist of the PUMA 560 robot.

Other components of the apparatus are shown in Fig. 3. The profile scanner is connected via a parallel link to both the IBM PC/AT compatible, and the Harmony¹ system. The experiments described in this paper were implemented on the PC, but the target machine is the Harmony system which is capable of real time feedback control of the PUMA robot [3]. We have found it extremely useful to have the PC in the apparatus to test methods which are intended to be run on the Harmony multiprocessor system.

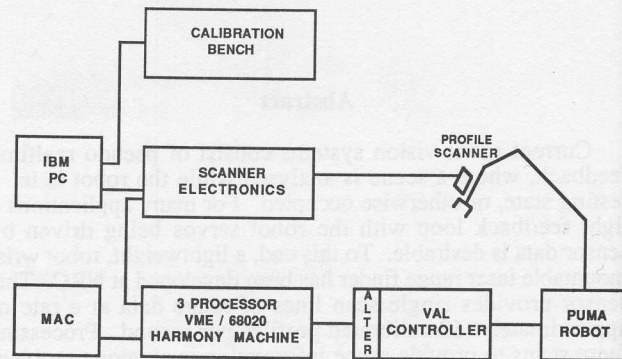


Fig. 3 Schematic of apparatus.

The calibration of the profile scanner is a rather complicated process of converting the raw data from the sensor to Euclidean 2-space. Because of the scanner's geometry, the profile collected has various nonlinear distortions. We have decided to use an experimental method of calibration to transform the profile to an orthographic projection. During this process the PC drives the camera up and down a calibration bench collecting many samples of raw data. Fig. 4 shows the vertical calibration bench which is accurate to .005 mm. The calibration procedure currently in use is described in [5]. It uses B-splines to convert data collected at the bench into lookup tables. These LUTs are stored in RAM and accessed in real time so that the sensor can provide 13 *calibrated* scans of 255 surfels (surface elements) per second. We have found that this B-spline method of calibration LUT generation is not as reliable as a linear interpolation which uses more densely packed input data, i.e., sample data is collected at very small intervals. We have implemented such a linear interpolation calibration but a comparison is not yet available.

1. **Harmony** is a mark reserved for the exclusive use of Her Majesty the Queen in right of Canada by Canadian Patents and Development Ltd./ Societe canadienne des brevets et d'exploitation.




Fig. 4 Vertical calibration bench.

Accuracy and resolution of the range finder is a very difficult issue for several reasons. Because the sensor is scanning by an oscillating mirror, there is some spherical distortion in the data. Although the calibration procedure effectively *flattens* the curvature, the accuracy of each surfel is no longer equal, due to the fact that the surfels near the ends of the scan are farther apart in Euclidean space than the surfels directly below the scanning mechanism. It is also noted that as the Z increases, i.e., the sensor moves farther away, the resolution decreases nonlinearly. Thus the accuracy of range finding devices vary nonlinearly in both X and Z. A standard method of determining and expressing the accuracy of range finders in real world, calibrated terms has not yet been devised. For our purposes, we have recorded the deviation of a profile from a straight line, when the sensor is viewing an ideal flat surface. The accuracy in Z decreases from approximately .2 mm to 10 mm over our calibrated range of 11 cm to 70 cm. We are certain that we can improve on this using the linear interpolated calibration, but we do not know by exactly how much.

3. Tools for Processing Calibrated Data

An interactive system using a library of utilities has been written on an IBM/PC AT compatible which is interfaced to the range finder. These tools allow us to try new operators, and combinations of operators, on the profiles to determine which are useful, what size of operator is required, etc. Fig. 5 is a menu for this interactive system. The profiles that are stored internally are numbered for reference. Any of the operators can specify any desired profile. For example, to grab a new calibrated profile from the sensor and store it as scan#2 the user simple enters: "1 2". The operators on the menu are commonly used in intensity image processing, and are implemented here specifically for X-Z range profiles. Fig. 6 is a photograph of a step pyramid which is used to demonstrate

Scan Utilities

1. Get scan	Scan#
2. Plot	Scan#
3. Plot title	Scan# title xlabel ylabel
4. Repair	Scan1# Scan2#
5. Add	Scan1# Scan2# Scan3#
6. Subtract	Scan1# Scan2# Scan3#
7. Greater Than	Scan1# Scan2# Threshold
8. Less Than	Scan1# Scan2# Threshold
9. Translate	Scan1# Scan2# X Z
10. Centroids	Scan#
11. Volume	Scan#
12. First Deriv	Scan1# Scan2#
13. Sobel	Scan1# Scan2# op_size
14. Smooth	Scan1# Scan2# op_size op_val1
15. Erode	Scan1# Scan2# op_size
16. Dilate	Scan1# Scan2# op_size
17. Open	Scan1# Scan2# op_size
18. Close	Scan1# Scan2# op_size
19. EXIT.	

Fig. 5 Menu of profile scan operations.




Fig. 6 Step pyramid being scanned.

the operators. This pyramid has 7 steps of approximately 6 mm. each, and the base is approximately 12mm. thick. The base is 112 mm. square. Note that this is a piece of wood, not a carefully machined target that would be used for examining closely the accuracy of the sensor. Fig. 7 demonstrates the capability of the range finder to collect data over a large field of view. The samples in the plot are collected at various standoffs measured from the range finder's frame to the background plane. As should be expected, the detail is lost as the sensor moves away. Also the length of the scan is considerably longer at the more distant samples. The data has deteriorated to the point where some surfels are lost at a standoff of 70 cm. The plot may appear upside down, but remember that the distances are measured from the sensor to the surface, not from a background plane as range data is usually displayed.

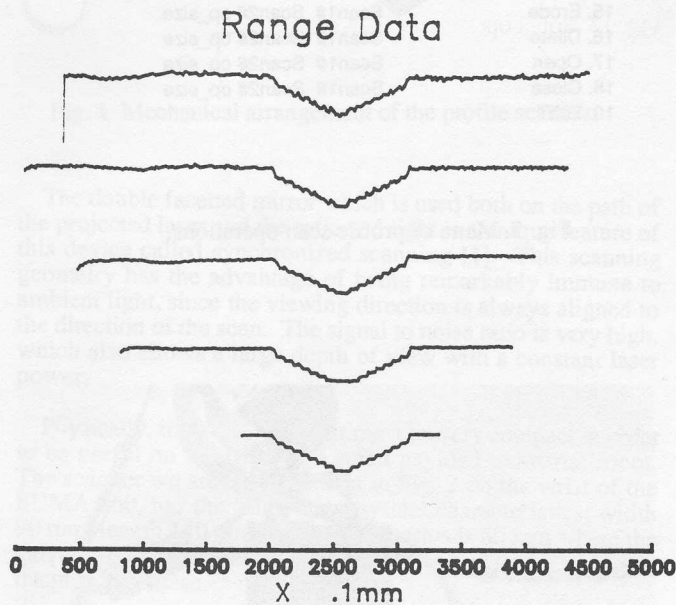


Fig. 7 Profiles of the pyramid from standoff distances of 20 cm., 40 cm., 50 cm., 60 cm., and 70 cm.

We have chosen a standoff of 22 cm. to demonstrate the functionality of some of the operators. We do not specify at this time for what applications these operators are useful. The objective of this exercise has been to provide tools to investigate the applicability of various operations to different applications. As these operations are moved to the Harmony system, the real time computational requirements will become critical. These tools are the first step in a procedure that can determine how many Harmony processors are required for a particular application. Fig. 8 is the calibrated data which will be used for the subsequent processing. The first operation is to repair the shadow effect which is evident in this plot. We have chosen to close this gap by replacing the missing surfels with one of the usable Z values which borders the gap. The one which is farther from the sensor is used. This is the obvious choice. Although we do not know that this value is in fact a correct representation of the scene, we do know that the higher value can not be correct because the gap would not exist at this position, i.e., valid data would continue to be available. The surfels that are replaced in this manner are flagged as being suspect in the structure representing a profile. Fig. 9 is the graphical result of this repair.

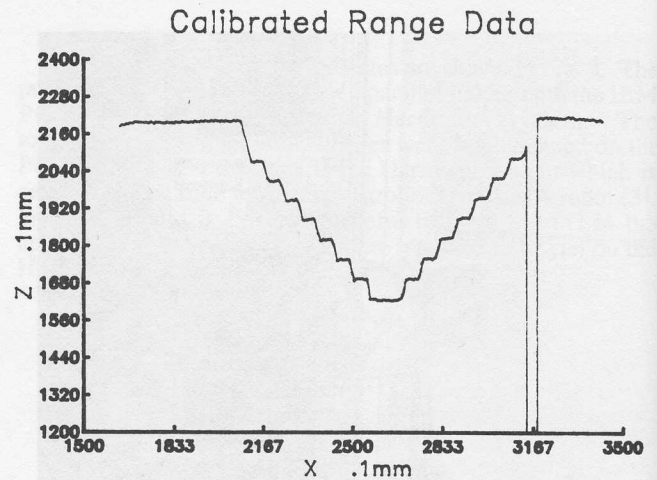


Fig. 8 Calibrated range profile of the step pyramid at a standoff of 22 cm.




Fig. 9 Shadow effect is removed.

The next 3 operators are quite well known. A 'smooth' operator allows the user to enter the size of the neighbourhood to be convolved and the coefficients for the operator. This is demonstrated in Fig. 10 by a simple local averaging. This operator is written so that the user can specify any size of neighbourhood, and enter the floating point coefficients representing the operator. For example, to arrive at Fig. 10, we entered "14 1 2 5 .2 .2 .2 .2 .2" which means "transform scan#1, store result as scan#2, use an operator 5 surfels wide, and a weighting factor of .2 on each surfel".

Fig. 11 labeled 'First Derivative' is the result of subtracting each surfel from its following neighbour. Fig. 12 is the result of the now familiar Sobel operator using a neighbourhood of 5 surfels. Although these operators have been used in greyscale image processing for decades, the result of a derivative operator is measured in grey levels that are thresholded in one way or another, while the result of a derivative operator in range data is measured on the same coordinate system. In our case in Fig. 11 and Fig. 12 where the derivative operators are applied we note that the step sizes on the pyramid are approximately correct.

Morphological operators have been applied to range images with some success [7]. A recent paper formalizing many aspects of mathematical morphology can be found in [6]. The erode and dilate operators are included in these utilities in a simplified one dimensional form. An erosion in this sense is performed by taking the minimum value of the surfels over a specified size neighbourhood and replacing the centre surfel with this value. Similarly, a dilation replaces the centre surfel with the maximum value. An opening is an erosion followed by a dilation, and a closing is a dilation followed by an erosion. Figs. 13 and 14 are the results of performing an erosion and a dilation respectively. Fig. 15 is an opening. As mentioned previously the interactive system is designed to allow the user to combine operators by passing the results of a particular operation to a subsequent operation. Fig.16 is the result of a morphological opening passed to the Sobel operator. The results are only slightly better for the sample case than the Sobel operator in Fig. 12, but this difference may easily be a significant difference in some applications where the profile contains finer details.




Fig. 12 Sobel operator detecting step edges.

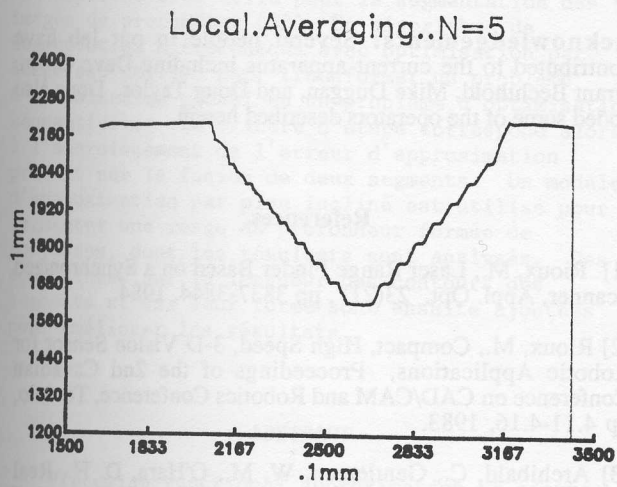


Fig. 10 Local averaging.

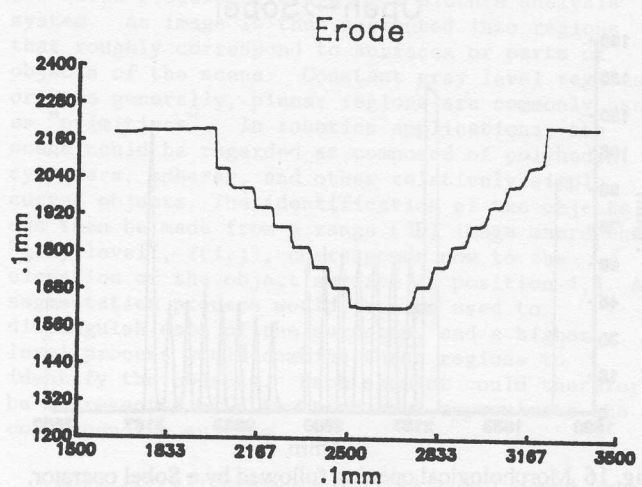


Fig. 13 Morphological erosion.

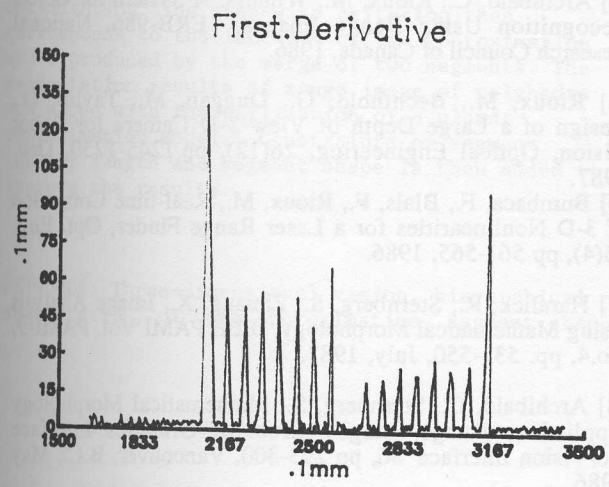


Fig. 11 First derivative.

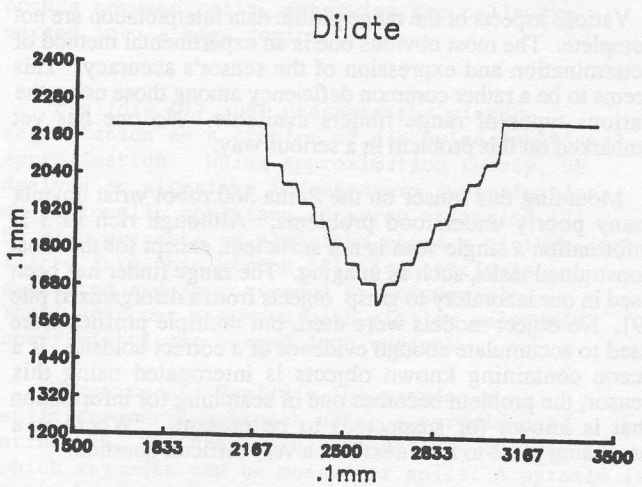


Fig. 14 Morphological dilation.

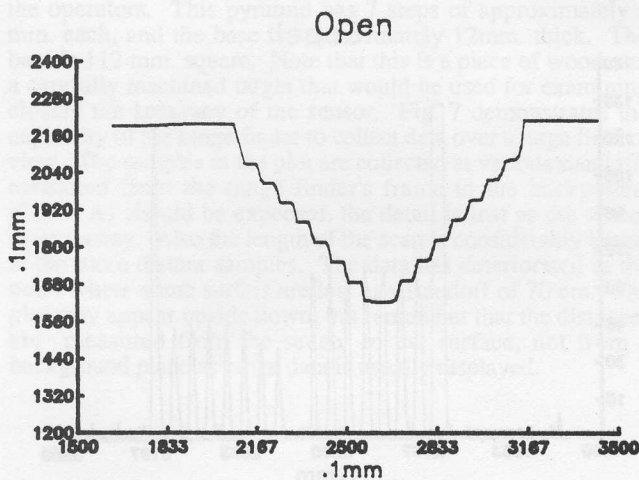


Fig. 15 Morphological opening.

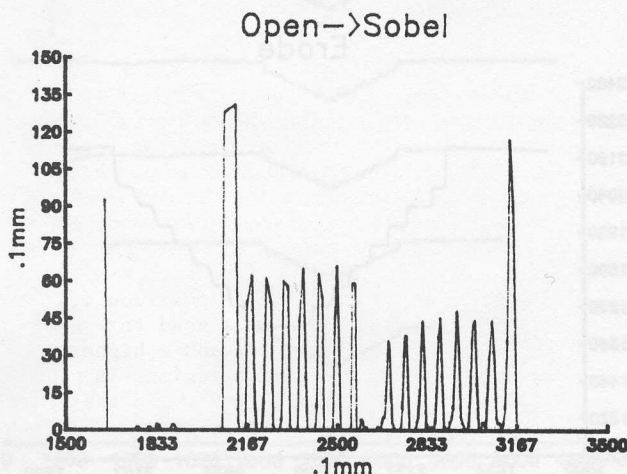


Fig. 16 Morphological opening followed by a Sobel operator.

4. Conclusions and Further Work

Various aspects of the range profile data interpretation are not complete. The most obvious one is an experimental method of determination and expression of the sensor's accuracy. This seems to be a rather common deficiency among those using the various types of range finders available. No one has yet embarked on this problem in a serious way.

Mounting this sensor on the Puma 560 robot wrist unveils many poorly understood problems. Although rich in 3-D information a single scan is not sufficient, except for the most constrained tasks, such as gauging. The range finder has been used in our laboratory to grasp objects from a disorganized pile [9]. No object models were used, but multiple profiles were used to accumulate enough evidence of a correct holdsite. If a scene containing known objects is interrogated using this sensor, the problem becomes one of searching for information that is known (or suspected) to be present. "Where is a promising place to look next?" is a very difficult question.

Because there are only 13 scans of 255 pixels arriving per second, it is possible to process all the data without special purpose image processing hardware. Our target system is the multi-tasking multi-processing Harmony operating system, running on a series of M68020 processors. We have

demonstrated this system's capability for real time tracking of moving objects [3]. Real time 3-D visual data for feedback control will make many interesting robotics applications feasible.

Several industrial applications of the profile scanner as it exists now have been suggested. 3-D inspection of manufactured parts, including electronics inspection, is a promising area. There are also various volume estimation problems for which this sensor is suitable. Specific applications in the space station program have been suggested, and we are now experimenting with locating a satellite grapple fixture with the profile scanner. A sparse range image is used.

The utilities for processing the range profiles have not been presented as a method of solving any particular application but will serve as tools for investigating any application. Some processing methods are obviously better suited for one application than another for reasons such as accuracy traded off with speed, etc. By arranging these utilities as an interactive package, it is very easy to experiment until the desired results can be obtained. We have demonstrated with figures only a small sampling of what can be achieved by combining these operators.

Acknowledgements: Several people in our lab have contributed to the current apparatus including Dave Green, Grant Bechthold, Mike Duggan, and Doug Taylor. Doug also coded some of the operators described herein.

References

- [1] Rioux, M., Laser Range Finder Based on a Synchronized Scanner, *Appl. Opt.* 23(21), pp 3837-3844, 1984.
- [2] Rioux, M., Compact, High Speed, 3-D Vision Sensor for Robotic Applications, *Proceedings of the 2nd Canadian Conference on CAD/CAM and Robotics Conference*, Toronto, pp 4.11-4.16, 1983.
- [3] Archibald, C., Gentleman, W. M., O'Hara, D. H., Real Time Feedback Control Using a Laser Range Finder and Harmony, *Proc. of 7th Canadian CAD/CAM & Robotics Conference*, Toronto, June 1988.
- [4] Archibald, C., Rioux, M., Witness: A System for Object Recognition Using Range Images, ERB-986, National Research Council of Canada, 1986.
- [5] Rioux, M., Bechthold, G., Duggan, M., Taylor, D., Design of a Large Depth of View 3-D Camera for Robot Vision, *Optical Engineering*, 26(12), pp.1245-1250, Dec. 1987.
- [6] Bumbaca, F., Blais, F., Rioux, M., Real-time Correction of 3-D Nonlinearities for a Laser Range Finder, *Opt. Eng.* 25(4), pp 561-565, 1986.
- [7] Haralick, R., Sternberg, S., Zhuang, X., Image Analysis Using Mathematical Morphology, *IEEE PAMI Vol. PAMI-9*, No.4, pp. 532-550, July, 1987.
- [8] Archibald, C., Sternberg, S., Mathematical Morphology Applied to Range Images, *Proc. of Graphics Interface '86/Vision Interface '86*, pp 293-300, Vancouver, B.C., May 1986.
- [9] Roth, G., O'Hara, D., A Holdsite Method for Parts Acquisition Using a Laser Range Finder Mounted on a Robot Wrist, *1987 IEEE International Conference on Robotics and Automation*, Raleigh, NC, pp 1517-1523.