

[REDACTED]

DATE: [REDACTED]
TIME: [REDACTED]
BY: [REDACTED]
FOR: [REDACTED]

[REDACTED]
[REDACTED]
[REDACTED]
[REDACTED]

[REDACTED]
[REDACTED]

[REDACTED]

**RANGE FROM
TRIANGULATION USING AN
INVERSE Perspective
METHOD TO DETERMINE
RELATIVE CAMERA POSE**

Karen Chaconas

**U.S. DEPARTMENT OF COMMERCE
National Institute of Standards
and Technology
Robot Systems Division
Intelligent Controls Group
Bldg. 220 Rm. B124
Gaithersburg, MD 20899**

August 1990



**U.S. DEPARTMENT OF COMMERCE
Robert A. Mosbacher, Secretary
NATIONAL INSTITUTE OF STANDARDS
AND TECHNOLOGY
John W. Lyons, Director**

**RANGE FROM TRIANGULATION USING
AN INVERSE PERSPECTIVE METHOD TO DETERMINE
RELATIVE CAMERA POSE**
Intelligent Controls Group
Robot Systems Division

Karen Chaconas

Date: August 14, 1990

Document number: ICG-#28

Document version: 1.0

Scope of the Document

This document describes a technique to obtain three-dimensional range from two arbitrarily-placed, stationary cameras. The method uses an inverse perspective algorithm to determine the position of each of the cameras with respect to a set of four coplanar points. Using the two transformations obtained, the relationship between the two cameras is determined. Subsequently, the range to a corresponding feature point that appears in both of the camera images can be calculated using triangulation.

1. Introduction

The determination of position and orientation of an object from extracted features can be accomplished using a variety of techniques applied to camera images [2]. Of these methods, two in particular seem computationally efficient and accurate for use with static cameras. The first, stereo vision, uses triangulation between two camera images which view the same object features to compute the object's pose. In order to determine pose from stereo, the focal length of each camera lens and the relative position and orientation between cameras is needed. Only with this information can triangulation be performed. The second method uses an inverse perspective transformation. By using the inverse perspective method, the relative position and orientation between cameras can be determined in order to perform subsequent triangulation calculations. Inverse perspective methods use the known geometry of a pattern to compute position and orientation of the pattern. There are several algorithms that implement inverse perspective techniques [1][2][3][4][6]. Of these, the Hung-Yeh-Harwood algorithm has an advantage in that it uses a closed-form solution to solve for pose from a unique projection of four coplanar points in an image.

The system used to demonstrate this idea is composed of two cameras each mounted on a tripod. There is no restriction on the initial placement of the two cameras with respect to each other. They need not be at the same height nor do they need to have parallel optical axes. In addition, they do not need to be in any known location or orientation. The cameras' poses are arbitrary with the only restrictions being that they must remain stationary after each camera's transformation to a common surface is determined and that the surface appear in both camera images. In addition, any feature point to which range is to be determined must also appear in both camera images.

The work presented in this paper is a useful step to using a two-camera system to determine range. It determines range from triangulation and requires no a priori knowledge of extrinsic camera parameters. The sensitivity of the extrinsic parameters are calculated for a two-camera system, and it is shown how this initial parameter estimation affects the subsequent range calculation.

2. Initializing the Transformation

The transformation from a planar surface to each camera's image plane is determined using an inverse perspective method for four coplanar points which define this surface [1]. The extrinsic camera parameters that quantize the image plane's position and orientation with respect to the planar surface are computed by comparing the known distance between the four points with their relationship as they appear in the camera image. The four points on the planar surface can be thought of as having positional vectors p_0 , p_1 , p_2 , and p_3 defined as the distance between the lens




Figure 1. Four Coplanar Feature Points in the Camera Coordinate Frame

center and each point on the surface (see figure 1). The rectangle formed by these four points can be described by the equation:

$$p_0(1 - \alpha - \beta) + \alpha p_1 + \beta p_2 = p_3. \quad (1)$$

The values for α and β are computed by knowing the width and height of the rectangle defined by the four points. The points have projected positions on the image plane v_0 , v_1 , v_2 , and v_3 defined as the distance between the lens center and each point on the image plane. The relationship between the positional vectors and the image coordinates is:

$$p_i = k_i v_i \quad (2)$$

where the k_i s are the unknowns that are used to scale the points in the image plane frame of reference to the three-dimensional coordinates of the points in the planar surface frame of

reference. Substituting this relationship into equation (1) yields:

$$\frac{k_0}{k_3} (1 - \alpha - \beta) v_0 + \frac{k_1}{k_3} \alpha v_1 + \frac{k_2}{k_3} \beta v_2 = v_3. \quad (3)$$

The v_i s are determined from the image coordinates, and the k_i s can be computed from equation (3) by using the relationship:

$$k_3 = \frac{\| \overline{P_0 P_3} \|}{\left\| \frac{k_0}{k_3} v_0 - v_3 \right\|}. \quad (4)$$

Next, the three-dimensional coordinates of the points in the planar surface's reference frame can be found from solving equation (2). Knowing the three-dimensional coordinates in both the image plane's reference frame and the planar surface coordinate system, the transformation between the two frames can be computed using the relationship:

$$\begin{bmatrix} x \\ y \\ z \end{bmatrix}_c = \begin{bmatrix} r_{11} & r_{12} & r_{13} \\ r_{21} & r_{22} & r_{23} \\ r_{31} & r_{32} & r_{33} \end{bmatrix} \begin{bmatrix} x \\ y \\ z \end{bmatrix}_p + \begin{bmatrix} t_x \\ t_y \\ t_z \end{bmatrix}_c \quad (5)$$

where the subscripts c and p denote quantities observed in the camera and planar surface frames respectively, and the matrix elements r represent the rotation between the two frames and the matrix elements t represent the translational component. To estimate the extrinsic camera parameters of translation and rotation with respect to the planar surface, the origin of the reference frame on the planar surface is chosen to be at p_0 , and p_1 and p_2 are at $(0, a, 0)$ and $(0, 0, b)$ respectively. The value a represents the height of the rectangle formed by the four points and the value b represents the width. Then, the second and third columns of the rotational matrix are:

$$\begin{bmatrix} r_{12} \\ r_{22} \\ r_{32} \end{bmatrix} = \frac{p_1 - p_0}{\|p_1 - p_0\|} \quad (6)$$

$$\begin{bmatrix} r_{13} \\ r_{23} \\ r_{33} \end{bmatrix} = \frac{p_2 - p_0}{\|p_2 - p_0\|} \quad (7)$$

Knowing the two columns of the rotational matrix, the first column can be found by taking the cross product of the second and third columns. The translation between the two frames can be found by solving equation (5).

Once the transformation between the surface and each camera's image plane is known, the relationship between the two cameras can be deduced. Knowing the transformation matrix between

each camera's image sensor and the surface, the relative transformation between the two cameras' image sensors can be obtained using the relationship:

$$S_{\lambda_{cam_2}} = S_{\lambda_{cam_1}} \cdot {}^{cam_1}\lambda_{cam_2} \quad (8)$$

$${}^{cam_1}\lambda_{cam_2} = (S_{\lambda_{cam_1}})^{-1} \cdot S_{\lambda_{cam_2}} \quad (9)$$

where $S_{\lambda_{cam_1}}$ represents the transformation between the surface and camera 1 ,

$S_{\lambda_{cam_2}}$ represents the transformation between the surface and camera 2 ,

and ${}^{cam_1}\lambda_{cam_2}$ represents the transformation between camera 1 and camera 2.

This relationship is shown below in figure 2:




Figure 2. Obtaining the Transformation Between Two Cameras

As a result, any point in one camera image can be transformed to a reference frame defined with respect to the other camera. This fact is used to convert the origins, or focal points, of both cameras as well as feature points on the image planes into the same coordinate system. Therefore, one camera's focal point coincides with the origin of the coordinate system chosen, and all other points are converted to this frame of reference. The transformation in equation (9) is then used to convert all points from camera 2's frame of reference into the same coordinate system with respect to camera 1.

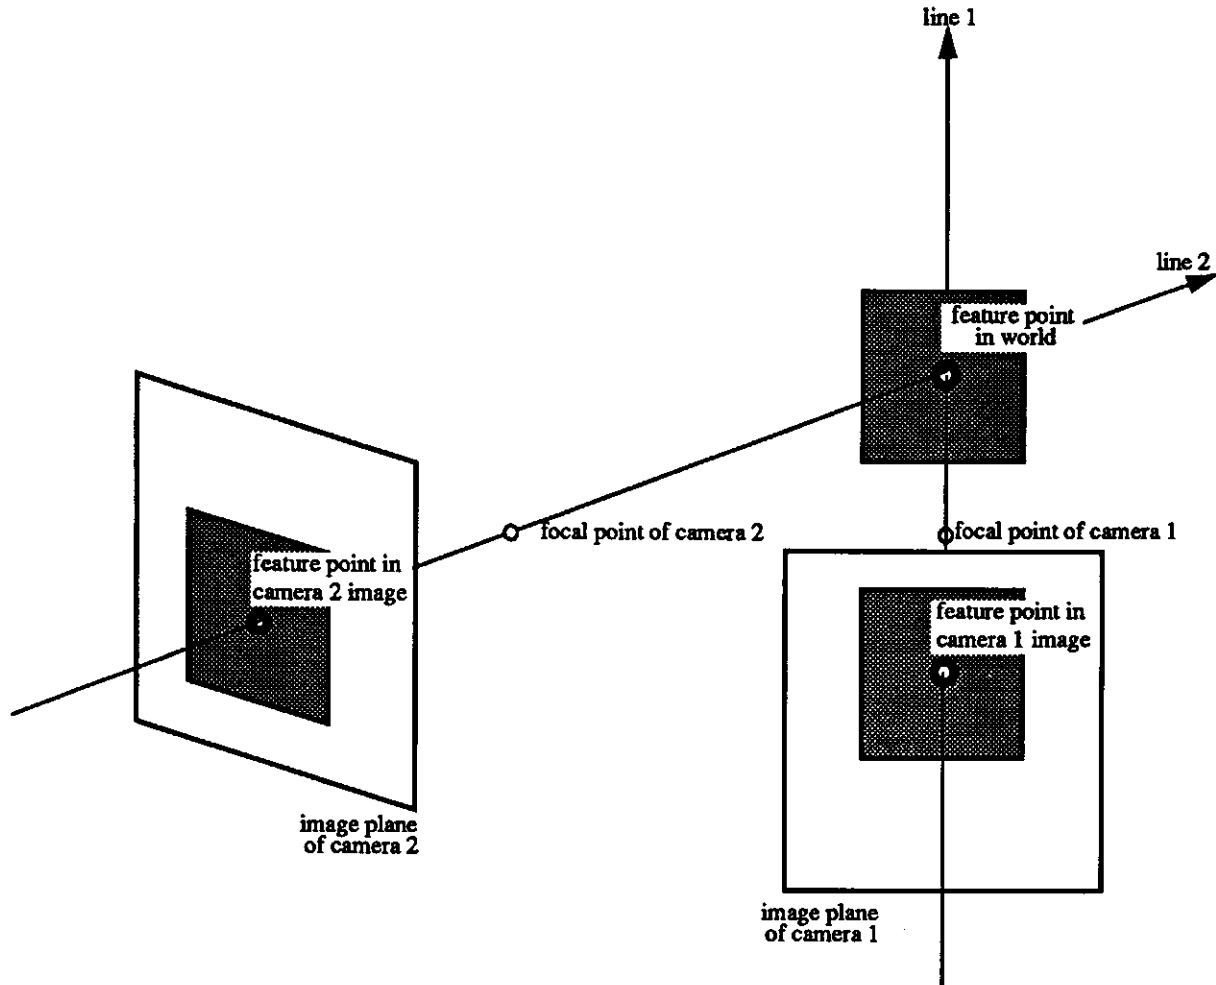


Figure 3. Intersecting Lines to Obtain Range

3. Computing Range

The feature points used in this set of experiments are computed using the centroid of a single object in the camera's field of view. This application computes the centroid of a thresholded intensity image. Alternatively, the centroid of a thresholded motion image could be used as could other extracted feature points. The only constraint on this process is that any feature point in the world to which range is to be calculated must appear in both camera images. A line which extends out from the camera and contains the feature point on the image chip and the camera's focal point is computed for each camera. The three-dimensional position where these two lines intersect corresponds to the world position of the imaged feature point, as shown in figure 3.

Realistically, these two lines will not actually intersect but rather cross over each other in a skewed relationship as in figure 4. To compute the range using two non-intersecting lines, a more complicated solution is necessary. For each camera, a plane is found which contains the line that passes through the focal point and feature point and which is perpendicular to the X,Z plane. The equation of the plane can be represented by the general point-normal plane equation:

$$a(x - x_0) + b(y - y_0) + c(z - z_0) + d = 0 \quad (10)$$




Figure 4. Finding Range Using Skewed Lines

where the coordinates for the camera's focal point are represented by (x_0, y_0, z_0) , the coordinates for the feature point on the image sensor are (x, y, z) , and a, b, c , and d are coefficients describing the plane and which are found in the following manner.

The coefficients of each of the planes are determined by first finding two lines that are contained in the plane. One of the lines is the line that extends through the focal point of the camera and the feature point on the image plane. The other line is the projection of this line onto the X,Z plane. The vector that is normal to these two lines is found by taking their cross product. Since the resulting vector is normal to the plane containing the two lines, the vector's parameters define the plane coefficients a, b , and c . The remaining coefficient, d , is found by using two points on the plane, the camera focal point and the feature point on the image plane, and the other known coefficients and substituting them into equation (10).

The intersection of these two planes is found, and the solution forms a line. Using the equations of the two planes:

$$a_1x + b_1y + c_1z + d_1 = 0 \quad (11)$$

$$a_2x + b_2y + c_2z + d_2 = 0 \quad (12)$$

the line is found by solving the following equations:

$$z = t \quad (13)$$

$$x = \frac{-d_2 - (c_2 t - b_2 y)}{a_2} \quad (14)$$

$$y = \frac{a_1 c_2 t + a_1 d_2 - (a_2 c_1 t - a_2 d_1)}{a_2 b_1 - a_1 b_2} \quad (15)$$

where any value can be substituted for t to find a point on the line of intersection.

The line where the planes intersect is in turn intersected with each of the other lines that pass through the cameras' focal points. Each of these intersections produces a point that can be solved for using the parameterized form of a line, which is detailed in equations (16) and (17).

$$\frac{x - x_{10}}{x_1 - x_{10}} = \frac{y - y_{10}}{y_1 - y_{10}} = \frac{z - z_{10}}{z_1 - z_{10}} = t_1 \quad (16)$$

$$\frac{x - x_{p0}}{x_p - x_{p0}} = \frac{y - y_{p0}}{y_p - y_{p0}} = \frac{z - z_{p0}}{z_p - z_{p0}} = t_p \quad (17)$$

where the points (x_1, y_1, z_1) and (x_{10}, y_{10}, z_{10}) are on line 1, the points (x_p, y_p, z_p) and (x_{p0}, y_{p0}, z_{p0}) are on the line formed by the intersection of the two planes, and the point (x, y, z) is the intersection point of these two lines. By rearranging these equations, the parameter t_p can be solved for as shown below:

$$t_p = \frac{(x_1 - x_{10})(y_{p0} - y_{10}) - (x_{p0} - x_{10})(y_1 - y_{10})}{(x_p - x_{p0})(y_1 - y_{10}) - (x_1 - x_{10})(y_p - y_{p0})} \quad (18)$$

By knowing the parameter, the intersection point for line 1 can be solved for using the parameterized line equations below:

$$x = t_p \times (x_p - x_{p0}) + x_{p0} \quad (19)$$

$$y = t_p \times (y_p - y_{p0}) + y_{p0} \quad (20)$$

$$z = t_p \times (z_p - z_{p0}) + z_{p0} \quad (21)$$

The same procedure is used to solve for the intersection between line 2 and the line formed by the intersection of the two planes. In this case, the points (x_2, y_2, z_2) and (x_{20}, y_{20}, z_{20}) on line 2 are substituted for the points on line 1. The computed range to the feature point is found by taking the mid-point of the two resulting points, as shown in figure 4. This method of determining range will be referred to as the triangulation method in the remainder of this paper.




Figure 5. Experimental Setup

4. Experimental Setup

In order to test the accuracy of the range calculation using the triangulation method, the system was configured so that the actual range could be measured and compared to the computed range. Two CCD cameras were arranged in a configuration shown in figure 5. They were fixed to an optical bench so that the optical axis of camera 1 was approximately perpendicular to a planar surface containing five white dots. The other camera, camera 2, was affixed to the optical bench approximately 45 cm to the left of camera 1. Camera 1 was mounted on a rigid platform beneath a string potentiometer. The string potentiometer was used to measure distances up to 50 inches using voltage readings that were measured in 0.001 mV increments up to 10 V. Camera 2 was mounted on a stand that allowed for pan and tilt adjustment. The two cameras were placed at different heights above the optical bench; camera 1 at 13 cm and camera 2 at 25 cm. This arrangement was used to test both the error in range calculation due to the computation of the extrinsic camera parameters and due to the computation using the triangulation method.

For the first set of experiments, the four outermost white dots on the planar surface were used to compute the transformations to both cameras. The transformations are those referred to by the Greek letter lambda in equations (8) and (9). The planar surface was moved parallel to the optical axis of camera 1 over a range of positions from 300 cm to 1400 cm from camera 1 at intervals of 75 cm. Each time the surface was moved, the pan of camera 2 was adjusted to keep the four dots in the field of view. Two different planar surfaces were used during the experiments. The first one was a small black cube with five white dots where the four outermost dots formed a 92.5 cm square. As the distance between the cameras and the surface grew larger, the separation of the white dots in the image became less distinct. To minimize error in the transformation, it was necessary that

the four outermost white dots be as disparate as possible. Therefore, at ranges greater than 80 cm, a cube that was twice the size of the small one was used. Every time the cube was moved to a new position, range was computed to the fifth dot at the center of the square in three different ways. The first method measured the actual range using the string encoder. This measurement was used as the "true" range to which the ranges from other methods were compared. The second method, referred to as the lambda method, was to compute range using the translation vector from the transformation between camera 1 and the planar surface. This provided range to the center of the four dot pattern, which is where the fifth dot was located. The third method was to compute range using the triangulation method described in sections 2 and 3.

The next set of experiments compared the accuracy of range calculations in a different manner. The transformation was computed to the surface at a fixed range. Then, leaving the cameras in the same position and orientation, range was computed to the center dot at different ranges which varied from 300 cm to 1300 cm. The range was computed using the string potentiometer and the triangulation method. Then, the initial transformation was taken at a different range, at increments of 75 cm, and ranges to the feature point were again varied. This entire process of taking the transformation at one position and measuring the range to the feature point at different distances was repeated for transformations taken from 300 cm to 1300 cm. Analysis of the data taken during these experiments is discussed in the next section.

5. Analyzing Sources of Error

The data from the two sets of experiments described in the previous section is graphed to determine the affect that the transformation computation has on the accuracy of the triangulation method. First, the accuracy of the transformation computation is quantified separately from the accuracy of the triangulation calculation. The first set of graphs quantifies the amount of error produced when using different transformations to estimate the position of a point. This measure is useful to determine the amount of error present in the transformation. Next, the relationship of the transformation algorithm and the triangulation algorithm can be determined. The next set of graphs compares the actual range to the range computed using the translation component of the camera-to-surface transformation and to the range computed from triangulation. This comparison is useful to discover the influence of the transformation on the triangulation computation, since the second algorithm depends on the first to compute range using two cameras. The last set of graphs shows the accuracy of the range calculation when the transformation is computed at different distances. These graphs demonstrate the sensitivity of the triangulation method to the range at which the transformation is computed.

The amount of error present in the computation of the transformation can be determined by using the transformation to estimate the position of one of the four coplanar points and to see how much of a discrepancy exists. It is also useful to determine if there is any correlation between a transform's accuracy and its range from or angle to the surface defined by the coplanar points. In figure 6, the graph shows the error in the computed position of the lower left-hand point on the surface of the small cube. Transformations from the two different cameras over a range of distances were used. The true position of the lower left-hand point is shown as the vertex of the two dark lines. The scale to show the errors is reflected in tenths of a millimeter and the actual size of the




Figure 6. Error in Transformation for Small Cube

small square is 92.5 mm.

The estimated dot positions show that orientation of the camera with respect to the planar surface is the largest contributing factor to error in the transformation computation. The plotted points fall into two basic groups. One set shows the transformations as computed using an image from camera 1, which was oriented perpendicular to the surface. This set shows less than 0.05 mm error in x position and 0.025 mm error in y position. The second set uses images from camera 2, which was oriented at approximately a 45° angle from the surface. This set shows almost consistently an error of -0.1 mm in x position and -0.025 mm in error in y position. The different errors in x position between the two groups are due significantly to error in orientation estimation.




Figure 7. Error in Range Using Three Different Methods

The angle between the camera and the surface contributes more to error than the distance between the two. The position estimates are consistent even though the range from each of the cameras to the surface varied between 300 and 900 mm.

Next, it can be shown how this error in the transformation contributes in error in the calculated range using the triangulation method. The graph in figure 7 plots the difference between the range computed using the position vector of the transformation between camera 1 and the planar surface and the actual range between the two measured by the string potentiometer. This error is plotted in the lower sets of lines. The graph permits comparison of this accuracy with the difference between

the actual range and the range computed using triangulation. This last measure of error is plotted in the upper pair of lines.

The accuracy between the measurements using the small cube and those made using the large cube can be compared using the sets of data that plot the error in the range from the transformation. The lefthand set of data plots range estimates when the transformation was computed over closer ranges using the small cube. The righthand set of data plots range estimates when the transformation was computed over farther distances using the large cube. The increase in accuracy evident between 800-900 mm occurs because each of the four circles on the large cube occupies more area in the image than a small circle. This increases the ability to accurately compute the circle's centroid in the image. In addition, the four dots on the large cube are more widely separated in the camera image. This increases the accuracy of the transformation because the dimensions of the spacing between the dots is less sensitive to error.

From this graph, it can be seen that the error in the position computed by the transformation increases from roughly 3 to 15 mm over the distance from 300 to 1400 mm. The difference in position between the actual range and the range from the triangulation method remains generally larger than the error in the range computed using the transformation. The conclusion is that, since the triangulation method depends on this transformation, the method is limited in its accuracy by the transformation accuracy.




Figure 8. Percentage Error in Range Using Three Different Methods

By looking at the next graph in figure 8, we see that the percentage error of the transformation range over this distance is relatively constant at about 1.0% of the actual distance. The error for the triangulation methods actually decreases from 1.5% to 1.0% over this same distance. Therefore, it is shown that the error in the triangulation method approaches the inaccuracy of the transformation method. This occurrence is explained by the fact that as the planar surface is moved farther away from the cameras, the angle between camera 2 and the surface decreases. Therefore, the increasing inaccuracy in one transformation (with respect to camera 1) is compensated for by the increasing accuracy in the other transformation (with respect to camera 2). The combination of the transformations between each of the cameras is an integral part of the triangulation method.




Figure 9. Error in Range Using Different Transformations

Lastly, it is useful to determine the effect of distance between the camera and the planar surface at which the transformation is computed so that subsequent range is calculated with tolerable accuracy. The graph in figure 9 shows 5 different lines which each reflect errors in range computations. The error measurements computed on a specific line use one transformation, where the transformation was computed to the planar surface at a fixed range. The distance at which the transformation was obtained is indicated by a black dot. Then, leaving the cameras in the same position and orientation, range was computed to a feature point at different ranges. The ranges to each position of the feature point were computed using the triangulation method. These ranges were compared to the actual ranges to obtain the error, and the result was plotted to yield one of the lines in figure 9. Then, the initial transformation was taken at a different range, and ranges to the feature point were again varied to produce each of the other four curves.

From this graph, we see that the error in the computed range usually increases beyond the point




Figure 10. Percentage Error in Range Using Different Transformations

where the transformation was computed and decreases in front of this point. This generalization can be made by comparing the effective baseline between the two cameras to the actual distance to the point. When the baseline is large compared to the actual distance, the triangulation method is more accurate than when the actual range increases and the effective baseline diminishes. The exception shown in the line through transformation number 1 can be explained by the likely influence of lens aberrations. Here, the first point is less accurate because its location appears very far to the right in the camera 2 image where lens distortions more greatly influence the accuracy of the image.

Figure 10 plots the percentage error instead of the absolute error. In this graph, the range curves become more level as the distance where the transformation is computed increases. This phenomenon is possible since the angle between camera 2 and the planar surface decreases as the range between the two increases, producing a more accurate transformation. Even though the triangulation calculation becomes less accurate with increasing distance, the increasing rate of its inaccuracy is not as rapid. Though the accuracy improves as the transformation is computed at

larger distances, it can be expected that at some point it would diminish. The accurate ability to locate the four dots defining the planar surface decreases with increasing distance. This occurs since the location of the centroid of a dot is more susceptible to pixel error when the area of the dot is small. Also, the projected sides of the rectangle formed by the four dots decrease in length and therefore become more affected by pixel error. Both of these inaccuracies contribute to incorrect transformation computation as the distance between the planar surface and the camera increases.

6. Conclusions

The use of two cameras to calculate range to a known feature point can be done using two arbitrarily-placed, stationary cameras. There are very few restrictions on the initial placement of the two cameras with respect to each other. They need not be at the same height nor do they need to have parallel optical axes. In addition, they do not need to be in any known location or orientation. The cameras' poses are arbitrary with the only restrictions being that they remain stationary after each camera's transformation to a common surface is determined and that the surface appear in both camera images. In addition, any feature point to which range is to be determined must also appear in both camera images and is assumed to be the same feature in both images.

When using this method to compute range, the accuracy of the range calculated depends on the accuracy of the transformation computed for each of the two cameras and the effective baseline between the cameras. The angle between each camera and the planar surface defined by the four dots has a significant impact on the accuracy of the transformation. The effect of this angle can be seen by describing the relationship between a fixed image plane and a simple rotation about the y axis of the planar surface. As the angle between the image plane and the planar surface increases, the projected width of the rectangle defined by the four points, or equivalently $p_1 - p_0$, changes as a function of $\cos\theta$. Therefore, this projected width decreases as θ approaches 90° . As this width measurement decreases, any error caused by pixel inaccuracy has a greater impact on the calculation of the second column in the rotational matrix in equation (6) and, as a result, in the computation of the first column. Analogously, it can be seen that rotation about the x axis produces error in computation in equation (7) which therefore impacts the transformation computation.

Regardless of transformation error, it is possible to obtain range with less than 2% error by choosing an appropriate distance at which to compute the transformation to a shared planar surface. The greater the distance between the camera and the planar surface, the less inaccurate the transformation computation due to decreasing effects of the angle of rotation between the two surfaces. By using a planar surface that is placed 700 mm from the camera, the error in range calculation using the triangulation method remains under 2% for ranges between 0.4 and 1.0 m. In general, an optimal location for the planar surface is at a distance that is both in the middle of the desired work volume and where the area of the rectangle formed by the four points occupies one-third of the image. This amount of error makes range using two cameras a reliable approach, especially at ranges between 0.5 and 1.5 m, where the effective baseline is greater.

7. Acknowledgments

I would like to acknowledge the efforts of Laura Kelmar and Marilyn Nashman. Discussions with them and their contributions in vision and world modelling have been essential to this work. I would also like to thank Ron Lumia and Jim Albus for constructively commenting on this paper.

8. References

- [1] Fischler, M., R. Bolles, "Random Sample Consensus: A Paradigm for Model Fitting with Applications to Image Analysis and Automated Cartography", *Communications of the ACM*, Vol. 24, No. 6, 1981, pp. 381 - 395.
- [2] Haralick, R., "Using Perspective Transformation in Scene Analysis", *Computer Graphics and Image Processing*, Vol. 13, 1980, pp. 191 - 221.
- [3] Hung, Y., P. Yeh, D. Harwood, "Passive Ranging to Known Planar Point Sets", 1985 IEEE International Conference on Robotics and Automation, March 25 - 28, 1985, St. Louis, Missouri, pp. 80 - 85.
- [4] Lowe, D., "Three-Dimensional Object Recognition from Two-Dimensional Images", *Artificial Intelligence*, Vol. 31, March 1987.
- [5] Nitzan, D., "Three-Dimensional Vision Structure for Robot Applications", *IEEE Pattern Analysis and Machine Intelligence*, Vol. 10, No. 3, May 1988, pp. 291 - 309.
- [6] Tsai, R., "A Versatile Camera Calibration Technique for High Accuracy 3D Machine Vision Metrology Using Off-the-Shelf TV Cameras and Lenses", *IEEE Journal of Robotics and Automation*, Vol. RA-3, No. 4, August 1987, pp. 323 - 344.

BIBLIOGRAPHIC DATA SHEET

1. PUBLICATION OR REPORT NUMBER NISTIR 4385
2. PERFORMING ORGANIZATION REPORT NUMBER
3. PUBLICATION DATE AUGUST 1990

4. TITLE AND SUBTITLE Range from Triangulation Using an Inverse Perspective Method to Determine Relative Camera Pose

5. AUTHOR(S) Karen Chaconas

6. PERFORMING ORGANIZATION (IF JOINT OR OTHER THAN NIST, SEE INSTRUCTIONS) U.S. DEPARTMENT OF COMMERCE NATIONAL INSTITUTE OF STANDARDS AND TECHNOLOGY GAITHERSBURG, MD 20899	7. CONTRACT/GRANT NUMBER
	8. TYPE OF REPORT AND PERIOD COVERED

9. SPONSORING ORGANIZATION NAME AND COMPLETE ADDRESS (STREET, CITY, STATE, ZIP)

10. SUPPLEMENTARY NOTES

☐ DOCUMENT DESCRIBES A COMPUTER PROGRAM; SF-185, FIPS SOFTWARE SUMMARY, IS ATTACHED.

11. ABSTRACT (A 200-WORD OR LESS FACTUAL SUMMARY OF MOST SIGNIFICANT INFORMATION. IF DOCUMENT INCLUDES A SIGNIFICANT BIBLIOGRAPHY OR LITERATURE SURVEY, MENTION IT HERE.) This document describes a technique to obtain three-dimensional range from two arbitrarily-placed, stationary cameras. The method uses an inverse perspective algorithm to determine the position of each of the cameras with respect to a set of four coplanar points. Using the two transformations obtained, the relationship between the two cameras is determined. Subsequently, the range to a corresponding feature point that appears in both of the camera images can be calculated using triangulation.
--

12. KEY WORDS (6 TO 12 ENTRIES; ALPHABETICAL ORDER; CAPITALIZE ONLY PROPER NAMES; AND SEPARATE KEY WORDS BY SEMICOLONS) binocular; camera pose; inverse perspective; range; stereo; triangulation
--

13. AVAILABILITY <input checked="" type="checkbox"/> UNLIMITED FOR OFFICIAL DISTRIBUTION. DO NOT RELEASE TO NATIONAL TECHNICAL INFORMATION SERVICE (NTIS). <input type="checkbox"/> ORDER FROM SUPERINTENDENT OF DOCUMENTS, U.S. GOVERNMENT PRINTING OFFICE, WASHINGTON, DC 20402. <input checked="" type="checkbox"/> ORDER FROM NATIONAL TECHNICAL INFORMATION SERVICE (NTIS), SPRINGFIELD, VA 22161.	14. NUMBER OF PRINTED PAGES 21
	15. PRICE A02

