3D Pose Estimation of Beef Carcasses using Symmetry

David W. Penman and Nawar S. Alwesh
Imaging and Sensing Team, Industrial Research Limited,
P.O. Box 2225, Auckland, New Zealand
d.penman@irl.cri.nz, n.alwesh@irl.cri.nz

Abstract

Robotic tasks in meatworks are faced with the inevitable high variation between carcasses. If the orientation of the carcass in 3D can be determined, subsequent handling and feature identification tasks can be simplified. This paper presents a method for estimating the pose of an object where bilateral symmetry is the only shape property that is assumed. The method operates with incomplete 3D range data suffering from self occlusion, as is obtained from a single viewpoint.

Keywords: orientation, pose estimation, symmetry, range data, beef carcass

1 Introduction

The technique presented in this paper was motivated by our work with animal carcasses. However there is nothing in the technique that is specific to these particular shapes and it is thus more generally applicable. To assist in identifying features on the carcass, we wished to estimate the pose of part of the carcass. The example used here is the head of the carcass. Each head has a different shape, it may only have some of the dominant features such as ears and horns, and its pose was only constrained to within some tens of degrees of a nominal orientation. To further complicate the task, only an incomplete 3D surface representation of the carcass was available.

Heads removed from carcasses are typically suspended from a chain conveyor. A laser range sensor with a vertical scan direction was mounted in a fixed position to one side of the path of the suspended heads. The motion of the heads past the scanner enabled 3D data to be obtained of the surface facing the scanner. The 3D sensor data is naturally represented in a cylindrical coordinate system with the principal axis parallel to the direction of motion of the heads and passing through the sensor centre. In this configuration, points on the surface are visible to the sensor at most once, and thus any point that is occluded will not be viewed again. Although not strictly a single viewpoint sensor, the limitations imposed by this arrangement are equivalent to a single viewpoint.

We have developed a technique to exploit the natural symmetry of animal carcasses, for the estimation of pose.

Symmetry is present in many natural and manufactured objects [1], and is an important cue to the human visual system both for recognising objects and determining their relative position and orientation. In contrast to many manufactured objects, the symmetry of natural objects is often far from perfect. Despite this, we readily detect the presence of an underlying symmetry, and can subjectively assess departures from perfect symmetry, even if we are unable to quantify such imperfections.

In the field of computer vision, symmetry has been used in a variety of applications. The human face is a classic example of a symmetrical object and symmetry has been exploited in a number of ways for face detection and orientation estimation, mainly using 2D images [2, 3]. Continuous symmetry measures have been used in the study of molecular structure [4] and more recently the techniques were applied to improving 3D reconstruction from sequences of 2D images [5] where symmetrical features of a chair are utilised. The plane of symmetry has been used in range images to estimate the pose of human faces [6]. If an object is bilaterally symmetric, the principal axes can be used to find the plane of symmetry, although this is not reliable if only partial data is present [7, 8].

Techniques that exploit symmetry for pose estimation typically require either a complete representation of the object, or the identification of particular features that can be used in forming correspondences. A further technique that is often used in estimating the pose is the registration of a set of data with a model, and various approaches have been used to solve this problem [9, 10], but such techniques obviously begin with prior knowledge of a model.

Our technique makes no particular assumption about the shape of the object other than that a plane of symmetry exists. It does not require any other prior knowledge of the shape such as the identification of features that exist in symmetrical pairs. We have also avoided any use of local curvature measures as these are sensitive to noise or, if filtering is employed, assumptions must be made about the curvature characteristics.

2 Basis of pose estimation technique

The method estimates the pose of an object by locating the plane of symmetry. The object is assumed to be symmetrical and this is exploited both for efficiency and to make the method more robust.

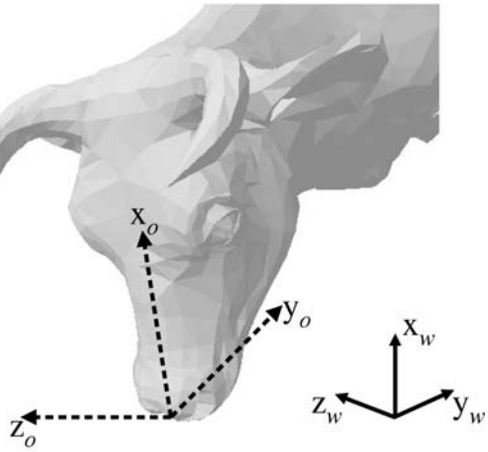

Figure 1: World and object coordinate systems.

Consider an object as in Figure 1, where there is a dominant plane of symmetry, and profiles parallel to the y_oz_o plane in the object coordinate system, are symmetrical about a line parallel to the y_o axis. In the following discussion, orientations are defined in terms of rotations about axes, such as X-rotation (about an X-axis), and a curve representing the surface boundary of an object in a single cross-section through the object is referred to as a profile.

If the data set is reflected about a plane, the position of that plane producing the best measure of similarity between the original and reflected data sets, gives an estimate of the orientation of the object. Strictly, this is true only under certain circumstances. If an object is only X-rotated, the angle of X-rotation can be determined from a single profile parallel to the y_wz_w plane. If an object is both X-rotated and Y-rotated, the symmetry is not necessarily preserved in profiles parallel to the y_wz_w plane, and thus except in special cases, such as a prism, it is not possible to determine either component from any number of such profiles. If the object is only Z-rotated, it will still appear symmetrical about the x_wy_w plane, although the profiles obtained will be different to those before rotation.

Notwithstanding these limitations, early tests using both synthetic and real data obtained from suspended heads of cow carcasses gave encouraging results. In the presence of 10 or 20 degrees of Y-rotation, the X-rotation could be estimated from the apparent X-rotation of a number of profiles spaced down the head, even though some individual profiles might give a result that was not near the true X-rotation. We imposed a requirement that the method should work with data obtained from a single viewpoint. This means that a limit of operation is reached when an object is rotated such that symmetry is no longer visible from the single viewpoint.

To develop the basic principle into a practical implementation, we considered two basic approaches. One was to perform the entire pose estimation process in a single step on the whole 3D data set (albeit an incomplete representation of the object) and the other broke the process down into operations on 2D profiles taken through the data set, followed by a combined analysis of the results from the individual profiles. The second approach, presented here, reduces the dimensionality of each step, giving a corresponding reduction in computational complexity.

3 Estimation of X-rotation using a single profile

This section covers the basic principles of X-rotation estimation using a single profile parallel to the y_oz_o plane. We will then extend this to the use of multiple profiles, and the estimation of additional rotations about other axes.

When a profile is reflected, translated and rotated (called a reflected profile for conciseness), a measure of similarity can be made between the original and reflected profiles. Simplistically, the orientation yielding the greatest similarity between the original and reflected profiles is the estimated pose of the object, and the measure of similarity is a measure of symmetry.

3.1 Symmetric and geometric constraints

The assumption that there is a plane of symmetry can be used to restrict the number of poses considered, because some do not propose a symmetrical object. We call this the symmetric constraint. Figure 2 shows a profile (solid) and three different translated and rotated reflections of the profile (dashed). The example in Figure 2(a) violates the symmetric constraint. We are also not interested in an arbitrarily symmetrical object where the profiles are separated as in Figure 2(b). The reflected profile may be generated by reflecting, translating and rotating the original profile, or by reflecting the original profile about a line that is translated and rotated, but still intersects the original profile, as in Figure 2(c). The second option automatically satisfies the symmetric constraint and guarantees that the profile and its reflection are not separated.

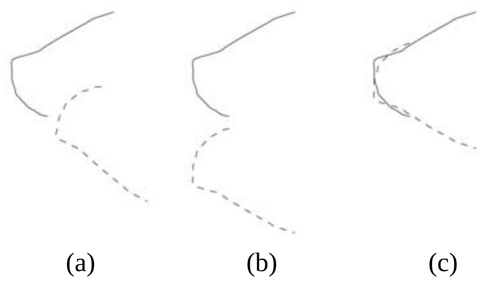

Figure 2: Original profile (solid) and its translated and rotated reflections (dotted) resulting in (a) asymmetrical, (b) symmetrical but separated and (c) symmetrical objects.

Knowledge of the sensor position, together with the assumption of symmetry, can be used to test the validity of the recovered object shape. This constraint is illustrated in Figure 3 where the original profile ABC, is reflected about line PQ giving the reflected profile A′B′C′. In this proposed pose, the section of the reflected profile between A′ and B′ implies that the sensor should also have collected data beyond point C towards point A′. Clearly, for this proposed pose, there is a region bounded by CA′PC that is geometrically inconsistent with the premise of symmetry. We call this the geometric constraint. In Figure 3, any part of the reflected profile that is in the region below line PC and any part of the original profile that is in the region above line PC′ violates the geometric constraint. The geometric constraint is used when assessing the measure of similarity between a profile and its reflection. It heavily penalises impossible solutions, and is a powerful constraint in forming the measure of similarity.

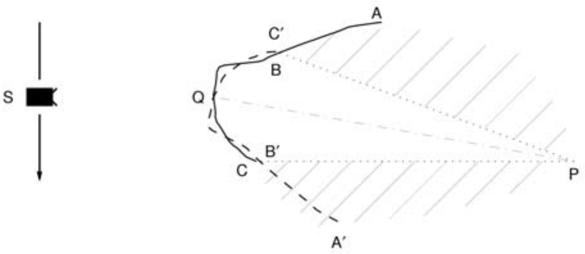

Figure 3: Proposed pose that violates the geometric constraint.

3.2 Measure of similarity

We use a simple measure of similarity to indicate how symmetrical an object would be in a given proposed pose, using the absolute area between the original and reflected profiles, and the length of the profile. In Figure 3, the measure of area between the corresponding portions of the profiles in the region between points B and C is simple, but the portions AB and A′B′ have no corresponding portion. Considering the portion A′B′, we do not know where the corresponding part of the original profile is because of self occlusion, but we do know that at least it does not extend below the line CP, so we take the area between A′B′ and the line CP. The measure of similarity Sim_{Prof}, for a single profile,

$$Sim_{\text{Prof}} = Area_{\text{Cor}} \frac{l_{\text{Prof}}}{l_{\text{Cor}}} \quad + \quad Area_{\text{GeomViol}} \tag{1}$$

where in Figure 3; $Area_{\text{Cor}}$ is the area between the profile and its reflection between points B and C, l_{Prof} is the entire length of the profile ABC, l_{Cor} is the length of the overlapping portion between points B and C, and $Area_{\text{GeomViol}}$ is the area between A′B′ and the geometric constraint line CP, and between AB and the geometric constraint line C′P. $Area_{\text{GeomViol}}$ is zero if the geometric constraint is not violated. Scaling the area between the corresponding portions by the ratio of the profile length to the overlapping length, balances the natural increase with increasing overlap length which would give a preference to small overlap.

This measure of similarity gives a value of zero for perfect symmetry, and increases as the symmetry degrades. In general we will not see perfect symmetry and so the orientation giving the lowest value of the measure of similarity is chosen as the orientation of best symmetry.

4 Estimation of X-rotation using multiple profiles

In practice we can use more than a single profile of an object to estimate the pose. Estimating the X-rotation of a number of individual profiles results in a range of values. There are several factors that contribute to this. Natural objects that are not perfectly symmetrical will give different values at different locations down the axis of the object, even if there is only one dominant plane of symmetry. The rotation estimated by each of a set of profiles spaced over the length of a suspended cow's head is typically grouped about several values corresponding to the multiple planes of symmetry that can be detected from the incomplete data that is available from a single viewpoint. To estimate the rotation of the object from the individual profile rotation estimates, we performed a series of relatively simple steps.

1. Estimate the X-rotation for each profile as in Section 3
2. Cluster the values of X-rotation and determine the cluster boundaries
3. For each cluster
   (a) within the bounds of the cluster obtain the best measure of symmetry for every profile
   (b) ignore abnormally poor measures of similarity from ends of object

(c) ignore measures of similarity that are grossly different to the others for this particular object

(d) calculate a weighted measure of similarity for the whole cluster

4. Choose the cluster with the best weighted measure of similarity

5. Estimate the rotation

The calculation of the measure of symmetry for each cluster, combining measurements from all profiles, was

$$Sim_{\text{Clus}} = \frac{\sum_{\text{profiles}} (Sim_{\text{Prof}} \cdot l_{\text{Prof}})}{\sum_{\text{profiles}} l_{\text{Prof}}} \quad (2)$$

where Sim_{Prof} is the similarity measure for a single profile as defined in equation 1; and l_{Prof} is the length of a single profile. This weights the similarity measure giving the greatest weighting to the longest profiles; those for which there is the most information. The cluster with the best measure of similarity is chosen and the X-rotation estimate is obtained from the median of all of the values taken from individual profiles.

5 Estimation of pose in the presence of multiple rotations

The method covered in earlier sections will not necessarily give valid results in the presence of Y-rotation. However, in the presence of a significant amount of Y-rotation, a reasonable estimate of X-rotation can still be obtained. Furthermore, from this first estimate, we are able to use an iterative approach to refine both the X and Y-rotation estimates. This avoids what would otherwise be a computationally intensive search for the pose in a higher dimensional space.

The method may be summarised in the following sequence of steps:

1. Choose a small number of widely spaced Y-rotations

2. Determine the X-rotation and measure of symmetry for each Y-rotation

3. Choose the Y-rotation and X-rotation giving the best measure of symmetry

4. Determine a correction to the Y-rotation estimate

5. Estimate the X-rotation at the corrected Y-rotation

6. Iterate over steps 4 to 5

The key to this process is the ability to determine a correction to the Y-rotation estimate using the X-rotation estimate.

A simple case involving two profiles of an object that has been X, Y and Z-rotated is illustrated in Figure 4 with the top profile offset behind the other due to the Z-rotation. Points p and q on profiles P and Q are the points where the proposed plane of symmetry intersects the profiles. Figure 4(a) and (c) show the two profiles in plan and elevation. This elevation suggests a misleading correction to the Y-rotation estimate because of the presence of both the X-rotation and the offset between the two profiles. Correcting for the estimated X-rotation as in Figure 4(b) and (d), the effect of the Z-rotation is removed, and the angle of the line pq gives an estimate δ of the Y-rotation error.

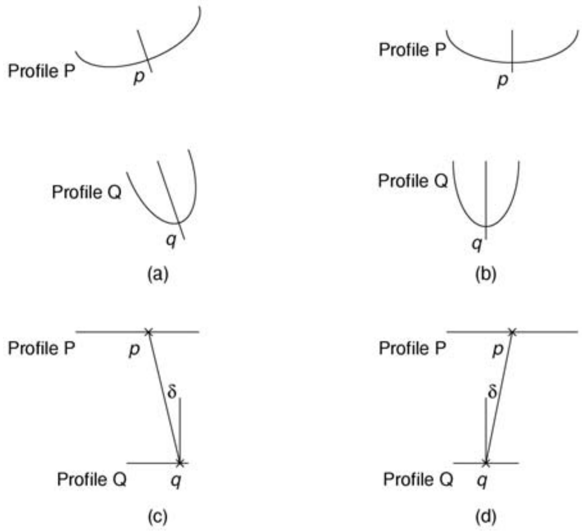

Figure 4: Correction of Y-rotation estimate. (a) plan view, (b) plan corrected for X-rotation, (c) elevation, (d) elevation corrected for X-rotation.

The third rotation, Z-rotation, can not be determined without attaching a meaning to the shape of the object. Although it can not be determined by symmetry alone, having estimated the pose and the position of the plane of symmetry we are able to create a profile in the plane of symmetry. This profile can then be used directly to determine the third of the three rotation angles.

6 Results

The method has been tested using both synthetic data sets [11] and data scanned from real cows. The quantitative results presented here are based on the synthetic data set because ground truth data could be accurately determined over a comprehensive range of orientations. When working with real cows' heads, the ground truth orientation was independently manually estimated.

Figure 5 shows errors in estimating the orientation of cows' heads that had been rotated to 343 different ori-

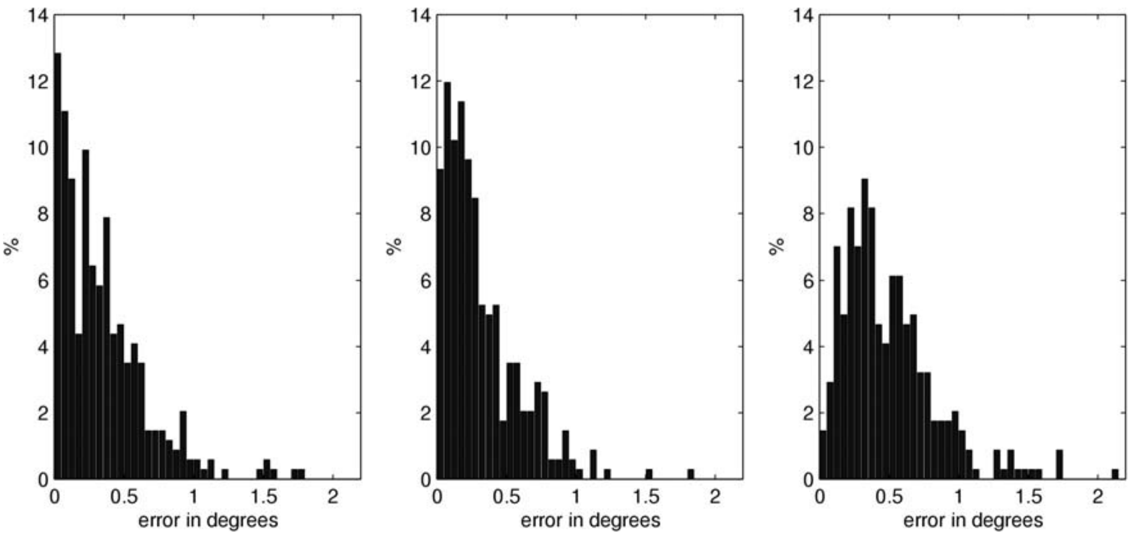

(a) (b) (c)

Figure 5: Distribution of errors in estimation of rotation of 343 cows' heads that have been simultaneously X, Y and Z-rotated in the range $\pm 30°$; (a) error in estimation of X-rotation, (b) error in estimation of Y-rotation, (c) error in estimation of direction of normal to plane of symmetry.

entations involving simultaneous X, Y and Z-rotations in the range of $\pm 30°$. The three histograms show the distribution of the errors in estimation of the individual components of rotation and the error in estimation of the direction of the normal to the plane of symmetry.

Finally, Figure 6 shows data scanned from a real cow's head complete with horns, ears and a flap of hide visible at the right rear of the head where the head was separated from the neck. Figure 6(a) shows the head in the orientation in which it was scanned, and Figure 6(b) shows a view with the head re-oriented according to the estimated X and Y-rotations. In Figure 6(b) the appearance to the left of the nose region is deceptive and it does not look symmetrical. This is because some of the surface was not visible from the viewpoint of the sensor. The geometric constraint enables a correct interpretation to be made in the presence of such missing data. A further example of missing data is evident near the base of the ear on the left and inside the ear on the right, as neither region was visible from the viewpoint of the sensor.

There are several fundamental limits on the application of the method, and how close an object must be to the *a priori* assumption of correct orientation. The object must possess a degree of bilateral symmetry as this is the only property that is exploited, and the solution must be unique within the orientation bounds considered, because no preference is given in the presence of multiple planes of symmetry.

A further limitation exists, particularly in the presence of Y-rotation. Profiles corresponding to an incorrect proposed Y-rotation may not possess sufficient symmetry to enable a reliable result to be achieved. This

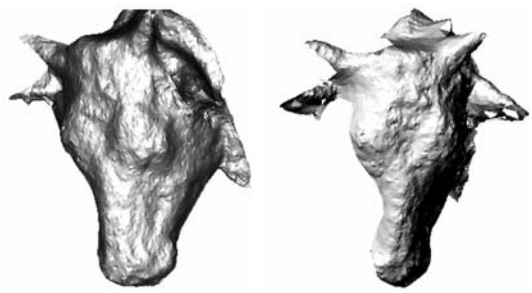

(a) (b)

Figure 6: View of a real cow's head, (a) oriented as scanned, (b) re-oriented by X-rotation of 39° and Y-rotation of 3°.

is particularly evident when nearby cross-sections perpendicular to the plane of symmetry are significantly different. Such a region exists in the eye region of a cow's head where the shape of the forehead above the eyes is markedly different from the shape of the nose region below the eyes. When multiple profiles are used to estimate the pose, this can be resolved.

7 Conclusions

This paper proposed a method for determining the pose of a symmetrical object from 3D surface data when no features or other knowledge of its shape was available. We have developed this method to the point where it works reliably on very irregular objects. It enables the orientation to be determined and expressed as rotations about the coordinate system axes. The underlying principle, estimating symmetry by finding correspondences in the reflection of a data set, is a problem of high dimensionality. We have reduced the problem to a series of steps of lower dimensionality. We have also explicitly incorporated a solution to the problem of incomplete data due to self occlusion.

8 Acknowledgements

This work was supported by the New Zealand Foundation for Research, Science and Technology.

The synthetic cow data used in some of our work was obtained from the POVLAB download site [11].

References

[1] Francois Gieres. About symmetries in physics. In F. Gieres, M. Kibler, C. Lucchesi, and O. Piguet, editors, *Symmetries in Physics, Proceedings of the fifth "Seminaire Rhodanien de Physique", Dolomieu 1997, (Editions Frontieres, 1998)*, 1998.

[2] E. Hjelmas and Boon Kee Low. Face detection: a survey. *Computer Vision and Image Understanding*, 83(3):236–274, September 2001.

[3] J. Zhou, X. G. Lu, D. Zhang, and C. Y. Wu. Orientation analysis for rotated human face detection. *Image and Vision Computing*, 20(4):239–246, April 2002.

[4] Hagit Zabrodsky, Shmuel Peleg, and David Avnir. Continuous symmetry measures IV: Chirality. *Journal of the American Chemical Society*, 117:462–473, 1995.

[5] Hagit Zabrodsky and Daphna Weinshall. Using bilateral symmetry to improve 3D reconstruction from image sequences. *Computer Vision and Image Understanding*, 67(1):48–57, 1997.

[6] K. Hattori, S. Matsumori, and Y. Sato. Estimating pose of human face based on symmetry plane using range and intensity images. In A. K. Jain, S. Venkatesh, and B. C. Lovell, editors, *Proceedings Fourteenth International Conference on Pattern Recognition*, pages 1183–1197. IEEE Computer Society, 1998.

[7] Changming Sun and Jamie Sherrah. 3D symmetry detection using the extended Gaussian image. *IEEE Transactions on Pattern Analysis and Machine Intelligence*, 19(2):164–168, February 1997.

[8] David O'Mara and Robyn Owens. Measuring bilateral symmetry in digital images. In *Proceedings of TENCON '96 - IEEE Region Ten Conference: Digital Signal Processing Applications, Perth, Western Australia*, pages 151–156, November 1996.

[9] P. J. Besl and N. D. McKay. A method for registration of 3-D shapes. *IEEE Trans. on Pattern Analysis and Machine Intelligence*, 14(2):239–256, February 1992.

[10] Andrew E. Johnson and Martial Hebert. Surface matching for object recognition in complex three-dimensional scenes. *Image and Vision Computing*, 16(9-10):635–651, July 1998.

[11] Denis Olivier. Cow.zip, raw file containing a cow data mesh. http://pdelagrange.free.fr/povlab/store.html, visited on 18/8/2003.