

Calibration and control of a robotic arm using a range imaging camera

Cameron B. D. Kelly, Adrian A. Dorrington*, Michael J. Cree and Andrew D. Payne
Department of Engineering, The University of Waikato, Hamilton, New Zealand

ABSTRACT

Time of flight range imaging is an emerging technology that has numerous applications in machine vision. In this paper we cover the use of a commercial time of flight range imaging camera for calibrating a robotic arm. We do this by identifying retro-reflective targets attached to the arm, and centroiding on calibrated spatial data, which allows precise measurement of three dimensional target locations. The robotic arm is an inexpensive model that does not have positional feedback, so a series of movements are performed to calibrate the servos signals to the physical position of the arm. The calibration showed a good linear response between the control signal and servo angles. The calibration procedure also provided a transformation between the camera and arm coordinate systems. Inverse kinematic control was then used to position the arm. The range camera could also be used to identify objects in the scene. With the object location now known in the arm's coordinate system (transformed from the camera's coordinate system) the arm was able to move allowing it to grasp the object.

Keywords: Time of flight, range imaging, calibration, robotics, target recognition

1. INTRODUCTION

Time of flight (TOF) range imaging cameras obtain a three-dimensional representation of a scene by measuring distance for each pixel. This three-dimensional information is useful for machine vision applications to allow computers to determine the shape, size and location of objects in a scene. As robotic limbs do not innately know their position in space, they require a method of feedback for accurate calibration and control. Previous publications demonstrate methods of calibrating a robotic arm using a TOF camera or standard cameras attached to the arm's end effector^{1,2}. In this paper we introduce a method of calibrating and controlling a simple robotic arm using a range imaging camera external from the arm. The arm used in our experiments was a Lynxmotion AL5D (Lynxmotion, Pekin, Illinois, USA). The AL5D arm has no positional feedback and comes with no method of accurately calibrating its servos. The goal was to be able to move the arm's end effector to a position defined in the range imaging camera's coordinate system. To do this the arm's location and orientation needed to be found in the scene, as well as the relationship between each limb's position in space, and the controllers input. A commercially available Mesa Imaging Swissranger 4000 TOF (Mesa Imaging, Zürich, Switzerland) range imaging camera was used to accomplish this goal.

2. ROBOT CONTROL

The AL5D is a small consumer scale robotic arm. It can be simplified to a 2D three degrees of freedom (DoF) kinematic chain (figure 1a). The three joints that make up the chain are referred to as the shoulder (figure 1a, origin), the elbow (figure 1a, P1), and the wrist (figure 1a, P2). The arm is mounted on a rotating base that increases the arm's workspace from 2D to 3D. The servos in the AL5D are controlled by setting a Pulse Width Modulation (PWM) value. PWM is a common form of motor control that works by varying the duty cycle of a square wave. The PWM value, controls the angle servo to which the turn should move. As the goal is to move the arm to a point in space, inverse kinematics were used to control the arm. Inverse kinematic control systems are well documented³ as is the task of computing the different servo angles from a given target point. The downside is that inverse kinematic problems very rarely have unique solutions and can often have no solutions.³ To simplify the inverse kinematic algorithm, it was modified to require four inputs, the x, y, z target coordinates and the angle of the gripper relative to the ground. This reduces the inverse kinematic problem to a two DoF chain (figure 1b). The basic principle in the method we selected to calculate the servo angles was to find the intersection of two circles, one centered at the base and the other centered at the target point (figure 1c). In most cases this produces two results, therefore a decision was made as to which result to use, and after calibration this was based on the limits of the servos' movements.

*adrian@waikato.ac.nz; phone: +64-7-858-5062; fax: +64-7-838-4835




Figure 1. Arm Kinematics

3. TIME OF FLIGHT RANGE IMAGING

Time of Flight range imaging is an increasingly common method of producing range images (images in which pixel values represent distance). TOF is the measurement of the time light takes to travel from an active source to an object and back to an image sensor. The time light takes to travel the short distances typically found in indoor scenes is in the order of pico-seconds. Making direct measurements at this time scale is difficult for modern electronics. One solution is to use an indirect method utilizing an Amplitude Modulated Continuous Wave (AMCW) illumination source. AMCW works by modulating the light source then measuring the phase change in the modulation envelope when the light is reflected back to a sensor. This phase change typically is measured by modulating the sensor at the same frequency to the light source (homodyne detection), producing an amplitude directly related to range. The camera used in this experiment was the Swissranger 4000 (SR4K). The SR4K is a homodyne system and that acquires four image frames ($C_0 - C_3$) each with a $\frac{\pi}{2}$ phase step between the illumination and the sensor modulation. Using these four captures, the amplitude (intensity, figure 2:left), DC offset (ambient lighting), phase (φ), and therefore distance, can be calculated with equations (1)-(4).⁴

$$A = \frac{\sqrt{(C_0 - C_2)^2 + (C_1 - C_3)^2}}{2} \quad (1)$$

$$B = \frac{C_0 + C_1 + C_2 + C_3}{4} \quad (2)$$

$$\varphi = \tan^{-1} \left(\frac{C_0 - C_2}{C_1 - C_3} \right) \quad (3)$$

$$d = \frac{\varphi c}{4\pi f_{mod}} \quad (4)$$

The SR4K can be used to produce standard range images containing radial distances (figure 2:right), however it can also produce three separate images containing the calibrated spatial Cartesian x, y, z coordinates (figure 3). The calibration adjusts the radial distance values into distance calibrated in meters, performs a perspective projection and corrects for lens distortion (note: the line where the table meets the wall in figure 2 right compared to figure 3). The x, y, z coordinates were used to calibrate and position the robotic arm in our experiments. The manufacture stated a typical 1σ precision of the SR4K is $\pm 4 \text{ mm}$ ⁵ The stated conditions for this precision are 100% target reflectivity at 30 FPS, at a 2 m working distance. To improve the precision each capture used was made by averaging 50 standard captures.

4. ARM IDENTIFICATION

The AL5D arm measures 146.5 mm from the shoulder servo to the elbow (L1), 186.5 m from the elbow to the wrist (L2), and 100 mm from the wrist to the middle of the end effector (L3). The base (point) of the arm is defined as the intersection of the vertical axis of rotation (rotated around the base servo) and the shoulder servos axis of rotation. In order to calibrate




Figure 2. Left: Standard black & white intensity, Right: Range image (black close to camera white distant)




Figure 3. Surface plot of the calibrated x, y, z data

the robotic arm the computer needs to be able to identify the arm. Retro-reflective targets were used as identifiable features on the arm to simplify the task of segmenting the desired features. The retro-reflective spots were mounted 44 mm from the center of the arm. A major advantage of using retro-reflectors is that they are effectively immune to multi-path effects. Multi-path effects result in incorrect depth measurements, due to light scattering from one object to another object before returning to the camera. The light measured from the second object therefore, contains mixed phase information resulting in incorrect depth.⁶ Measurements using retro-reflective targets are more accurate because they appear significantly brighter in the scene than anything else meaning any interfering multi-path will be of insignificant brightness. The retro-reflective spots were placed on the axis of rotation of the shoulder, elbow and wrist servos. This was helpful as it meant the centroid of the target, tasked to identify a particular joint, did not move when the joint moved. In a typical setup the camera was fixed in place approximately 900 mm from the base of the arm. The arm was also covered with dark colored paper, to




Figure 4. Intensity images of arm with clearly visible retro-reflective targets; Left: Higher integration time saturating the targets; Right: Lower integration time with non-saturated targets

attenuate any specular reflections and make segmentation easier. As this system used active illumination, with the source based around the image sensor, a significant amount of light was returned from the retro-reflective surfaces. A concern was that the retro-reflectors would saturate the sensor making any data returned invalid. Use of a low integration time, limited the light received by the sensor, therefore reducing the chance of the sensor saturating. To further decrease the chance of saturation, black retro-reflectors were used (3M, Scotchlite™ Plus, 680 Series Reflective Sheeting), as shown in figure 4. Even with the limitations on the retro-reflector brightness, the targets were still by far the brightest objects in the scene (figure 4 right). Basic intensity thresholding was used to rapidly segment the targets, leaving only the three targets shown in figure 4. The targets were analyzed using the range data individually to find their 3D center. To accurately identify the center of the targets a gray scale centroiding algorithm was used. A traditional pixel based centroiding algorithm common in photogrammetry,⁷ is:

$$\bar{x} = \frac{\sum_{j=1}^m \sum_{i=1}^n i \cdot I_{i,j}}{\sum_{j=1}^m \sum_{i=1}^n I_{i,j}} \quad (5)$$

where $I_{i,j}$ is the intensity value of the i, j pixel, and m and n are dimensions of the window in which the centroid is being calculated. The gray scale centroiding gives particularly good results with smaller target sizes,⁷ which is beneficial as the SR4K has a low resolution of only 176×144 pixels. This approach has been applied to radial range images before,⁸ but we now apply it to the calibrated Cartesian coordinate data using equation (6) repeated for x, y, z to produce a 3D centroid.

$$\bar{x} = \frac{\sum_{j=1}^m \sum_{i=1}^n x_{i,j} \cdot I_{i,j}}{\sum_{j=1}^m \sum_{i=1}^n I_{i,j}} \quad (6)$$

The modification replaces the index pixel value i with the calibrated spatial x value produced from the SR4K. The results shown in table 1 are the standard deviation of the x, y, z coordinates taken of four targets centroided compared over 100 consecutive captures (each capture averaged over 50 low integration time captures) of a static scene. Targets 1 and 3 were 28 mm (8 pixels) in diameter which is the same size as the spots used on the arm. Target 4 had a smaller diameter of 20 mm (5 pixels) and target 2 was smaller again with a diameter 14 mm (4 pixels). For all the targets, submillimeter 1σ precision was achieved.

5. CALIBRATING THE ROBOTIC ARM

Calibration is a set of operations that establishes, under a specific set of conditions, a relationship between sets of values.⁹ In this case the relationship was between the servo's PWM value and the arm's position in space, as defined by the TOF

Dimension	Point 1	Point 2	Point 3	Point 4
X Standard Deviation (mm)	0.134	0.073	0.165	0.054
Y Standard Deviation (mm)	0.121	0.074	0.153	0.040
Z Standard Deviation (mm)	0.237	0.241	0.226	0.251

Table 1. Table of the Standard Deviation in x, y, z , from 100 captures of the same static scene

range imaging camera. Each servo had a PWM range from $500 \mu s$ to $2500 \mu s$ that spanned an angle of approximately π radians. The goal was to find the relationship between the servo angle, relative to a zero axis defined by the orientation of the camera with the origin centered at the base of the arm, and the servo PWM value. To achieve this, each servo was swept through a range of specific PWM values from which the change in the position of the centroid dots was measured. Before any measurements could be made, a decision was required about which target belonged to which servo. This was achieved by measuring the separation between the targets. As the fixed lengths of the arm were known, the targets could be easily identified. The target on the elbow was separated from the shoulder target by 146.5 mm and the wrist from the elbow by 186.5 mm. This technique fails in one particular configuration; when the arm is positioned in such a way that the wrist and the base are separated by either 146.5 mm or 186.5 mm. It was possible to overcome this issue, by identifying the targets based on where they should be in space, when the ideal position of the arm was known. The ideal position being the kinematic model of the arm mapped on to the base of the arm in the camera's coordinate system. However correct mapping of the ideal position is dependent upon the calibration. To avoid this problem, and to ensure the correct the identification of the targets, the position of the arm was selected during the early calibration stages.

The first step in calibration is to find the base (figure 1(a), origin), the center around which the arm rotated (intersection of the shoulders rotational axis and the base servo's rotational axis). This was done by rotating the base into three different positions. The distance between the target on the shoulder and the base is known to be 44 mm, therefore as the base rotates the target will move in a circle of radius 44 mm. If three circles of radius 44 mm are placed on the target's circular path, in plane with the movement, they will intersect over the axis of rotation. The centroid of the target aligns with the center of rotation for the shoulder servo therefore the point of intersection is the base. In order to simplify the task of calibrating each servo, the servos were individually calibrated from the base up. The different calibration movements are shown in figure 5.




Figure 5. Arm movements for calibrating the three servos. Each servo is moved sequentially, allowing the camera to make individualized measurements

6. PERFORMANCE ANALYSIS

The calibration was repeated multiple times and the results compared to validate the accuracy of the measurements. The plots in figure 6 (a)-(c) show the relationship between each PWM value and the corresponding specific limb's angle (figure

6 (d) is referred to later). The angle of the base is measured against the camera's x-axis. The angles of the shoulder servo and the elbow servo are measured relative to the ground (parallel to the x,z plane). The linear fit in all three cases had an r^2 value of 0.99. The dashed line represents the linear fit to the data.

The linear relationship was used to calculate the required PWM value to achieve a desired angle. Tests were run to see if the calibrated arm moved to the targeted position. Figure 7 shows the measured position/orientation of the targets on the arm (dashed line), transformed from their position on the side of arm to be centered in the arm (square marked line), compared to the ideal kinematic model of the arm (x marked line).

6.1. Identifying target object

Depth information can be extremely useful when using a computer to identify objects in a scene. Background subtraction is a technique used to segment changes with standard imaging. One problem that occurs with standard imaging is that changes in the background lighting can cause sufficient differences across the entire scene, thus the entire scene can be included in the segment. However as the depth in the scene is not sensitive to such changes, direct background subtraction can be performed using range data to detect new objects entering the scene. We tested moving the calibrated arm to grab a simple object in a scene identified in this manner. The detected object was located by applying the centroiding algorithm to the x,y,z data in the segmented section. We were successfully able to identify a new target object in the scene, a round piece of wood small enough to fit in side the gripper. The arm was then able move the arm to the target's location, however our control lacked dynamics, so the arm would push the target aside rather than moving and grasping it.

6.2. Limitations

In figure 6 (b), at approximately 1900 μs in the shoulder calibration there is a 0.07 radians deviation from the linear relationship. This is believed to be a measurement error as this is the first point measured during the calibration movement. Another test of the same servo over a PWM range from 1600 - 2200 μs shows the same linear relationship with a similar error (0.05 radians) in the first measurement at 2150 μs yet no error between 1800 - 2000 μs figure 6 (d). The cause of this anomaly has not yet been established.

The retro-reflectors are very useful for easy image segmentation however they only help if they do not saturate. The SR4K would replace the range data from a saturated pixel with a Not-a-Number (NaN). When the centroiding algorithm came across a NaN pixel, it ignored it. If even one pixel saturated, on any given target, it could skew the centroid's position by as much as a centimeter. It was found that in some positions and orientations, the arm would always saturate a specific few pixels. The particular orientations and pixels were dependent on the position of the camera relative to the arm. This is believed to be the result of changing the position of the camera, and in doing so changing the illumination of the arm.


As the retro-reflective targets were additions to the arm, it meant that our method only works when the camera can view all the targets. Another approach that would better address this, is to identify structural features of the arm. The use of structural features should allow the arm to be calibrated from any orientation.

7. CONCLUSIONS AND FUTURE WORK


We have developed a method of externally calibrating a robotic arm using a TOF range imaging camera. With the exact relationship between the camera and arm unknown, to achieve calibration, the computer needed to identify the arm's location and orientation. This was achieved by locating retro-reflective targets positioned on the arm. The center of each target was found by using a centroiding algorithm on the x,y,z data obtained from the range imaging camera. With the position of the targets on the arm identified, the location and orientation of the arm's movement could be inferred. From this data, the arm was then calibrated. The calibration involved finding the relationship between each servo's control signal and corresponding limb angle. This was achieved by moving each servo independently to a known PWM value then measuring the position or change in position. The relationship between the servos and their corresponding angles was found to be linear. In future work we hope to identify the arm based on geometric features, improving the scope of the method.

ACKNOWLEDGMENTS


The authors would like to thank the Chronoptics group for their support through out the project. Funding for this work was provided through the University of Waikato Strategic Investment Fund.




(a)



(b)



(c)



(d)

Figure 6. Calibration graphs: Measured angle vs control PWM, with linear fit. (a)-(c) show the results of multiple runs with similar settings. (d) shows two different sets of shoulder calibration data, highlighting continuing linear relationship as the PWM value increases




Figure 7. Ideal kinematic chain of the arm, in the camera's coordinate system, compared to the measured target positions transformed to have the same base as the ideal model

REFERENCES

1. Kim, M.-Y., Cho, H., and Kim, J. H., "Hand/eye calibration of a robot arm with a 3D visual sensor;" *Optomechatronic Systems II* **4564**(1), 123–134, SPIE (2001).
2. Gumundsson, S., *Robot Vision Application using the CSEM SwissRanger Camera*, Master's thesis, Technical University of Denmark (2006).
3. Bekey, G. A., [Autonomous robots : from biological inspiration to implementation and control], The MIT Press, Cambridge, Massachusetts (2005).
4. Payne, A. D., Dorrington, A. A., Cree, M. J., and Carnegie, D. A., "Characterization of modulated time-of-flight range image sensors," *Three-Dimensional Imaging Metrology* **7239**(1), 723904, SPIE (2009).
5. Mesa Imaging, Zürich, Switzerland, *SR400 Data Sheet* (2009).
6. Godbaz, J. P., Cree, M. J., and Dorrington, A. A., "Undue influence: Mitigating range-intensity coupling in AMCW 'flash' lidar using scene texture," *Proc. Image and Vision Computing New Zealand 2009* (2009).
7. Shortis, M. R., Clarke, T. A., and Short, T., "Comparison of some techniques for the subpixel location of discrete target images," *Videometrics III* **2350**(1), 239–250, SPIE (1994).
8. Dorrington, A., Cree, M., Carnegie, D., Payne, A., and Conroy, R., "Heterodyne range imaging as an alternative to photogrammetry," *Videometrics IX* **6491**(1), 64910D, SPIE (2007).
9. "International vocabulary of metrology basic and general concepts and associated terms, VIM." ISO/IEC Guide 99-12:2007 (2007).