Inspection of Free-form Surfaces Using Dense Range Data

Edvaldo M Bispo and Robert B Fisher

DAI Research Paper No. 700

June 1994

Submitted to the European Symposium on Optics for Productivity and Manufacturing Conference: Automated Vision Inspection, Frankfurt, June 94

Inspection of Free-form Surfaces Using Dense Range Data

Edvaldo M. Bispo and Robert B. Fisher
Dept. of Artificial Intelligence,
University of Edinburgh,
5 Forrest Hill, Edinburgh EH1 2QL

ABSTRACT

This paper presents some research on the use of dense range data for the automatic inspection of mechanical parts that have free-form surfaces. Given a part to be inspected and a corresponding model of the part, the first step towards inspecting the part is the acquisition of a range image of it. In order to be able to compare the part image and its stored model, it is necessary to align the model with the range image of the part. This process, called registration, finds the rigid transformation that superposes model and data. After the registration, the actual inspection uses the range image to verify if all the features predicted in the model are present and within tolerance. Free-form surfaces are particularly interesting in that few inspection processes can inspect surface shape across the whole surface.

We focuses in the inspection of free-form surfaces and present some results concerning the extraction of nominal shape models from dense range data using B-splines and the use of B-spline models of free-form surfaces for the purposes of registration.

1 INTRODUCTION

An essential contribution to the success and growth of manufacturing industries is the control of the quality of their products. This situation creates a great demand for cheap and reliable automatic inspection systems able to replace human inspection, who are not ideal for many reasons:

- The performance of humans in inspection is subject to many fluctuations due to human fatigue and tedium.
- The necessity of performing inspection in hazardous environments in many circumstances.
- Some situations, as in the case of hot steel slabs [12], in which the high throughput and accuracy required make use of human inspectors completely impossible.

We are investigating the use of dense range data for the inspection of mechanical parts. The main idea behind our research is the comparison between the three dimensional geometric information present in the range image and descriptions or models of ideal parts stored previously by the system.

Given a part to be inspected and a corresponding model of the part stored in the *model data base*, the first step of inspecting the part is the *acquisition of data* corresponding to the part which in our case means the acquisition of a range image of it. In order to be able to compare the part image and its stored model, it is necessary to align the model with the range image of the part. This process, called *registration*, finds the rigid transformation that superposes model and image. After the image and model are registered, the actual *inspection* uses the range image to verify if all the features predicted in the model are present and have the right pose and dimensions. As the manufacturing and measurement processes are never error free, the feature's pose and dimensions should be within specific tolerances, which means that the inspection models must convey explicit information about the acceptable variances of the part around the nominal value.

In many cases not all the features to be inspected are visible in one view of the object, and it is necessary

to acquire other images of the part being inspected. Therefore, an important element in the inspection process is the *planner*, which takes account of the current position of the part (known after the first registration), the model of the shape of the part, the model of the sensors, and the list of features to be inspected to guide the acquisition of images corresponding to different views of the object.

Most of the time the *model data base* will be built by using information about the shape of the ideal parts and the tolerances associated with these parts supplied by the part manufacturer. However, there may be situations in which this information is not available and one has just physical examples of good and bad parts. In these cases, a very interesting characteristic of an ideal system would be its ability to *learn models* of the part to be inspected just from the images of the good and bad parts.

From all the considerations above, we can envisage a diagram of a general inspection system, given in Figure 1, composed by the following modules: *data acquisition, model data base, registration, inspection, planer* and *model learning*. In this paper we start by discussing the acquisition of range images in Section 3. Then we present some results concerning the extraction of B-spline models from range images in Section 4. The problem of registering free-form surfaces is discussed in Section 5. Our finals comments and conclusions are presented in Section 6.

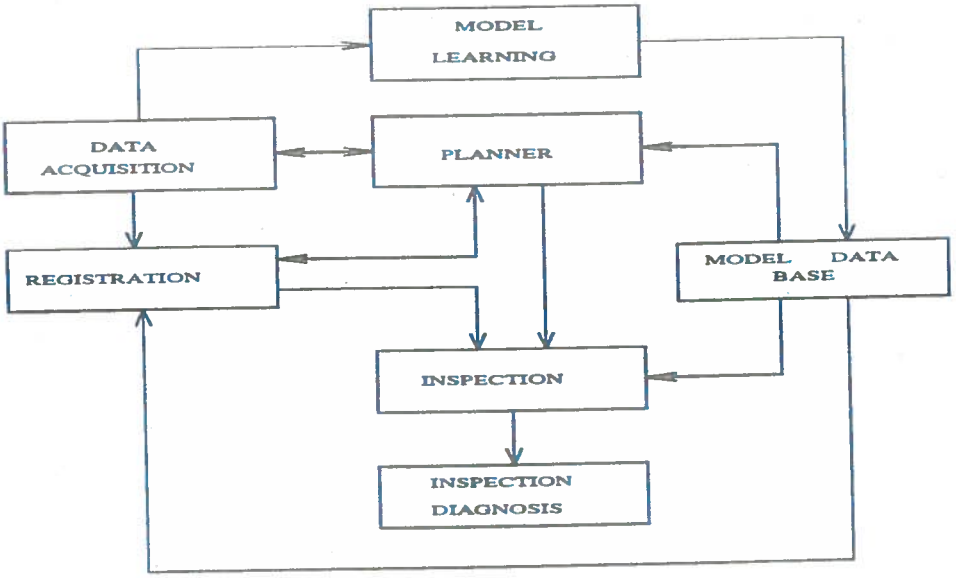

Figure 1: General system diagram

2 DATA ACQUISITION

One of the most common ways of acquiring 3D information for inspection involves tactile sensing, using co-ordinate measuring machines (CMM) (see [7] for an example of an application). CMMs machines are able to produce very accurate and precise results but, despite their popularity CMMs have some important drawbacks related to their slow rate of data capture and the difficulties of programming them for inspection tasks.

Among the non-contact techniques for producing range images the most important ones are (see [1] and [8] for extensive reviews):

- *Radars*: based in the emission and reception of electromagnetic waves, the most common types are the time of flight, amplitude modulation and frequency modulation.

- *Shape from shading*: surface orientation is recovered using a model of surface reflectance, knowledge of light source position and making the assumption of surface smoothness.
- *Photometric stereo*: uses knowledge about surface reflectance and multiple light source positions to recover surface orientation.
- *Passive triangulation*: pairs or more 2-D images obtained from different viewpoints are combined to determine the depth of surrounding surfaces. The depth calculation is based on finding the correspondences between the 2-D images' features and then using these correspondences and geometric principles (triangulation).
- *Active triangulation*: emission of some form of structured light (*e.g.* stripes, coded binary parts) and use of geometric principles for the calculation of depth. Depth is recovered using a camera model and knowledge about the geometry of the structured light.
- *Moirè techniques*: range measurements are obtained through the modulation and demodulation of light using special gratings, the working principle being based in the interference phenomena.

The radar techniques, in general, are very expensive and are applied to measuring big depth values. Moirè techniques are extremely accurate, but are also very slow and involve quite a lot of processing in calculating the depth information. Passive triangulation techniques suffer the burden of establishing the correspondence between the 2-D images, which is a very complicated and ill-posed problem. Shape from shading and photometric stereo do not recover depth information explicitly and depend on a large number of assumptions about the environment. Considering these facts and the difficulties involved in using CMMs, we decided to use a range finder based in the active triangulation principle. More specifically, we propose to investigate the viability of using a system like the Edinburgh's laser striper [2], which consists of a pair of cameras and a plane of laser light used to project stripes on the object.

3 EXTRACTING B-SPLINES FROM RANGE DATA

We extracted uniform B-splines from range data using the following procedure:

Smoothing: Application of conservative and average smoothing to the range image in order to eliminate outliers.

Thresholding: Thresholding of the original image for separating background and foreground and delimiting holes. Holes were delimited by thresholding the range image only in a window containing the hole. The thresholding was used for producing a bitmap indicating the interesting points in the range image.

Fitting: Fitting of a uniform B-spline to the interesting points indicated by the bitmap produced by the thresholding. The fitting was done using a algorithm proposed by DeeBoor [3] that minimises the average square distance between the spline model and the image points.

Results of the spline fitting are shown in Figures (4) and (5) that show spline models of the test objects **1** and **2** shown in Figures (2) and (3). The average distance between model points and data points in both cases were in the order of tenths of millimiters (withing image noise standard deviation equal to 15 mm). As one would expect the bigger errors happened near the discontinuities of the image. This fact explains the better results obtained with the test object number **2** which presents less discontinuities than the test object number **1**.

4 REGISTERING FREE-FORM SURFACES

Given a range image of a free-form surface, which describes the shape of this surface in the sensor coordinate system, and a model of the same surface in a different coordinate system, the problem of registering model and data consists in finding the rigid transformation (rotation and translation) that aligns or superposes both.

Some of the first research into matching free-form surfaces was done by Faugeras and Hebert [6]. They presented a matching algorithm able to successfully match a Renault auto part, but their algorithm had the drawback of depending on the existence of planar regions in the object being matched. Potmesil [13] developed a system for

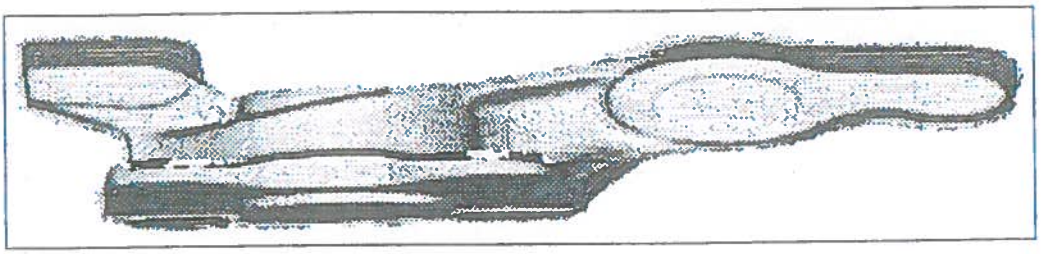

Figure 2: Range image of Test object number 1

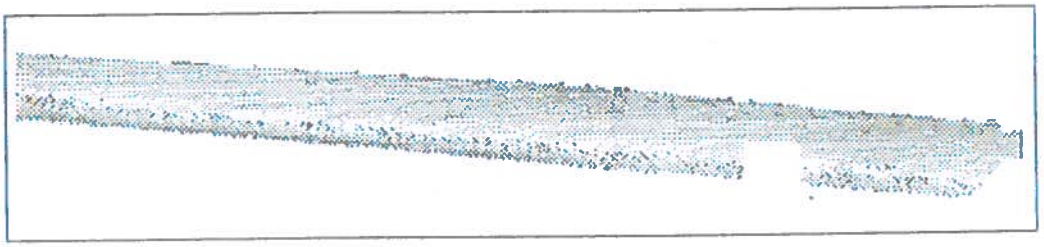

Figure 3: Range image of Test object number 2

modelling the complete surface of a object from different range images. The different range images were matched through heuristic search in the 6-D space of the translation and rotation. Chen and Medioni [11] also developed an algorithm for registering overlapping range images of an object. They assume an *a priori* knowledge of an estimated pose aligning the different range images and approximate one of the surfaces being matched using its tangent planes. Besl and McKay proposed the ICP (Iterative Closest algorithm) [10] for registering free-form surfaces using only the 3-D points on the surface of the object. Zhang [15] proposed an algorithm very similar to the ICP in which an *a priori* approximation of the registration is assumed to be known and is used to accelerate the convergence of the algorithm and make it more robust to outliers. Recently Stein [14] advocated the use geometric hashing together with an IT (Interpretation Tree) [9] for recognizing and aligning objects in range and gray level images.

We assume that the object being registered does not present any salient features (*e.g.* planar regions) that could help guiding the registration process and concentrate solely on the use of the 3-D points for obtaining the registration. We also assume the *a priori* knowledge of an approximation of the pose aligning model and data and investigate the use of algorithms as the ICP for improving the accuracy of the registration.

The final accuracy that can be obtained using only 3-D points on the data and model surface depends critically on the calculation of the pose aligning sets of correponding 3-D points on the surfaces of model and data. Therefore we start by evaluating the accuracies that can be achieved using sets of corresponding 3-D points in Section **4.1**. In Section **4.2** we discuss the ICP algorithm in detail and present some extensions for improving the robustness of the algorithm.

4.1 Registering two sets of corresponding 3-D points

Consider a set of data points $\mathbf{D} = \{d_i\}_{i=1,N}$ and a set of corresponding model points $\mathbf{M} = \{m_i\}_{i=1,N}$ such that $\mathbf{D}$ and $\mathbf{M}$ are related by a rigid transformation corresponding to a translation $\mathbf{T}$ and a rotation $\mathbf{R}$ as defined in the equation (1).

$$d_i = \mathbf{R} \cdot m_i + \mathbf{T} + w_i \quad (1)$$

$$\mathbf{E}[w_i] = [0 \;\; 0 \;\; 0]^t \quad (2)$$

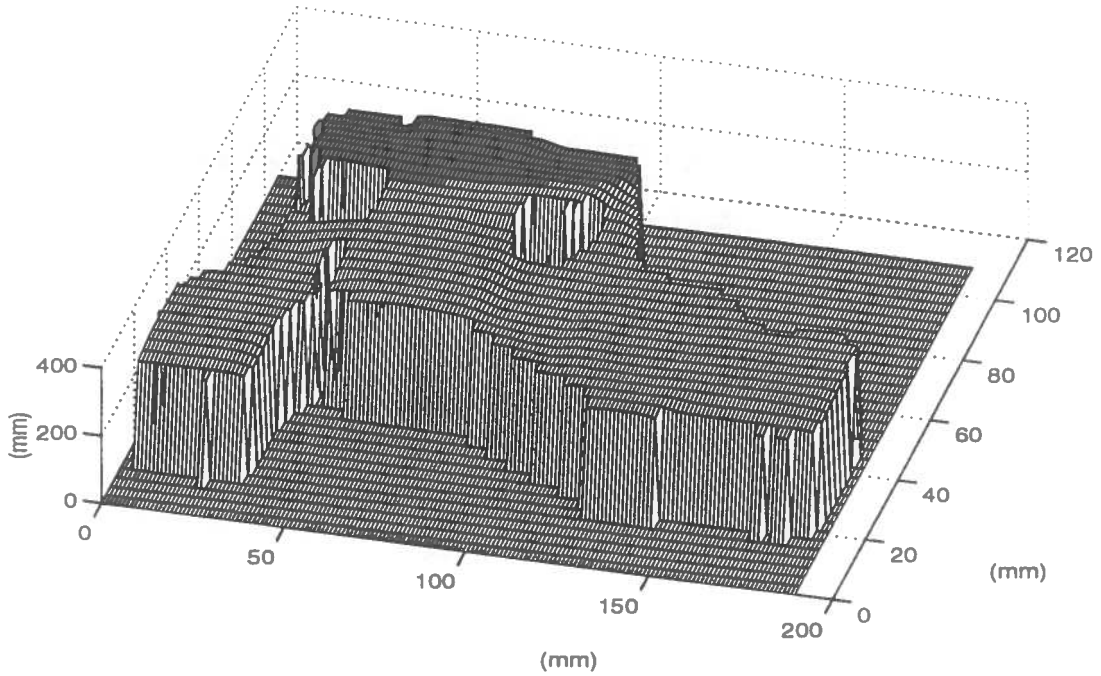

Figure 4: Uniform B-spline model of Test object number 1. Average distance between model and data equal to 0.4 mm. Uniform grid of 50X50 knots.

$$\mathbf{E}[w_i \cdot w_i{}^t] = \begin{bmatrix} {\delta_i}^2 & 0 & 0 \\ 0 & {\delta_i}^2 & 0 \\ 0 & 0 & {\delta_i}^2 \end{bmatrix} \tag{3}$$

where : w_i is a white noise modelling the sensor errors and $\mathbf{E}[\cdot]$ is the expectation operator

The problem of finding $\mathbf{T}$ and $\mathbf{R}$ that minimize the mean square distance E_M between $\mathbf{D}$ and $\mathbf{M}$ (equation (4)) is a well known problem to which there are several proposed solutions [6,4,5]. Under the the criterion of accuracy in finding the right values of $\mathbf{R}$ and $\mathbf{T}$ all these methods are equivalent, and we will concentrate our attention on Horn's algorithm [5].

$$E_M = {D_M}^2 = \sum_i \sum_{i=1,N} \|d_i - \mathbf{R} \cdot m_i - \mathbf{T}\|^2 \tag{4}$$

In order to investigate the degradation in the accuracy of the algorithm with increasing noise δ_i we extracted 3-D points from the spline models of test objects **1** and **2** and use them as model points. The data points were generated by applying a known rotation and translation to the model points and corrupting the result with noise. The experiment was run for different values of rotation and translation and different levels of noise δ_i. To each value of noise level 200 samples were considered.

As one could expect the results did not depend on the specific values of rotation and translation. Also, the final value of the mean square distance (E_M) found by Horn's algorithm tended assymptotically to $\mathbf{E}[{w_i}^t \cdot w_i] = 3 \cdot \delta^2$ as the number of matched points increase. Figures (6) and (7) show the variation of D_M (eq. 4) and the error of the translation estimate ($\hat{\mathbf{T}}$) for the test object number **1**. Figures (8) and (9) show the same for the test object number **2**. These graphs show the maximum accuracy that can be achieved for a given level of noise, which corresponds to the combined effects of sensor and modelling errors, and number of matched points. A careful look in the graphs shows that the accuracy in the estimation of translation will not change much with the shape of the surface being registered. The graphs also show that the error in the estimate of translation will decrease almost linearly with the number of matched points, as shown in equation (5). It is important to remember that the graphs show the very best accuracy that can be achieved, because other sources of errors as sampling errors and mismatchings were not

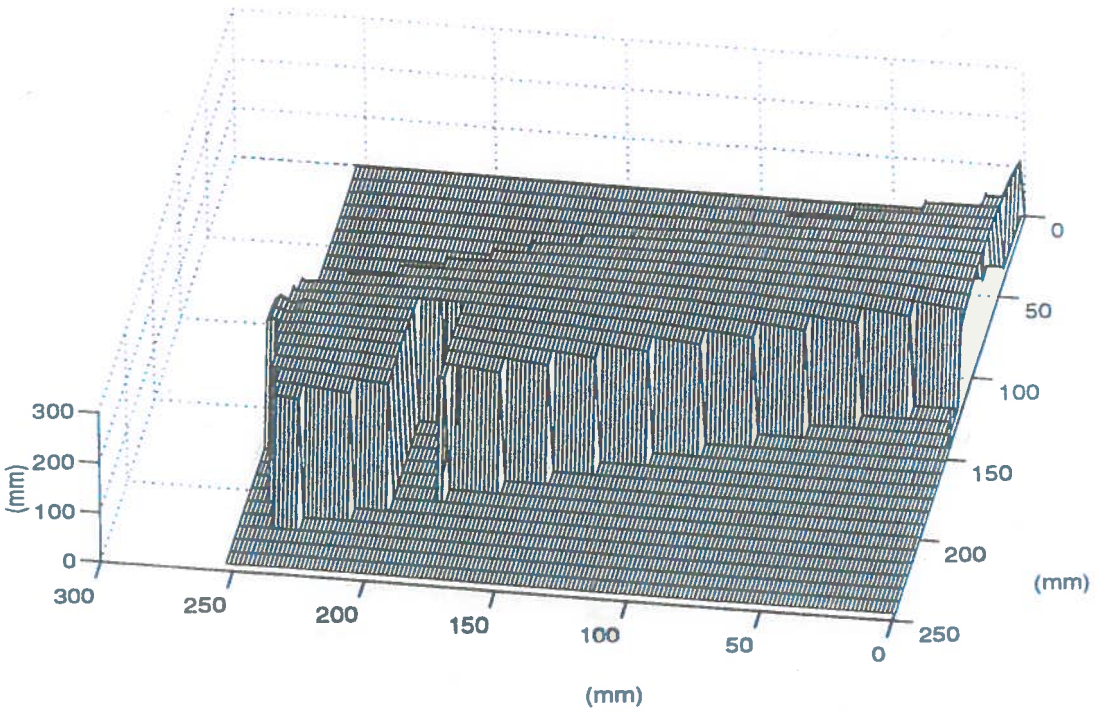

Figure 5: Uniform B-spline model of Test object number 2. Average distance between model and data equal to 0.08 mm. Uniform grid of 50X50 knots.

taken into account.

$$\|\hat{\mathbf{T}} - \mathbf{T}\| \propto \delta_i / N \tag{5}$$

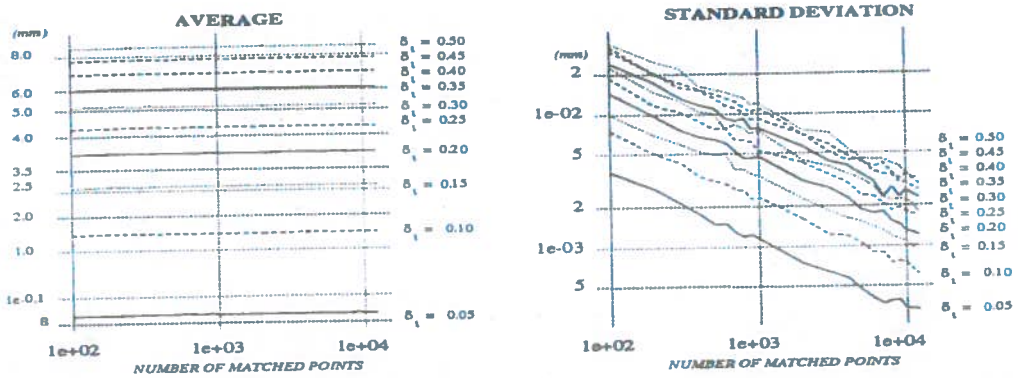

Figure 6: Variation of the mean distance (D_M) with the noise level δ_i for test object number 1. Average and standard deviations calculated using 200 samples (10 different curves for δ_i varying from 0.05 mm to 0.5 mm).

4.2 The ICP algorithm

Given a set of data points $\mathbf{D} = \{d_i\}_{i=1,N}$, a model shape $\mathbf{M}$ and an initial estimate $(\hat{\mathbf{R}}, \hat{\mathbf{T}})$ of the rigid transformation aligning $\mathbf{D}$ and $\mathbf{M}$, the ICP algorithm [10] is:

1. Apply the rotation $\hat{\mathbf{R}}$ and translation $\hat{\mathbf{T}}$ to the points in $\mathbf{D}$ obtaining the set of points $\mathbf{D^k} = \{d^k{}_i\}_{i=1,N}$
2. To each point $d^k{}_i$ find the point $m^k{}_i$ in $\mathbf{M}$ which is closest to $d^k{}_i$.

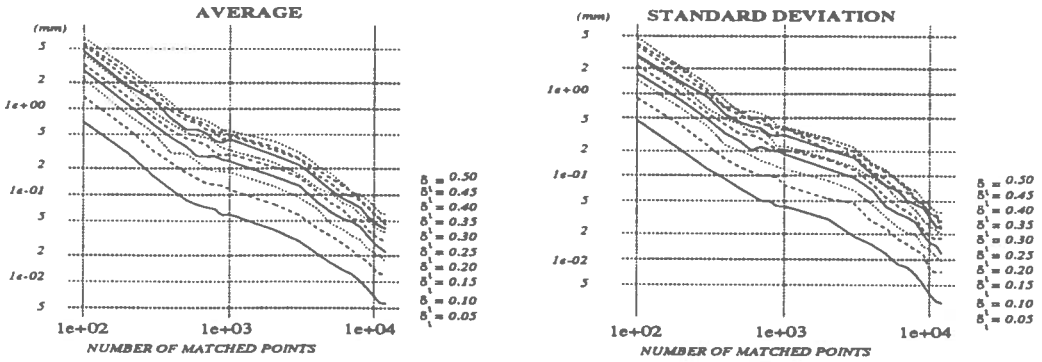

Figure 7: Variation of the translation error ($||T - \hat{T}||$) with the noise level δ_i for test object number 1. Average and standard deviations calculated using 200 samples (10 different curves for δ_i varying from 0.05 mm to 0.5 mm).

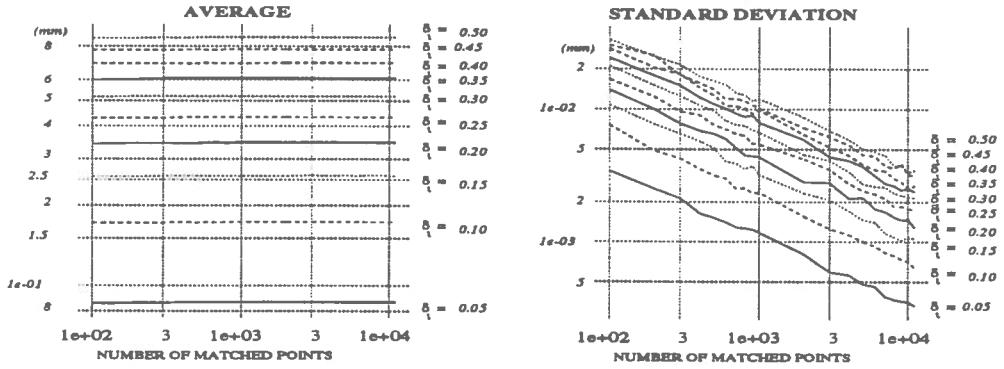

Figure 8: As Figure (5) but for test object number 2

3. Calculate new estimates of $\hat{\mathbf{T}}$ and $\hat{\mathbf{R}}$ by assuming the points d_i and $m^k{}_i$ correspondent, and applying Horn's algorithm (or equivalent).

4. Repeat 1 to 3 till convergence (E_M stops decreasing).

In our experiments the model shape M was extracted from a B-spline model of the surface being registered. The models were obtained as described in Section 3. Because the evaluation online of the points on the spline would be too computationally expensive, we used as the model shape M a dense grid of points extracted offline from the spline model.

A crucial step in the algorithm consists of finding the closest point to the point $d^k{}_i$ on the model shape M. Zhang [15] uses K-D trees. In our implementation we used a multiscale search for finding the closest point. This approach can occasionally fail when the region containing the real closest point is not detected during the coarse scale search. In our experiments this never happened because we start the search at a scale small enough to avoid this problem. Also, due to the consistency checks we added to the ICP algorithm (see below), occasional failures to find the true closest point do not degrade the performance of the algorithm. Figures 7 and 9 show that an accuracy of the order of a tenth of millimeter (*i.e.* approaching the noise variance) can be achieved using only 1000 points. Therefore, we could subsample the data points (range image) which reduced considerably the number of points to be considered. The sampling of the model and data shape together with the use of the multiscale shape allowed us to achieve a good speed without the use of K-D trees. As suggested by Figures 7 and 9, better accuracies can be achieved by running an extra iteration of the ICP algorithm with more data points, after its convergence using coarser grids of data points.

In our experiments with the ICP algorithm using the object tests, we verified most of the properties described by Besl [10]. Clearly, the biggest problem with the ICP is the problem of converging to a local minimum of E_M that

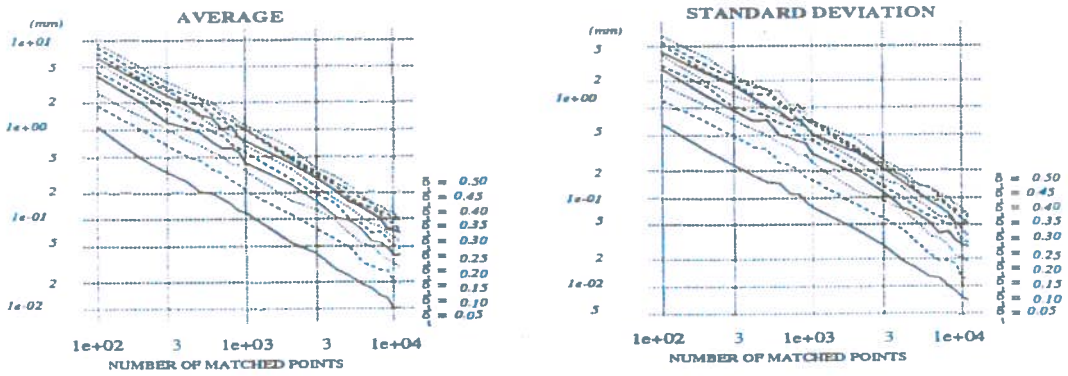

Figure 9: As Figure (6) but for test object number 2

does not correspond to the right registration between model and data. This problem may not be very significant with objects that have a large radius of convergence for the global minimum (*e.g* object test 1 to which the ICP will always converge to the right answer as long as the model is rotated less than 100 degrees around any axis), because is only necessary to run the ICP a few times using different initial estimates of $\mathbf{R}$ and $\mathbf{T}$. But this fact is a very serious problem in the case of objects like the object test 2 that has a very small radius of convergence due to the existence of other local minima near the desired global minimum.

A interesting aspect of the problem of the initial estimate is the influence of the original estimate of translation. As said in [10], when the data set is a subset of the model which covers a reasonable portion of the model, the ICP performance is reasonably insensitive to the original estimate of translation. However the final accuracy of the the registration found by the algorithm can change considerably as illustrated in Figure 10 (a). The situation becomes even worse when the data covers only a small part of the model as in Figure 10 (b).

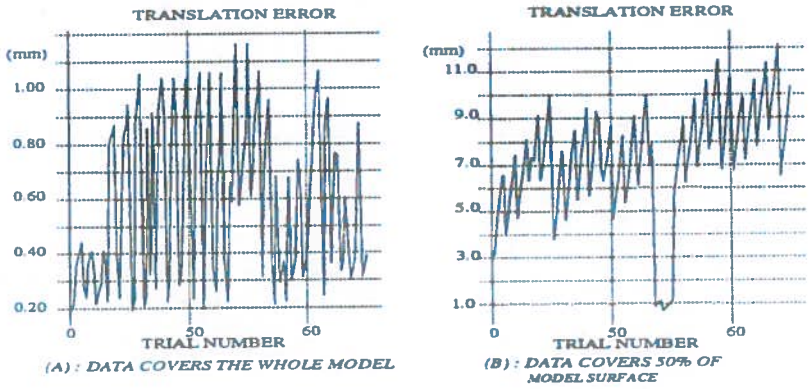

Figure 10: Variation of translation error ($||\hat{\mathbf{T}} - \mathbf{T}||$) of ICP with the value of the initial estimate of translation - Test object 1. Initial estimates of translation vary between -2 and 2 times the dimensions of the object. Part (a) covers the complete model and part (b) uses data from a subset of the model.

Another problem with the ICP is the fact that it can not cope with situations in which the data points are not a subset of the model points. This is a considerable limitation in applications such as model acquisition and inspection in which this assumption does not hold. In order to make the method more robust to outliers we altered the original algorithm and eliminated the pairs $(d_i, m^k{}_i)$ that did not satisfy a few consistency tests. The first of these tests, essentially used by Zhang [15] though in a different way, eliminates pairs of points $(d_i, m^k{}_i)$ if the distance between d_i and $m^k{}_i$ is bigger than three times the value of E_M found so far. We also used two local geometric constraints for rejecting wrong pairs: the distance between points in the grid and the angles between between the vectors joining pairs of points in the grid. Because of the surface sampling, a matching for a data point d_i is any model point m_i such that after the alignment between model and data m_i belongs to the portion of the data surface $\mathbf{D}$ that d_i represents. We approximated this region associated to d_i by a cuboid centered in d_i with the faces parallel to the

planes **xy**, **xz** and **yz** of the sensor coordinate system and dimensions corresponding to the size of the sampling grid.

The maximum values of distances and angles associated with the points d_i were calculated considering the angles and distances between the cuboids associated with each of the data points. Figure 11 illustrates the calculation for the 2-D case when the cuboids are rectangles. All the calculation of minima and maxima was done offline. In execution, whenever the relative distance between pairs of data points or relative angles between triples of data points was inconsistent with their corresponding model points, the correponding pair of data points and model points were rejected.

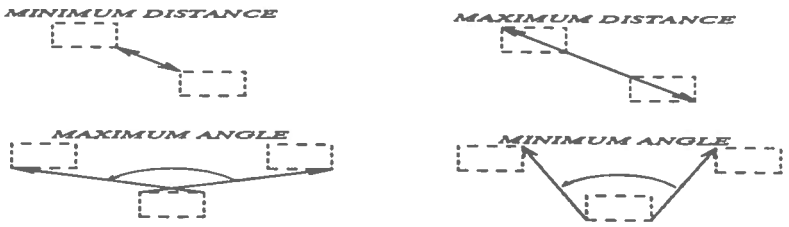

Figure 11: Local geometric constraints

These alterations make the method more robust to outliers but impose a couple drawbacks. All the tests proposed can only be used when the algorithm starts from a reasonable estimate of **R** and **T**. Besides that, the consistency tests make the algorithm lose its monotonicity in the presence of outliers, *i.e.* the value of E_M may increase sometimes. This happens because the tests may fail to eliminate all the outliers in some iterations. However, in our experiments with the modified ICP, the algorithm always found the right registration.

5 CONCLUSIONS

We discussed the problem of inspecting free-form surfaces and presented some results concerning the acquisition of B-spline models of free-form surfaces, as well as the use of such models for the purposes of registration. Despite some initial promising results there are still many issues to be further investigated: evaluating the robustness to outliers of the modified version of the ICP, investigating further on the monotonicity of the modified ICP algorithm, studying the use of techniques as the Hough Transform and search to obtain initial estimates of registration and, most of all, producing the final inspection diagnosis after the registation between model and data.

6 ACKNOWLEDGEMENTS

This work was supported by ACME under Grant GR/H/86905. We would like also to acknowledge the financial support of the CNPq in Brazil to Edvaldo Marques Bispo and to thank Andrew Fitzgibbon for his many suggestions and constructive criticising.

7 REFERENCES

[1] P. J. Besl: *Active optical range imaging sensors*, In J. L. C. Sanz, editor, *Advances in Machine Vision*. Springer-Verlag, New York, 1988.

[2] D. K. Naidu, R. B. Fisher, G. Cameron: *User guide for the laser range image acquisition package*, University of Edinburgh, Dept. of AI, Imagine software paper no. 18, 1990.

[3] C. DeBoor: *A practical guide to splines*, New York:Springer-Verlag, 1978.

[4] K. S. Arun, T. S. Huang and S. D. Blostein: *Least-squares fitting of two 3-D point sets*, IEEE Trans. Patt. Anal. Mach. Intell., vol. PAMI-9, no. 3, Sept 1987.

[5] B.K. Horn and J.G. Harris: *Closed-form solution of absolute orientation using unit quaternions*, J. Optical Society of America, vol. 4, no. 4, 1987, pp. 629 – 642.

[6] O.D. Faugeras and M. Hebert: *The representation, recognition, and locating of 3-D objects*, Int. J. Robotics Research, vol. 5, no. 3, 1986, pp. 27 – 52.

[7] C. Menq, H. Yau and G. Lai: *Automated precision measurement of surface profile in CAD-Directed Inspection*, IEEE Trans. Robotics and Automation, vol. 8, no. 2, April 1992.

[8] D. Poussart and D. Laurendau, *3-D Sensing for Industrial Computer Vision*, In J. L. C. Sanz, editor, *Advances in Machine Vision*. Springer-Verlag, New York, 1988.

[9] W. Grimson and T. Lozano-Perez: *Model-Based Recognition from Sparse Range and Tactile Data*, Int. Journal of Robotics Research, vol. 3, 1984, pp. 3 – 35.

[10] P. J. Besl and N. McKay: *A Method for Registration of 3-D Shapes*, IEEE Trans. Patt. Anal. Mach. Intell., vol. PAMI-14, no. 2, Feb 1992, pp.239 – 256.

[11] Y. Chen and G. Medioni: *Object modelling by registration of multiple range images*, Image and Vision Computing, vol. 10, no. 3 , April 1992.

[12] B. R. Suresh, R. A. F. Koushi, T. S. Levetti and J. E. Overlend: *A real-time automated inspection for hot steel slabs*, IEEE Trans. Patt. Anal. Mach. Intell., vol. PAMI-5, no. 6, Nov 1983.

[13] M. Potmesil: *Generating models for solid objects by matching 3D surface segments*, Proc. Int. Joint Conf. on Artif. Intell., Karlsruhe, Germany, August 1983, pp 1089–1093.

[14] F. J. Stein: *Structural indexing for object recognition*, PhD thesis, IRIS, University of Southern California, April 1992.

[15] Z. Zhang: *Iterative point matching for registration of free-form curves*, Research Report number 1658, INRIA, April 1992.